



Defense Science Board



SUMMER STUDY ON
AUTONOMY



JUNE 2016



Report of the
Defense Science Board
Summer Study on

Autonomy

June 2016

Office of the Under Secretary of Defense
for Acquisition, Technology and Logistics
Washington, D.C. 20301-3140

This report is a product of the Defense Science Board (DSB).

The DSB is a Federal Advisory Committee established to provide independent advice to the Secretary of Defense. Statements, opinions, conclusions, and recommendations in this report do not necessarily represent the official position of the Department of Defense (DoD). The DSB Summer Study on Autonomy completed its information gathering in August 2015. The report was cleared for open publication by the DoD Office of Security Review on June 1, 2016

This report is unclassified and cleared for public release.



**DEFENSE SCIENCE
BOARD**

**OFFICE OF THE SECRETARY OF DEFENSE
3140 DEFENSE PENTAGON
WASHINGTON, DC 20301-3140**

June 10, 2016

MEMORANDUM FOR UNDER SECRETARY OF DEFENSE FOR ACQUISITION,
TECHNOLOGY & LOGISTICS

SUBJECT: Final Report of the Defense Science Board (DSB) Summer Study on Autonomy

I am pleased to forward the final report of the DSB Summer Study on Autonomy. This report offers important recommendations to identify the science, engineering, and policy problems that must be solved to permit greater operational use of autonomy across all warfighting domains.

The study focused on three areas: institutional and enterprise strategies to widen the use of autonomy; approaches to strengthening the operational pull for autonomous systems; and an approach accelerate the advancement of the technology for autonomy applications and capabilities. The study concluded that action is needed in all three areas to build trust and enable the most effective use of autonomy for the defense of the nation.

This report provides focused recommendations to improve the future adoption and use of autonomous systems. Recommendations also include 10 example projects intended to demonstrate the range of benefits of autonomy for the warfighter. The study also provides thoughts on how to expand the available technology for the use of autonomy for defense through several innovative technology stretch problem challenges.

I fully endorse all of the recommendations contained in this report and urge their careful consideration and soonest adoption.

A handwritten signature in black ink, appearing to read "Craig Fields".

Craig Fields
Chairman



**DEFENSE SCIENCE
BOARD**

**OFFICE OF THE SECRETARY OF DEFENSE
3140 DEFENSE PENTAGON
WASHINGTON, DC 20301-3140**

June 9, 2016

MEMORANDUM FOR CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: Final Report of the Defense Science Board (DSB) Summer Study on Autonomy

The final report of the Defense Science Board 2014 Summer Study on Autonomy is attached. In accordance with its terms of reference, the study reviewed the applicability of autonomy across a broad array of DoD missions and concluded that there are both substantial operational benefits and potential perils associated with its use.

The study was informed by briefings describing a sampling of related DoD programs spanning the spectrum from deployed capabilities to research investments; relevant efforts in the commercial sector; and international activities. While evident that the DoD is moving forward in the employment of autonomous functionality, it is equally evident that the pull from diverse global markets is accelerating the underlying tech base and delivering high-value capabilities at a much more rapid pace.

The study provides recommendations aligned with three over-arching vectors:

- Accelerating DoD's adoption of autonomous capabilities
- Strengthening the operational pull for autonomy
- Expanding the envelope of technologies available for use on DoD missions

The first vector focuses on enterprise-wide recommendations that target barriers to increased operational use of autonomy. In providing recommendations the study focused on issues including the need to build trust in autonomous systems while also improving the trustworthiness of autonomous capabilities, and identified a number of enablers to align RDT&E processes to more rapidly deliver autonomous capabilities to DoD missions. The study concluded that action on this set of interdependent enterprise-wide recommendations is of far greater importance—and urgency—than the implementation of any single program of record.

The study observed that autonomy can deliver value by mitigating operational challenges including:



**DEFENSE SCIENCE
BOARD**

**OFFICE OF THE SECRETARY OF DEFENSE
3140 DEFENSE PENTAGON
WASHINGTON, DC 20301-3140**

- Rapid decision-making
- High heterogeneity and/or volume of data
- Intermittent communications
- High complexity of coordinated action
- Danger of mission
- High persistence and endurance

Given the current budget environment, the study opted not to recommend major new programs. Instead, to strengthen the operational pull for autonomy, the study recommends a set of experiments/prototypes that would demonstrate clear operational value across these operational challenges. The recommended projects are intended also to serve as pilots to help refine and institutionalize the enterprise-wide recommendations.

Finally, given that commercial market drivers are spawning rapid advances in the underlying tech base, the report recommends that DoD take steps to engage non-traditional R&D communities in novel ways to both speed DoD's access to emerging research results and identify areas in which additional DoD investment is needed to fully address DoD missions.

While difficult to quantify, the study concluded that autonomy—fueled by advances in artificial intelligence—has attained a 'tipping point' in value. Autonomous capabilities are increasingly ubiquitous and are readily available to allies and adversaries alike. The study therefore concluded that DoD must take immediate action to accelerate its exploitation of autonomy while also preparing to counter autonomy employed by adversaries.

A handwritten signature in black ink that reads "Ruth David".

Dr. Ruth David
Co-Chairman

A handwritten signature in black ink that reads "Paul Nielsen".

Dr. Paul Nielsen
Co-Chairman

TABLE OF CONTENTS

Executive Summary 1

1 Introduction 4

 About autonomy 4

 Autonomy accelerates enterprise performance 6

 Military value and current DoD uses 11

 Study approach 13

2 Trustworthiness and Trust in Autonomous Systems 14

 Key issues and barriers to trust 14

 An integrated approach is needed 21

3 Accelerating the Adoption of Autonomous Capabilities 24

 Tackling engineering, design, and acquisition challenges 24

 Mitigating cyber issues 27

 Creating new test, evaluation, modeling, and simulation paradigms 29

 Integrating technology insertion, doctrine, and concepts of operation 35

 Developing an autonomy-literate workforce 36

 Improving technology discovery 38

 Improving DoD governance for autonomous systems 40

 Countering adversary use of autonomy 42

4 Strengthening Operational Pull for Autonomy 45

 Autonomy for battlespace awareness 46

 Autonomy for protection 53

 Autonomy for force application 60

 Autonomy for logistics 69

5 Expanding the Envelope 76

 Generating future loop options 79

 Enabling autonomous swarms 83

 Intrusion detection on the Internet of Things 87

 Building autonomous cyber-resilient military vehicle systems 91

 Planning autonomous air operations 94

Summary 98

Defense Science Board Summer Study on Autonomy

Terms of Reference 102

Members of the Study 104

Briefers to the Study..... 107

List of Tables, Figures, and Boxes in the Report

Figure 1 DoD is increasingly employing autonomous capabilities across a diverse array of systems..... 5

Figure 2 Global autonomy startups are mapped (top); startup opportunity targets are categorized (bottom)..... 7

Figure 3 Machine intelligence ecosystem 8

Figure 4 Autonomy derives operational value across a diverse array of vital DoD missions 12

Figure 5 Combat veterans refresh unmanned aircraft skills..... 18

Figure 6 Oversight “on-the-loop” provides additional opportunities for human-machine partnership 19

Figure 7 Establishing an appropriate calibration of trust in autonomous systems 22

Figure 8 On-Line processor for system V&V and performance enhancement 34

Figure 9 Both inexpensive systems, such as the Flight Red Dragon Quadcopter (left) and more expensive ones, such as the Haiyan UUV (right), are becoming more capable and more available. 43

Figure 10 The Airborg (center top) capabilities are shown in the red boxes. The maximum gross take off weight of unmanned aircraft is compared with payload (left) and endurance (right)..... 44

Figure 11 The study evaluated many candidate projects and selected those that encompassed the range of benefits of autonomy 46

Figure 12 Elements of the ARGUS-IS Wide Area sensor are shown (left), along with the pace of technology change in sensor capabilities that can enable onboard autonomy (right)..... 50

Figure 13 Examples of seized media are shown (left), along with tools that can make sense of the stored information in real time (center). The resulting social network can reveal a real-time threat (right). 52

Figure 14 Current Mine countermeasure capabilities utilize two separate vehicles – an autonomous UUV for search and detection (left) and a vehicle remotely operated by a manned ship in the mine field for disposal (right). 56

Figure 15 Illustration of a cascading unmanned undersea vehicle concept. 62

Figure 16 Illustration of a concept for organic tactical ground vehicle support with an unmanned aircraft system. 66

Figure 17 Raft built entirely of fire ants, where the building follows a few simple rules and results in a buoyant structure that allows ants to survive until they reach dry land. 84

Figure 18 The Internet of Things is increasing rapidly in both numbers and types of smart objects. 88

Figure 19 Schematic of a drone collecting data from the Internet of Things in a typical neighborhood. 89

Figure 20 MAAP team responsibilities within the Joint Air Tasking Cycle 95

Defense Science Board Summer Study on Autonomy

Table 1	Projected capabilities for autonomous systems	11
Table 2	Value in participating in Stretch Problems.....	78
Table 3	Forms of Robotic Swarms.....	85
Table 4	Summary of Recommendations	99
Box 1	Amazon: A Commercial Enterprise Example	10
Box 2	Logistics as a Testbed for Software that Learns.....	31

Acronyms and Abbreviations

3D	three-dimensional
A2/AD	anti-access and area denial
AAR	air-to-air refueling
ACO	airspace control order
ACRS	autonomous cyber resilient systems
AFRL	Air Force Research Laboratory
AEODRS	Advanced EOD Robotics System
AFOSR	Air Force Office of Scientific Research
AI	artificial intelligence
ALLOREQ	allotment request
ALLOT	allotment
AOD	air operations directive
APU	auxiliary power unit
ARCIC	Army Capabilities Integration Center
ARGUS-IS	Autonomous Real-Time Ground Ubiquitous Surveillance Imaging System
ASA(ALT)	Assistant Secretary of the Army for Acquisition, Logistics, and Technology
ASD(R&E)	Assistant Secretary of Defense for Research & Engineering
ATO	Air Tasking Order
ATR	automatic target recognition
BGP	border gateway protocol
C2	command and control
CASCOM	Combined Arms Support Command
CCTV	closed circuit television
CERDEC	Communications-Electronics Research, Development and Engineering Center
CIA	Central Intelligence Agency
COI	community of interest
CONOPs	concepts of operation
CRASH	Clean-Slate Design of Resilient, Adaptive, Secure Hosts
DARPA	Defense Advanced Research and Projects Agency
DCA	defensive counter air
DIA	Defense Intelligence Agency
DISA	Defense Information Systems Agency
DLA	Defense Logistics Agency
DNS	domain name service
DoD	Department of Defense
DOMEX	document and media exploitation
DOT&E	Director of Operational Test and Evaluation
DOTMLPF	Doctrine, organization, training, materiel, leadership and education, personnel, and facilities
DRFM	digital radio frequency memory
DSB	Defense Science Board
DTCWC	Dynamic Time Critical Warfighting Capability
DTE	Office of Developmental Test and Evaluation
EMBERS	Early Model Based Event Recognition using Surrogates

Defense Science Board Summer Study on Autonomy

EO	electro-optical
EW	electronic warfare
EXCOM	executive committee
FAA	Federal Aviation Administration
FARP	forward arming and refueling point
FBI	Federal Bureau of Investigation
FFRDC	federally funded research and development center
GPS	global positioning system
IARPA	Intelligence Advanced Research and Projects Agency
IC	intelligence community
ICEWS	Integrated Conflict Early Warning System
IED	improvised explosive device
IHMC	Institute for Human and Machine Cognition
IoT	Internet of Things
IP	Internet protocol
IPA	Intergovernmental Personnel Act
ISR	intelligence, surveillance, reconnaissance
JAOP	joint air operations plan
JIDA	Joint Improvised-Threat Defeat Agency
JIEDDO	Joint Improvised Explosive Device Defeat Organization
JIPTL	joint integrated prioritized target list
JTCB	joint targeting coordination board
KI/CAS	killbox interdiction/close air support
LO/CLO	Low Observable and Counter Low Observable
M&S	modeling and simulation
MAAP	Master Air Attack Plan
MCCDC	Marine Corps Combat Development Command
MCM	mine countermeasures
MDAR	Mobile Detection Assessment and Response system
MRO	maintenance repair overhaul
MTRS	Mobile Tactical Robotic System
NAVSUP	Navy Supply Systems Command
NIPRNet	Non-classified Internet Protocol (IP) Router Network
NSA	National Security Agency
OASD(R&E)	Office of the Assistant Secretary of Defense for Research and Engineering
OCA	offensive counter air
ONR	Office of Naval Research
OSI	open source indicators
OTI	Office of Technical Intelligence
PCPAD-X	Planning & Direction, Collection, Processing & Exploitation, Analysis & Production, and Dissemination Experimentation
PEO-LCS	Program Executive Office Littoral Combat Ships
PNT	positioning, navigation, and timing
RAM	random access memory
R&D	research & development
REMUS	Remote Environmental Monitoring UnitS
RF	radio frequency

Defense Science Board Summer Study on Autonomy

RFID	RF identification
ROE	rules of engagement
SEAD	suppression of enemy air defenses
SIGINT	signals intelligence
SMART	self-monitoring, analysis, and reporting technology
SPAWAR	Space and Naval Warfare Systems Command
STIX	Structured Threat Information Expression
TACDOMEX	tactical document and media exploitation
TACE	Test Automation Center of Excellence
TAXII	Trusted Automated eXchange of Indicator Information
T&E	test and evaluation
TRADOC	United States Army Training and Doctrine Command
TRL	technical readiness level
TRMC	Test Resource Management Center
TTPs	tactics, techniques, and procedures
UARC	University Affiliated Research Center
UAS	unmanned aircraft system
UA	unmanned aircraft
UGV	unmanned ground vehicle
UOES	user operational evaluation system
USCYBERCOM	United States Cyber Command
USD(AT&L)	Under Secretary of Defense (Acquisition, Technology, and Logistics)
USD(I)	Under Secretary of Defense for Intelligence
USD(P&R)	Under Secretary of Defense for Personnel & Readiness
USSOCOM	United States Special Operations Command
USV	unmanned surface vehicle
UUV	unmanned undersea vehicle
V&V	verification and validation

Executive Summary

At the request of the Under Secretary of Defense for Acquisition, Technology, and Logistics (USD(AT&L)), the Defense Science Board (DSB) conducted a study on the applicability of autonomy to Department of Defense (DoD) missions. The study concluded that there are both substantial operational benefits and potential perils associated with the use of autonomy.

Autonomy delivers significant military value, including opportunities to reduce the number of warfighters in harm's way, increase the quality and speed of decisions in time-critical operations, and enable new missions that would otherwise be impossible. Autonomy is by no means new to the DoD. Fielded capabilities demonstrate ongoing progress in embedding autonomous functionality into systems, and many development programs already underway include an increasingly sophisticated use of autonomy.

Autonomy also delivers significant value across a diverse array of global markets. Both enabling technologies and commercial applications are advancing rapidly in response to market opportunities. Autonomy is becoming a ubiquitous enabling capability for products spanning a spectrum from expert advisory systems to autonomous vehicles. Commercial market forces are accelerating progress, providing opportunities for DoD to leverage the investments of others, while also providing substantial capabilities to potential adversaries.

This study concluded that DoD must accelerate its exploitation of autonomy—both to realize the potential military value and to remain ahead of adversaries who also will exploit its operational benefits.

Major recommendations

The issue of trust is core to DoD's success in broader adoption of autonomy. On the one hand, an autonomous system must be designed to operate in a trustworthy fashion with respect to the missions for which it was designed. On the other hand, an autonomous system must be designed so that humans (and/or machines) can straightforwardly determine whether, once it has been deployed, it is operating reliably and within its envelope of competence — and, if not, that appropriate action can be taken. Establishing trustworthiness at design time and providing adequate capabilities so that inevitable variations in operational trustworthiness can be assessed and dealt with at run time is essential, not only for operators and commanders, but also for designers, testers, policymakers, lawmakers, and the American public. The broad topic of trust shaped many of the recommendations that follow.

The first set of recommendations focuses on accelerating DoD's adoption of autonomous capabilities and includes:

- Tackling the engineering, design, and acquisition challenges
- Mitigating cyber issues introduced by increasingly autonomous and networked systems
- Creating new test and evaluation and modeling and simulation paradigms
- Integrating technology insertion, doctrine, and concepts of operations

- Developing an autonomy-literate workforce
- Improving technology discovery
- Improving DoD governance for autonomous systems
- Countering adversary use of autonomy

These interdependent recommendations are intended to build trust in autonomous systems, while at the same time accelerating DoD's progress.

The next set of recommendations focuses on strengthening the operational pull for autonomy, both by better understanding how others may use autonomy against the U.S., and by equipping our forces to counter such capabilities. These recommendations take the form of a series of demonstrations and experiments that demonstrate near-term military value while also building warfighter trust:

- Autonomous agents to improve cyber-attack indicators and warnings
- Onboard autonomy for sensing
- Time-critical intelligence from seized media
- Dynamic spectrum management for protection missions
- Unmanned undersea vehicles (UUVs) to autonomously conduct sea mine counter-measures missions
- Automated cyber response
- Cascaded UUVs for offensive maritime mining
- Organic tactical unmanned aircraft (UA) to support ground forces
- Predictive logistics and adaptive planning
- Adaptive logistics for rapid deployment

The final set of recommendations is intended to expand the envelope of technologies available for use on DoD missions. "Stretch problems" are proposed as a means to both strengthen the operational pull and mature the underlying technologies, such that they would be trusted for application on DoD missions:

- Early warning system for understanding global social movements
- Autonomous swarms that exploit large quantities of low-cost assets
- Intrusion detection on the Internet of things
- Autonomous cyber resilience for military vehicle systems
- Autonomous air operations planning

The recommendations of the study are briefly outlined in Table 4 at the end of the report, on page 102. Details for each recommendation are included in the relevant chapter sections and provide additional information to assist with their implementation.

Summary comments

Autonomy has many definitions and interpretations. For this reason, the report begins with an introductory section that defines the term and its context for the purposes of this study. This section also includes examples drawn from the commercial sector that illustrate diverse applications with operational relevance to the DoD.

The second major section addresses the issue of trust, and highlights both similarities and differences between military applications and commercial uses of autonomy. The study argues that an integrated approach—one that spans the entire lifecycle of a system—is needed to establish, maintain, and act upon current evaluations of trustworthiness in autonomous systems.

The remaining three sections motivate and elaborate on the major recommendations summarized above.

In summary, the study concluded that autonomy will deliver substantial operational value across an increasingly diverse array of DoD missions, but the DoD must move more rapidly to realize this value. Allies and adversaries alike also have access to rapid technological advances occurring globally. In short, speed matters—in two distinct dimensions. First, autonomy can increase decision speed, enabling the U.S. to act inside an adversary's operations cycle. Secondly, ongoing rapid transition of autonomy into warfighting capabilities is vital if the U.S. is to sustain military advantage.

1 Introduction

In November 2014, the Under Secretary of Defense for Acquisition, Technology, and Logistics directed the Defense Science Board to conduct a study to identify the science, engineering, and policy problems that must be addressed to facilitate greater operational use of autonomy across all warfighting domains. The study team identified opportunities for DoD to enhance mission efficiency, shrink life-cycle costs, reduce loss of life, and perform new missions—in both physical and virtual domains. The team concluded that there are both substantial operational benefits and potential perils associated with the use of autonomy.

Imagine if....

We could covertly deploy networks of smart mines and UUVs to blockade and deny the sea surface, differentiating between fishing vessels and fighting ships...

...and not put U.S. Service personnel or high-value assets at risk.

We had an autonomous system to control rapid-fire exchange of cyber weapons and defenses, including the real-time discovery and exploitation of never-seen-before zero day exploits...

...enabling us to operate inside the “turning radius” of our adversaries.

We had large numbers of small autonomous systems that could covertly enter and persist in denied areas to collect information or disrupt enemy operations...

...a “sleeper presence” on call.

We had large numbers of low-cost autonomous unmanned aircraft capable of adaptively jamming and disrupting enemy PNT capabilities...

...destroying their ability to coordinate operations.

We had autonomous high performance computing engines capable of not only searching “big data” for indicators of WMD proliferation, but of deciding what databases to search...

...to provide early warning and enable action

And imagine if we are unprepared to counter such capabilities in the hands of our adversaries.

About autonomy

Autonomy results from delegation of a decision to an authorized entity to take action within specific boundaries. An important distinction is that systems governed by prescriptive rules that permit no deviations are *automated*, but they are not *autonomous*. To be autonomous, a system must have the capability to independently compose and select among different courses of action to accomplish goals based on its knowledge and understanding of the world, itself, and the situation.¹

¹ Definitions for intelligent system, autonomy, automation, robots, and agents can be found in L.G. Shattuck, *Transitioning to Autonomy: A human systems integration perspective*, p. 5. Presentation at *Transitioning to Autonomy: Changes in the role of humans in air transportation* [March 11, 2015]. Available at human-factors.arc.nasa.gov/workshop/autonomy/download/presentations/Shaddock%20.pdf (Accessed June 2016.)

Recognizing that no machine—and no person—is truly autonomous in the strict sense of the word, we will sometimes speak of autonomous *capabilities* rather than autonomous *systems*²

The primary intellectual foundation for autonomy stems from artificial intelligence (AI), the capability of computer systems to perform tasks that normally require human intelligence (*e.g.*, perception, conversation, decision-making). Advances in AI are making it possible to cede to machines many tasks long regarded as impossible for machines to perform.

Intelligent systems aim to apply AI to a particular problem or domain—the implication being that the system is programmed or trained to operate within the bounds of a defined knowledge base. Autonomous function is at a system level rather than a component level. The study considered two categories of intelligent systems: those employing *autonomy at rest* and those employing *autonomy in motion*. In broad terms, systems incorporating *autonomy at rest* operate virtually, in software, and include planning and expert advisory systems, whereas systems incorporating *autonomy in motion* have a presence in the physical world and include robotics and autonomous vehicles. As illustrated in Figure 1, many DoD and commercial systems are already operating with varying kinds of autonomous capability.

Robotics typically adds additional kinds of sensors, actuators, and mobility to intelligent systems. While early robots were largely *automated*, recent advances in AI are enabling increases in *autonomous* functionality.

One of the less well-known ways that autonomy is changing the world is in applications that include data compilation, data analysis, web search, recommendation engines, and forecasting. Given the limitations of human abilities to rapidly process the vast amounts of data available today, autonomous systems are now required to find trends and analyze patterns. There is no need to solve the long-term AI problem of general intelligence in order to build high-value applications that



Figure 1 DoD is increasingly employing autonomous capabilities across a diverse array of systems.

² See J.M. Bradshaw, R.R. Hoffman, M. Johnson, and D.D. Woods, “The Seven Deadly Myths of ‘Autonomous Systems,’” *IEEE Intelligent Systems* 28, no. 3. [May/June 2013], pp. 54-61. Available at jeffreymbradshaw.net/publications/IS-28-03-HCC_1.pdf (Accessed April 2016.)

exploit limited-scope autonomous capabilities dedicated to specific purposes. DoD's nascent Memex program is one of many examples in this category.³

Rapid global market expansion for robotics and other intelligent systems to address consumer and industrial applications is stimulating increasing commercial investment and delivering a diverse array of products. At the same time, autonomy is being embedded in a growing array of software systems to enhance speed and consistency of decision-making, among other benefits. Likewise, governmental entities, motivated by economic development opportunities in addition to security missions and other public sector applications, are investing in related basic and applied research. Applications include commercial endeavors, such as IBM's Watson, the use of robotics in ports and mines worldwide, autonomous vehicles (from autopilot drones to self-driving cars), automated logistics and supply chain management, and many more. Japanese and U.S. companies invested more than \$2 billion in autonomous systems in 2014, led by Apple, Facebook, Google, Hitachi, IBM, Intel, LinkedIn, NEC, Yahoo, and Twitter.⁴

A vibrant startup ecosystem is spawning advances in response to commercial market opportunities; innovations are occurring globally, as illustrated in Figure 2 (top). Startups are targeting opportunities that drive advances in critical underlying technologies. As illustrated in Figure 2 (bottom), machine learning—both application-specific and general purpose—is of high interest. The market-pull for machine learning stems from a diverse array of applications across an equally diverse spectrum of industries, as illustrated in Figure 3.

Autonomy accelerates enterprise performance

Commercial enterprises are enhancing performance through the use of autonomy to exploit advances in processing power, big data analytics, and networked systems that leverage diverse and distributed sensor arrays. Opportunities exist for DoD to enhance mission performance by employing autonomy at rest and autonomy in motion, both supporting human-machine collaboration. The commercial sector is a lucrative source of both basic capability and best practices relevant to many such opportunities.

Over the past several years, autonomous sensing and decision-support techniques have been demonstrated by the commercial sector. The following are only a few examples:

Footage from the estimated 52,000 government-operated closed circuit television (CCTV) cameras in the United Kingdom, along with the 1.85 million total cameras across the country, is used in as many as 75 percent of the 3.9 million criminal cases annually.^{5,6}

³ W. Shen, *Memex*. Available at www.darpa.mil/program/memex (Accessed June 2016.)

⁴ Quid, referenced in *Robot Revolution—Global Robot and AI Primer*. [Bank of America Merrill Lynch, December 16, 2015].

⁵ *The Price of Privacy: How local authorities spent £515m on CCTV in four years*. [Big Brother Watch, February 2012]. Available at www.bigbrotherwatch.org.uk/files/priceofprivacy/Price_of_privacy_2012.pdf. (Accessed January 2016.)

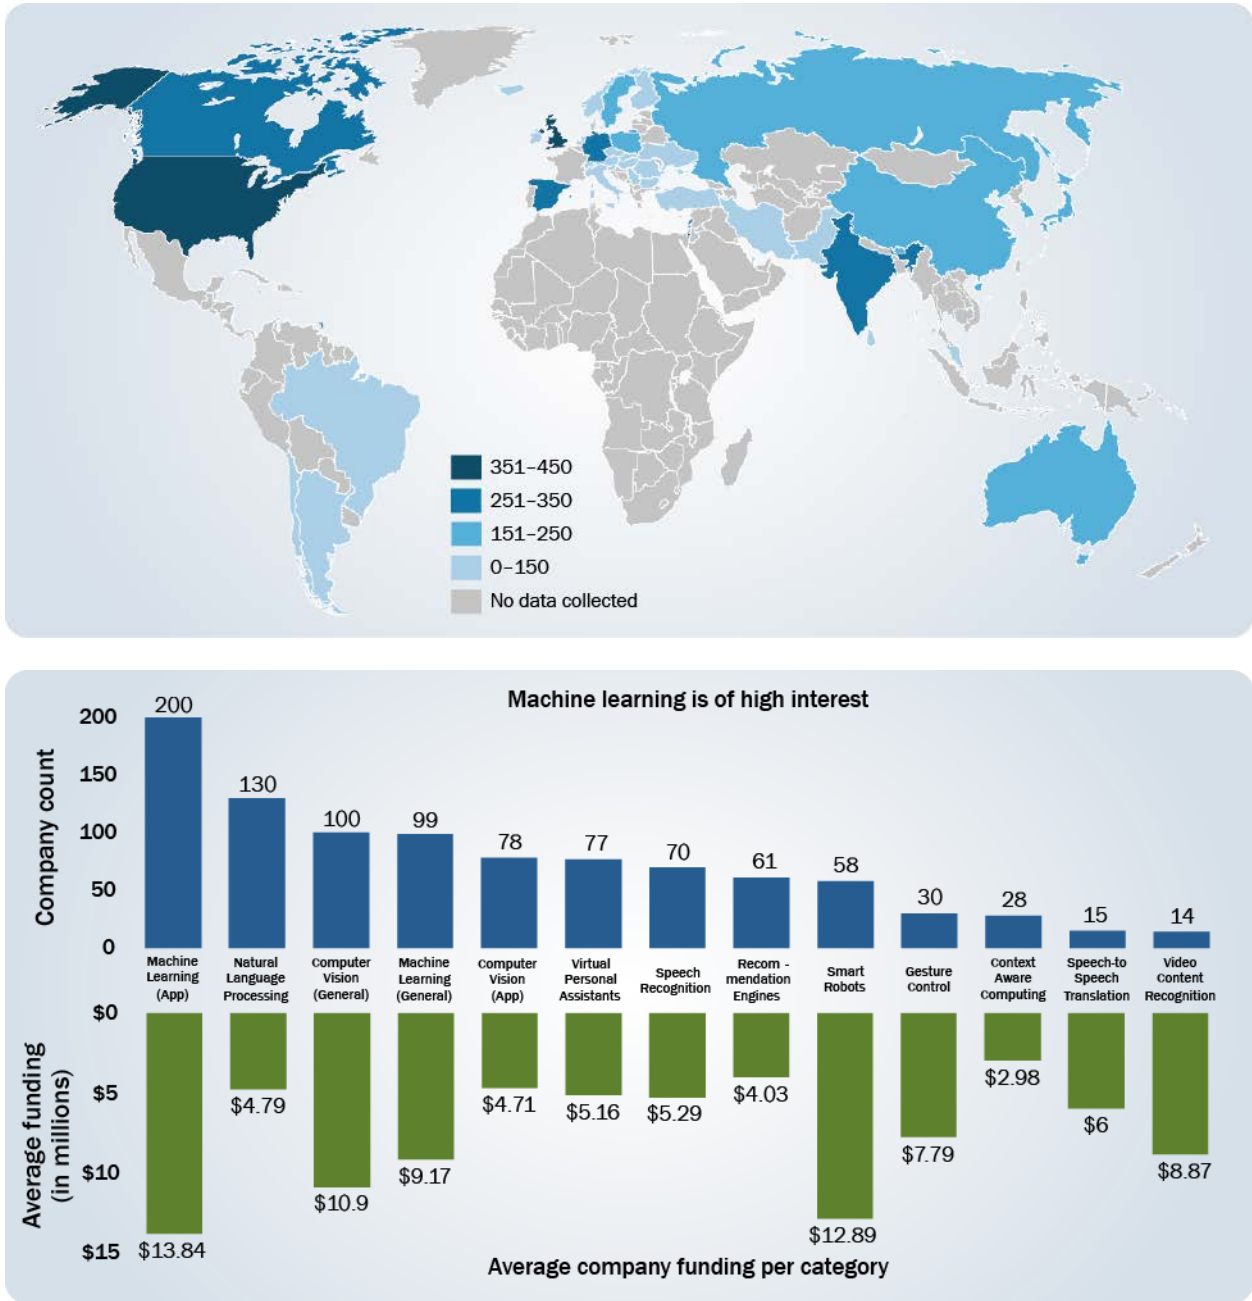


Figure 2 Global autonomy startups are mapped (top); startup opportunity targets are categorized

- Automated video content analysis software searches long videos for key events and can incorporate human facial and gait recognition.

⁶ P. Taylor and S. Bond, eds., *Crimes detected in England and Wales 2011/12*. [July 2012]. Available at www.gov.uk/government/uploads/system/uploads/attachment_data/file/116435/hosb0812.pdf. (Accessed January 2016.)



Figure 3 The machine intelligence ecosystem comprises a diverse array of applications and industries.

SOURCE: www.shivonzilis.com/machineintelligence

- The U.S. credit, debit, and prepaid card industry monitors more than 1,200 transactions per second with automated tools that identify fraudulent transactions within 10 milliseconds on a customer base of more than 1.3 billion cards.⁷
- IBM’s Watson for Oncology system has ingested over 12 million pages of medical content, including over 200 medical textbooks and 290 medical journals, and provides oncologists with recommended courses of treatment within 30 seconds, largely from unstructured records.⁸
- Google machine-learning password classifiers authenticate users via multiple signals, such as Internet protocol (IP) address, geolocation, and login time, resulting in 99.5 percent reduction in accounts compromised by spammers.⁹
- The mining industry has integrated autonomy into many systems. Both excavating and hauling vehicles are equipped with vehicle controllers, high precision global positioning system (GPS) sensors, and obstacle detection. These features allow safe operation through a complex load, haul, and dump cycle, and the ability to integrate with other vehicles, people, and obstacles that are also part of the autonomous system.¹⁰

⁷ The Nilson Report, *U.S. General Purpose Cards - Midyear 2015*, Issue 1069. [August 2015].

⁸ *IBM Watson for Oncology*. Available at www.ibm.com/smarterplanet/us/en/ibmwatson/watson-oncology.html. (Accessed January 2016.)

⁹ Google, *An update on our war against account hijackers*. [February 19, 2013]. Available at googleblog.blogspot.com/2013/02/an-update-on-our-war-against-account.html. (Accessed January 2016.)

¹⁰ Warner Norcross & Judd, *Trucking, Mining Industries Blazing a Path to Vehicle Autonomy*. [January 14, 2015]. Available at www.wnj.com/Publications/Trucking-Mining-Industries-Blazing-a-Path-to-Vehic (Accessed June 2016.)

Lucrative global markets are attracting ongoing investment, advancing basic technologies while also targeting delivery of capabilities relevant to DoD missions. The following are only a few examples:

- Vehicles that drive themselves in limited circumstances (*e.g.*, on the freeway and in traffic jams in good weather) will begin to enter the market in 2016 and may be on the road in large numbers by 2017. “The global market for autonomous vehicles is projected to grow from \$42 billion in 2025 to \$77 billion by 2035,” and Japan and Western Europe are likely to be early adopters of the technology.¹¹
- The advanced autonomous waterborne applications initiative was funded by Tekes (Finnish Funding Agency for Technology and Innovation) in 2015 and is led by Rolls-Royce. It will bring together “universities, ship designers, equipment manufacturers, and classification societies to explore the economic, social, legal, regulatory and technological factors which need to be addressed to make autonomous ships a reality.”¹²
- A prototype hybrid drone, capable of flying as a traditional helicopter as well as a fixed-wing aircraft, was demonstrated at a conference in Denmark in 2015. The drone, which stems from a project led by the Technical University of Denmark, can maneuver with an accuracy of five centimeters and is intended to enable public and private enterprises to collect data faster and more accurately than previously possible.¹³

Enterprises are already achieving accelerated performance through the integration of available autonomy-enabling technologies. At the same time, their successes are stimulating increased investment across a broad array of underlying technologies that together will spawn new types of autonomous systems. See Box 1 on page 10 for an example of how this strategy evolved at Amazon.

In this study, four broad categories are used to characterize underlying technologies critical to the development of autonomous systems: *Sense*; *Think/Decide*; *Act*; *Team*. The relative importance of each category varies by application, as do the drivers that spawn new capabilities. Advances to *Sense* are driven by a diverse array of applications (including, but not limited to, autonomous systems) that share a common need to reduce sensor size, weight, and power requirements. Artificial intelligence, which enables the *Think/Decide* functionality in autonomous systems, is benefiting from advances in computational power as well as availability of vast data sets. Demand for productivity growth via automation was an early driver of advances in actuators and mobility that *Act*; and that demand is growing as robotics become more intelligent and new applications are emerging. A growing number of applications require human-machine teaming and collaboration—letting each do what it does best, but also imposing new requirements on the underlying *Team* technologies. Table 1 summarizes

¹¹ Boston Consulting Group, *The Autonomous Vehicle: The Car of the Future*. Available at on.bcg.com/1HNAHKH (Accessed June 2016.)

¹² Rolls-Royce, *Rolls-Royce to Lead Autonomous Ship Research Project*. [July 2, 2015]. Available at www.rolls-royce.com/media/press-releases/yr-2015/pr-02-07-15-rolls-royce-to-lead-autonomous-ship-research-project.aspx. (Accessed January 2016.)

¹³ Homeland Defense & Security Information and Analysis Center, *Denmark Creates Hybrid Smart UAV*. [August 3, 2015] Available at www.hdiac.org/node/2069 (Accessed January 2016.)

Box 1: Amazon: A Commercial Enterprise Example

Commercial enterprises are now leveraging autonomy as a means to create entirely new business models. A well-known example is Amazon, which began as an online bookseller and has expanded to offer online retail, computing services, consumer electronics, and digital content, as well as other local services such as daily deals and groceries. While autonomy has not been the only factor in Amazon’s success, Amazon has employed autonomy effectively in areas where its business model does not allow human activity or where autonomy can break a bottleneck.

Amazon began as a dot-com company where its primary advantages were perceived to be elimination of bricks and mortar costs and the broad reach of the early Internet. It could sell and deliver physical books at costs lower than competitors.

The early differentiator for physical bookstores was their ability to know the tastes and interests of frequent customers, which added to loyalty and sales. Amazon moved quickly to develop a recommendation engine that captured and assessed large amounts of data about its customers. It also sought, aggregated, and published customer reviews, along with customer-generated questions and answers about its wide range of products. As Amazon expanded into other retail areas, it has expanded and refined this use of autonomy at rest.

Amazon also employs autonomy in motion by custom-packaging books for shipment using a fully autonomous line. When Amazon expanded into a broader range of retail items, warehouse operations became a bottleneck. In 2012, Amazon purchased Kiva Systems, a manufacturer of mobile robotic fulfillment systems.

The figure below displays the revenue performance of Amazon over the past twenty years. It is evident that Amazon leadership embraces the benefits of autonomy as part of that picture.

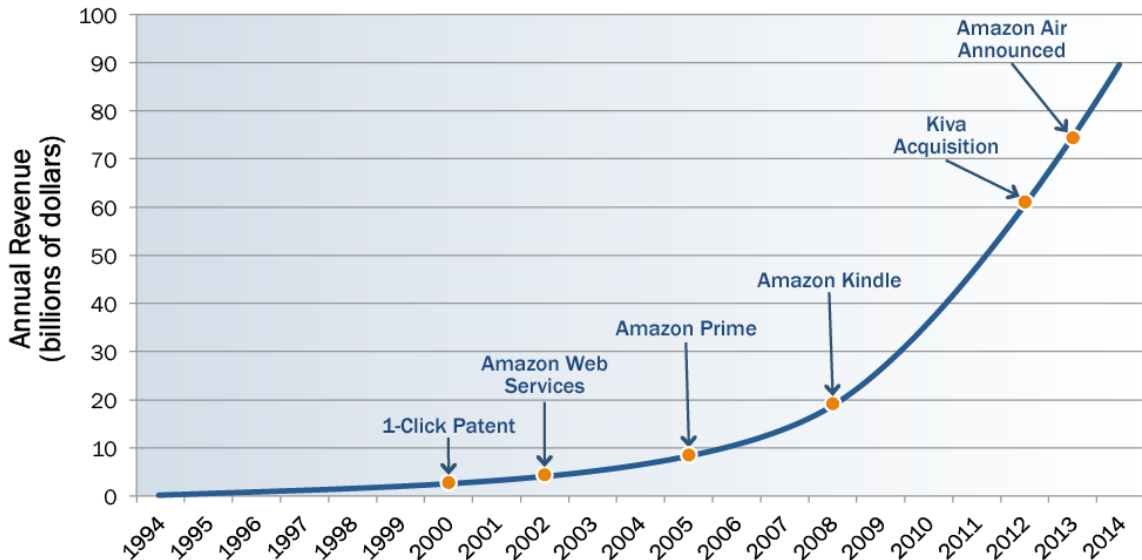


Table 1 Projected capabilities for autonomous systems

SENSE: Sensors, Perception, Fusion

- *Available today:* Full-spectrum sensing (EM, bits, vibration, chemical...); Object recognition
- *Likely available near term:* Human senses (sight, smell...); Integration of perception with motor skills
- *May be available long term:* High-fidelity touch; Scene understanding

THINK/DECIDE: Analysis, Reasoning, Learning

- *Available today:* High-volume computational throughput and data architectures; Algorithm variety and complexity; Task-specific, rule-based decision; Rules; Learning from training data, sentiment analysis
- *Likely available near term:* Explicit and testable knowledge representation; Anomaly recognition; Option generation, pruning; Social and behavioral models; Culturally informed, values-based reasoning; Transparent decision logic; C2 for many nodes; Learning by doing, watching
- *May be available long term:* Goal definition; Abstraction, Skills transfer; Inference; Empathy; General purpose, idea-based reasoning; Judgment, intuition

ACT: Motion, Manipulation

- *Available today:* Navigation (routing); Strength, endurance
- *Likely available near term:* Navigation (obstacle avoidance); Agility, dexterity
- *May be available long term:* Navigation (dense, dynamic domains); High degree of freedom actuator control

TEAM: Human/machine, Machine/machine, Info exchange

- *Available today:* High man:machine ratio; Rule-based coordination of multiple platforms; High-volume communications and data transfer
 - *Likely available near term:* Observability and directability; Provably correct emergent behavior; Trustworthiness and trust calibration under defined conditions; Natural language processing
 - *May be available long term:* Shared “mental models,” mutual predictability; Understanding intent; Fully adaptive coordination; Implicit communication
-

the study team’s assessment of the availability of critical underlying technologies in each of these four categories.

Military value and current DoD uses

The DoD has strategically increased its adoption of robotics and unmanned vehicle systems in the last decade, but the vast majority of the systems are remotely operated rather than autonomous. Recent programs show a progression from pre-programming and remote control to autonomous functionality, but progress has been slow. The Department is engaged in R&D across many aspects of autonomy, but has not yet addressed the R&D needed to overcome the systemic challenges to the widespread use of autonomy.

The growth in robotics and unmanned systems was largely driven by perceived improvements in performance and cost. The actual advantages are more complex. Safety improves by reducing the lethality of warfare and the ability to adopt riskier tactics because a system is unmanned. Accuracy also improves, with more endurance, range, and speed in comparison to manned vehicles. Systems are also more flexible and more mobile. Autonomy also enables the execution of new missions—particularly in domains such as cyber and electronic warfare, in which decision speed is critical to success. Figure 4 summarizes the relative value of autonomy against key mission parameters.

Defense Science Board Summer Study on Autonomy

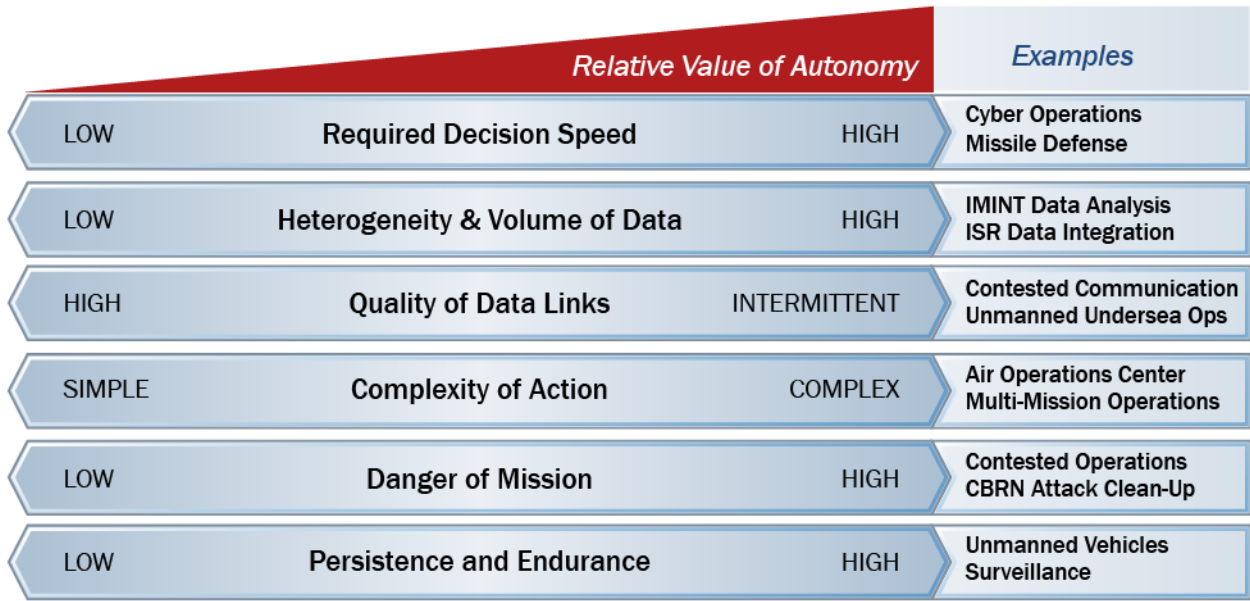


Figure 4 Autonomy derives operational value across a diverse array of vital DoD missions.

While more difficult to quantify for the DoD, commercial enterprises are increasingly demonstrating substantial cost savings through the adoption of autonomy. From a DoD perspective, UA are believed to result in cost reductions over time; other applications may deliver similar benefits—particularly in areas such as logistics in which DoD can leverage advances driven by large commercial markets. But, as indicated in Figure 4, the value of autonomy to DoD missions extends well beyond cost reductions.

U.S. defense spending on UA has grown to almost \$3 billion in 2016.¹⁴ The number of unmanned aircraft has grown to more than 11,000, or 40 percent of all aircraft.¹⁵ In the U.S., the Federal Aviation Administration (FAA) introduced a national framework for registering UA at the end of 2015, and over 181,000 drones were registered in the first two weeks, dwarfing the DoD population.¹⁶ The U.S. is not alone; ninety countries around the world operate UA, with thirty armed UA programs established or in development.¹⁷

While commercial market opportunities will advance many underlying technologies that are critical to DoD applications, there are areas in which DoD cannot rely on the commercial industrial base to develop needed capabilities, including the following:

¹⁴ U.S. Department of Defense, March 2016.

¹⁵ J. Gertler, *U.S. Unmanned Aerial Systems*. [Congressional Research Service CRS R42136, January 3, 2012] and *How many UAVs for DoD?* [CRS IN10317, August 27, 2015].

¹⁶ R. Marsh and H. Kelly, “181,000 drones registered with FAA in two weeks,” *CNN Money* [January 6, 2016]. Available at money.cnn.com/2016/01/06/technology/faa-drone-registration/index.html (Accessed June 2016.)

¹⁷ K. Saylor, *A World of Proliferated Drones: A Technology Primer* [Center for a New American Security, June 2015]. Available at www.cnas.org/world-of-proliferated-drones-technology-primer#.V1L4myGEBqM (Accessed June 2016.)

- The commercial sector can often structure the environment to simplify decision-making for an autonomous system. In many DoD missions, the operation will be conducted in an environment as encountered, with no opportunity to structure the environment or map it to make the problem easier for the autonomous system.
- The commercial sector can decompose a problem and first tackle the easy parts using autonomous systems; hard parts of the problem can still be solved using humans when this is more cost-effective (or until the required technology matures). In some DoD applications—for example, operating in contested environments—using humans to address the limitations of autonomous systems may prove challenging (or impossible).
- Commercial applications rarely deal with intelligent adversaries that try to actively defeat the technology in use, although cyberspace is one notorious counter example. Many DoD applications will need autonomy capabilities that are robust enough to cope with the deployment of deception, assault, and counter-autonomy technologies by adversaries.

In summary, the study concluded that autonomy has the potential to deliver substantial operational value across a diverse array of vital DoD missions, and that the DoD should speed its adoption to realize the potential benefits across a diverse array of missions.

Study approach

Throughout the course of the study, the team grappled with difficult questions, including the following:

- How can systems be “future-proofed” to enable autonomy technology advances to gracefully—and rapidly—upgrade?
- How can military advantages be sustained, given the rapidly advancing commercial global technology base?
- How can autonomous systems benefit from test and evaluation (T&E) when every time they are used they change themselves?... and how can we verify that systems are learning the right things?
- How can the U.S. compensate for adversaries who may have very different rules of engagement to employ lethal autonomous systems?
- How can autonomous systems be made more “cyber-proof”—given that every new capability introduces a vulnerability?
- Will we increase trust in autonomous systems if each includes a “black box” audit trail that can explain why they did what they did?
- How can we counter an adversary's autonomous systems? And how can we protect our autonomous systems from an adversary's interference?
- Can autonomous systems save lives of U.S. Service personnel as well as those of innocent non-combatants? ...and could the U.S., therefore, be less deterred from pursuit of foreign policy objectives?

2 Trustworthiness and Trust in Autonomous Systems

Most commercial applications of autonomous systems are designed for operation in largely benign environments, performing well-understood, safe, and repetitive tasks, such as routing packages in a fulfillment center warehouse. Design for commercial systems rarely considers the possibility of high-regret outcomes in complex, unpredictable, and contested environments. In military operations, these can include an adversary whose goal is to neutralize the use and effectiveness of such systems, either through deception, direct force, or increased potential for collateral damage or fratricide. Although commercial applications are gradually expanding beyond these controlled environments, *e.g.*, self-driving cars, delivery drones, and medical advisory systems, fielded autonomous systems do not yet face a motivated adversary attempting to defeat normal operations.

Trust is complex and multidimensional.¹⁸ The individual making the decision to deploy a system on a given mission must trust the system; the same is true for all stakeholders that affect many other decision processes. Establishing trustworthiness of the system at design time and providing adequate indicator capabilities so that inevitable context-based variations in operational trustworthiness can be assessed and dealt with at run-time is essential, not only for the operator and the Commander, but also for designers, testers, policy and lawmakers, and the American public.

Key issues and barriers to trust

Methods for ensuring trustworthiness include careful design and implementation of systems to assure key attributes including high levels of competence, reliability, and integrity. Of course, designers are expected to embed these attributes in development and manufacturing of autonomous weapons systems. However, such attributes may be undercut by characteristics associated with hybrid, multi-party human-machine teams, including:

Lack of human-analog sensing and thinking by the machine. Because an autonomous system may have different sensors and data sources than any of its human teammates, it may be operating on different contextual assumptions of the operational environment. In addition, for some specific algorithm choices—such as neuromorphic pattern recognition for image processing, optimization algorithms for decision-making, deep neural networks for learning, and so on—the “reasoning” employed by the machine may take a strikingly different path than that of a human decision-maker.

Lack of self- or environmental awareness by the machine. Self-awareness can be as simple as understanding its own system health, such as battery level, or more subtle, such as knowing when it is operating outside of its original design boundaries or assumptions. Environmental awareness includes conventional sensing of the environment; such as icing on a wing or jammed communications, as well as more subtle effects such as GPS spoofing. Of course, it is not sufficient

¹⁸ R. R. Hoffman, Matthew Johnson, J.M. Bradshaw, and Al Underbrink. “Trust in Automation.” *IEEE Intelligent Systems*, Vol. 28, Issue 1 [January/February 2013], pp. 84-88.

for a machine to be aware of changes in itself and its operating environment, it must also be able to adapt to those changes flexibly and effectively.

Low observability, predictability, directability, and auditability. Autonomous systems not only need to operate reliably and within their envelope of competence in dynamically varying and complex operational contexts, but also to be able to make relevant information *observable* to human and machine teammates. Moreover, even if machines are competently designed to enable observation of *current* state and effects, they may not incorporate sufficient *anticipatory* indicators to allow other human and machine teammates to ensure *predictability*. In addition, when something goes wrong, as it will sooner or later, autonomous systems must allow other machine or human teammates to intervene, correct, or terminate actions in a timely and appropriate manner, ensuring *directability*.¹⁹ Finally, the machine must be *auditable*—in other words, be able to preserve and communicate an immutable, comprehensible record of the reasoning behind its decisions and actions after the fact.²⁰

Low mutual understanding of common goals. If humans and autonomous machines are to work effectively together, they need common goals and a mutual knowledge of those common goals. Many of the commercial aircraft accidents in the 1990s associated with automation occurred when the flight crew had one goal (*e.g.*, staying on the glide slope during an approach) and the flight management computer had another (*e.g.*, executing a go-around). Improved training of personnel can help to address such issues. Dealing with future autonomous systems, however, may demand more than a one-sided approach, and may include increasing the machine's awareness of what the operator is trying to achieve.

Ineffective interfaces. Conventional computer interfaces, such as mouse point-and-click, can slow communications between humans and machines and inhibit the coordination and cooperation needed in time-sensitive or high-risk situations. Better interfaces, such as natural-language processing or more effective visualizations of complex information and situations can help mitigate such issues.

Systems that learn. Machines are being developed with experience that change their capabilities and limitations and adapt to their use and environment. Such systems will outgrow their initial verification and validation and will require more dynamic methods to perform effectively throughout their lifecycle.

Autonomy in support of command and control

One of the most contentious applications of autonomy is for command and control in military operations or warfighting, but the potential benefits are real. The time for concepts of operations

¹⁹ On observability, predictability, and directability, see M. Johnson, M., J.M. Bradshaw, P. J. Feltovich, C. M. Jonker, M. B. van Riemsdijk, and M. Sierhuis, "Coactive Design: Designing Support for Interdependence in Joint Activity," *Journal of Human-Robot Interaction*, Vol. 3, No. 1. [2014], pp. 43-69.

²⁰ M.R. Endsley, "Measurement of Situation Awareness in Dynamic Systems," *Human Factors*, 37(1) [March 1995], pp. 65-84; and M.R. Endsley, "Building Resilient Systems: Incorporating Strong Human-system Integration," *Defense AT&L Magazine* [January–February 2016, 2015]. Available at dau.dodlive.mil/2015/12/28/building-resilient-systems-via-strong-human-systems-integration (Accessed March 2016.)

(CONOPs) development, target selection, and mission assignments can be significantly reduced, and, during combat operations, commanders could be better equipped to respond to changing situations and redirect forces. While commanders understand they could benefit from better, organized, more current, and more accurate information enabled by application of autonomy to warfighting, they also voice significant concerns.

Implementation of autonomous capabilities will require significant changes in command and control concepts. A previous DSB study acknowledged the importance of addressing command and control of autonomous systems, but found that it was an unsolved problem and did not address it further.²¹

Whether mediated by man or machine, all acts, but especially acts related to warfighting, must be executed in accordance with policy and so, in some sense, there is no completely autonomous behavior. Any use of autonomy must conform to a substantive command and control regime laying out objectives, methods and express limitations to ensure that autonomous behavior meets mission objectives while conforming to policy.²²

In fact, most autonomous combat systems will and should act under the guidance and instructions of a field commander who will exercise direct oversight. Initial use of autonomous systems for combat will likely assist commanders and their staffs with developing situational awareness and planning missions. The volume and velocity of the data used by the underlying system will change the pace of operations. For example, current air operations planning often involves a several-day process, beginning with identification of objectives; moving to general target selection; to intelligence support identifying particular targets; to determination of final plans by the Commander balancing risks and potential value in achieving objectives; coordination of air, land, and sea assets; and, finally, mission execution and battle damage assessment. Given human limitations, each stage results in static point in time, and planning is generally a linear and time-consuming process. Because planning often needs to respond to new information, autonomous systems will greatly accelerate the pace of information update and can suggest significant plan changes far more quickly.

Commanders will not only need to develop models to calibrate understanding of machine generated information, but will need better automated tools to adjust to the pace of update. This requires careful design of autonomous systems so they can explain and justify recommendations in principled terms that build trust between commanders, staff, and machine-generated output as well as develop a concrete understanding of mission outcomes that depend critically on the data and methodology employed by the autonomous systems.

²¹ Defense Science Board, *Report of the Task Force on the Role of Autonomy in DoD Systems*. [2012], p 16. Available at www.acq.osd.mil/dsb/reports/AutonomyReport.pdf (Accessed March 2016.)

²² For one approach to this problem, see A. Uszok, J.M. Bradshaw, J. Lott, M. Johnson, M. Breedy, M. Vignati, K. Whittaker, K. Jakubowski, and J. Bowcock, "Toward a Flexible Ontology-Based Policy Approach for Network Operations Using the KAOs Framework," *Proceedings of the 2011 Military Communications Conference (MILCOM 2011)*. [New York City, NY: IEEE Press, November 2011], pp. 1108-1114.

Commanders and their staff need to both exercise oversight of fielded autonomous systems, and have access to efficient mechanisms to develop control datasets that teach autonomous systems so they can react in well-understood ways to unexpected and subtle changes. This requires new approaches to human factors that are informed by warfighting practice and tradition. It also requires development of commander and staff training.

Human-machine collaboration and combat teaming

Applications of autonomy engage both human and machine throughout the system lifecycle. Certain roles will remain the purview of the human, others will be shared, and some tasks will be implemented solely by machines. The specific roles of humans will vary by mission and over time.

The overall initial operational design will be done primarily by humans. This includes tasks such as defining the envelope and boundaries of potential autonomous behavior, identifying general behavioral parameters, and establishing the range of rules of engagement within which the system will be designed to operate.

Development of the system may be shared. Humans would define behavioral parameters for a range of missions and train the system for that range of missions; machines will learn behaviors in both standard and unplanned situations while adhering to the rules of engagement.

During deployment, the machine might continue to learn while executing the mission within established bounds. Alternately, the machine might be operationally limited to execution of its trained behaviors. Humans in the loop will set mission parameters and may continue to refine the rules of engagement by verifying operations across varied operational conditions. Some operational modes may dynamically specify what responsibilities humans and machines will perform and share throughout the mission. Only humans would retain the rights to change mission parameters or change the rules of engagement.

Effective human-machine collaboration requires that team members share common goals. This requires mutual understanding of the common goals, even though the goals may be expressed in different frameworks and semantics. It is critically important for the human operator to have a good understanding of the machine's goals, otherwise human-machine team failures are bound to occur. It may be equally important to provide a means for a machine to understand team goals. Significant autonomy capabilities will derive from a machine's ability to infer the commander's intent and to act adaptively in a non-pre-programmed fashion, and in doing so, being able to deal with unanticipated situations not foreseen by either the designer or the operator.

In addition, because teamwork between humans and autonomous machines may require various levels of communications, system architectures must support a variety of machine-to-machine and machine-to-human communications links, with the latter focused on improving the link through emerging technologies (*e.g.*, neurolinguistics programming). Finally, humans and systems must train together (as shown in **Figure 5**) to develop CONOPs, as well as to achieve skilled human-machine team performance across a wide range of missions, threats, environments, and operators.

It is important for the human operator to understand the basic competence of the machine, and, during operations, when it may be operating outside of its design assumptions or operational boundaries. A direct approach would simply encode the operational parameters; a more comprehensive approach would rely on in-line modeling and simulation to instruct behavior, including nominal and out-of-envelope behaviors. Models, tools, and datasets must be developed to deal with the fact that operational boundaries are situational, may evolve, and may violate the original system design assumptions.

Because many autonomous system behaviors will change over time due to learning, discrepancies may occur between actual system performance and operator expectations, possibly leading to teammate surprise during operations. One approach to ameliorate this situation is to provide any human teammate with a training history of the autonomous system's experience or competencies, analogous to current military practice based on a human teammate's service history and rank. Alternatively, a human and machine team rehearsing missions together could provide each teammate with implicit expectations of behavior based on current capabilities.

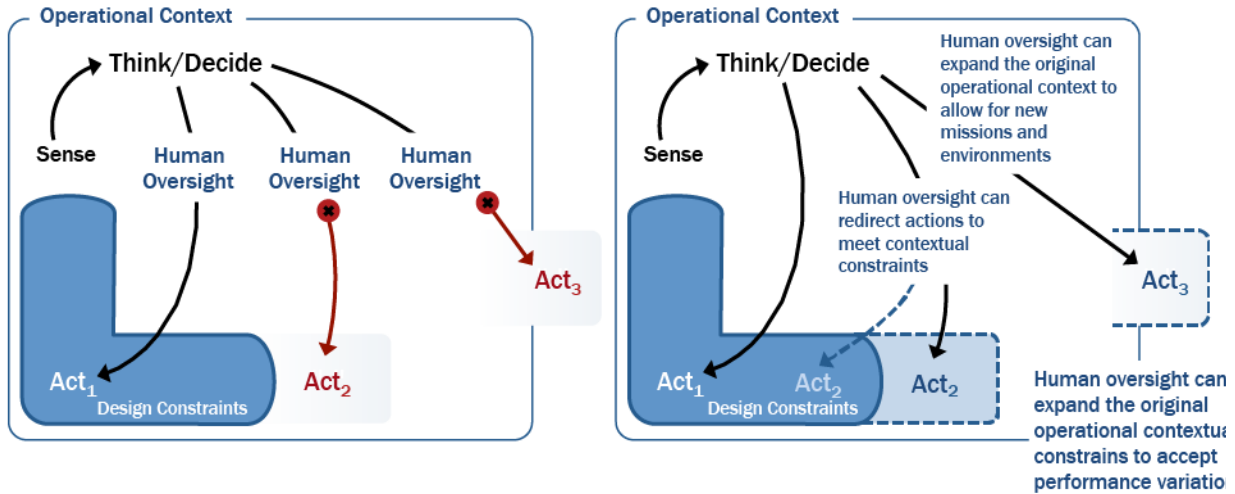
In addition, if a human or machine changes capabilities and limitations—learns—over time, more dynamic and adaptive methods are needed to reliably verify and validate system behavior for future engagements. This will necessitate a significant change in today's relatively static T&E practices.

A wide range of design approaches can be implemented to provide human and machine teammates with insight and foresight with respect to a system's integrity. The autonomous system design needs to support self-awareness (including an assessment of internal system health and external environmental and contextual factors) and anticipation of the future to support its own reasoning and adaptation processes. It also needs to provide this information in response to commands or requests from other machine or human teammates. This information, in conjunction with an understanding of the system's own design performance limitations or constraints, can be used for both self- and external assessment of a system's operation within its design envelope;



Figure 5 Combat veterans regularly refresh their unmanned aircraft skills.

SOURCE: News, February 12, 2013, www.army.mil



Human oversight is in the loop to

- confirm actions (Act1)
- deny actions outside designed constraints (Act2)
- deny actions outside the operational context (Act3)

Human oversight is on the loop as needed to

- allow actions outside designed constraints (Act2)
- allow actions outside the operational context (Act3) – and take advantage of evolving opportunities.

Figure 6 Oversight “on-the-loop” provides additional opportunities for human-machine partnership.

violations or anticipated violations of this envelope can be communicated to other teammates, human or machine, for appropriate mitigation. In addition, designing for transparency and traceability of a system’s situational awareness and decision-making can help to explain decisions taken, both operationally and forensically, given that adequate communications and semantic links are included. Figure 6 illustrates how the range of operational options might be dynamically expanded through the implementation of “human-on-the-loop” intervention.

Because of the potential for high-regret outcomes in complex scenarios, especially via adversary attacks against system vulnerabilities, designs need to be red teamed at all phases of concept evaluation, development, acquisition, and employment. One approach is to work with an in-line suite of consistent and related models and simulations employed and verified at all stages of the acquisition and deployment cycle.

By red-teaming early in design and engineering, concepts can be de-selected or augmented. With suitable modeling and simulation (M&S)

representation, the effectiveness of conventional adversary actions (e.g., kinetic weapons, electronic warfare) can be evaluated, as well as any potential vulnerabilities to more subtle attacks such as sensor spoofing, communications intercepts, or embedded cyber-attacks. As development progresses, the level of the M&S representation and tools should be made more sophisticated, providing greater fidelity of not only capabilities, but also of potential vulnerabilities, setting the stage for another round of red teaming, at a higher level of resolution. At

A red team is an independent group that challenges the organization—its doctrine, CONOPs, and its systems—with the mindset of an adversary to improve the process or the product.

a sufficient level of maturity, ideally soon after concept development, live and constructive simulations can be used to introduce humans in-the-loop and on-the-loop to help develop CONOPs and tactics for the human-machine teams. Later in the development cycle, more sophisticated M&S facilities can be used for both training and dealing with vulnerabilities, to provide additional opportunities for red teaming.

Cultural, policy, and legal issues

The overwhelming majority of potential military applications for autonomy are non-lethal and offer the potential for improved efficiencies or entirely new capabilities. Skepticism about the employment of autonomy in military operations is almost wholly focused on the use of autonomous weapons systems with potential for lethality. For this reason, any new autonomous capability may meet with resistance unless DoD makes clear its policies and actions across the spectrum of applications.

DoD recently undertook a comprehensive review of policy and procedures associated with autonomy in weapon systems, and produced Directive 3000.09 on “Autonomy in Weapon Systems” in November 2012. As stated in the purpose statement of the directive:

DoDI 3000.09 establishes DoD policy and assigns responsibilities for the development and use of autonomous and semi-autonomous functions in weapon systems, including manned and unmanned platforms; establishes guidelines designed to minimize the probability and consequences of failures in autonomous and semi-autonomous weapon systems that could lead to unintended engagements.

The Directive provides comprehensive guidance to all of the stakeholders concerned with deployment of autonomous systems. The most important policy points to be made from the Directive that are relevant to public concerns are that there are no proscriptions for the development of lethal autonomous weapon systems, but their development would require a much more rigorous review and approval process. Emphasis is placed on assurance that the system will perform as intended and be as immune as possible to unintended loss of control, capture, or compromise by the adversary. Moreover, appropriate use of human judgment over the use of force is required and use must be in accordance with all applicable domestic and international law, in particular, the law of war.²³

Distinctions are drawn among three types of weapons systems:

- Semi-autonomous weapon systems, which require human operator selection and authorization to engage specific targets (*e.g.*, human in-the-loop control)
- Human-supervised autonomous weapon systems, which allow human intervention and, if needed, termination of the engagement, with the exception of time-critical attacks on platforms or installations (*e.g.*, human on-the-loop control)

²³ Office of General Counsel, *Department of Defense Law of War Manual*. [June 12, 2015], p. 329. Available at www.dod.mil/dodgc/images/law_war_manual15.pdf (Accessed June 2016.)

The law of war does not specifically prohibit or restrict the use of autonomy to aid in the operation of weapons. In fact, in many cases, the use of autonomy could enhance the way law of war principles are implemented in military operations. For example, some munitions have homing functions that enable the user to strike military objectives with greater discrimination and less risk of incidental harm. As another example, some munitions have mechanisms to self-deactivate or to self-destruct, which helps reduce the risk they may pose generally to the civilian population or after the munitions have served their military purpose.

DoD Law of War Manual

Section 6.5.9.2. No Law of War Prohibition on the Use of Autonomy in Weapon Systems.

- Autonomous weapon systems, which upon activation can select and engage targets without human intervention

Extensive verifications, validation, test, and evaluation are required before fielding autonomous weapons systems. The use of automated regression testing of system software is also recommended, acknowledging the difficulty of full path regression testing in learning systems. The principles of military necessity, distinction, discrimination, and proportionality apply, according to the law of war.

In spite of the clarity provided in official documents, they alone are not sufficient for allaying public concerns about the use of autonomous weapons. Recent statements by prominent scientists and technologists are attempting to promulgate the notion of dire consequences due to the rapid adoption of artificial intelligence and autonomous robots.^{24,25,26} Some of the published statements are careful to use the qualifier fully autonomous weapons, but the distinction can be easily lost in communications. The potential for a backlash against the introduction of non-lethal and semi-autonomous systems for military use could grow.

An integrated approach is needed

Establishing—and maintaining—trust in autonomous systems requires a broad view of the entire lifecycle of the system. This principle is idealized in Figure 7 as a continuous process (purple arrows), which begins with experimentation and development of doctrine and CONOPs, followed by specification of operational requirements, and proceeds to system design, development, testing, training, operations, and maintenance. While all these functions are part of any normal process to

²⁴ R. Cellan-Jones, *Stephen Hawking warns artificial intelligence could end mankind*, BBC News. [December 2, 2014]. Available at www.bbc.com/news/technology-30290540 (Accessed June 2016.)

²⁵ Future of Life Institute, *Autonomous weapons: An open letter from AI & robotics researchers* [July 28, 2015]. Available at futureoflife.org/AI/open_letter_autonomous_weapons (Accessed June 2016.)

²⁶ International Human Rights Program at Harvard Law School, *Advancing the Debate on Killer Robots: 12 Key Arguments for a Preemptive Ban on Fully Autonomous Weapons* [May 2014]. Available at hrp.law.harvard.edu/wp-content/uploads/2014/05/Advancing-the-Debate_final.pdf (Accessed June 2016.)

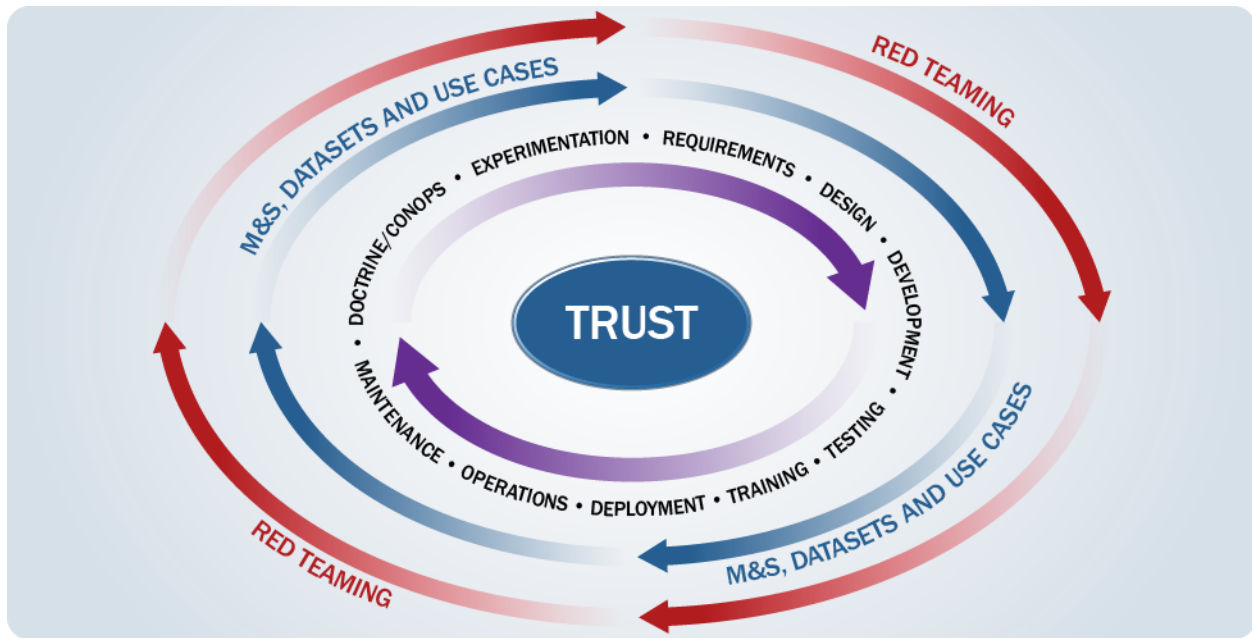


Figure 7 The lifecycle of the systems needs to establish an appropriate calibration of trust in autonomous systems.

field a new system, the distinction for an autonomous system is that the sequence is not linear, but a continuous process spanning the entire system lifecycle.

Because autonomous systems are sensor- and software-intensive, rapid upgrading of subsystems will be vital for maintaining military advantage. The full lifecycle process, as conceptualized in Figure 7, efficiently integrates upgraded subsystems into operational systems. Moreover, the learning that is inherent in autonomous system development and employment will couple into evolution of doctrine and CONOPs, modifications to the system, and additional testing. There may be many sub-loops in the realization of this process, but the concept of the continuous loop throughout the system lifecycle is of primary importance.

The surrounding circle of M&S (blue arrows) depicts the need for M&S integration throughout the cycle, from initial concept to operational test and evaluation and through operator training. A continuing and pervasive thread of M&S activities can support rapid evolution of system design and performance. This will require a transformation of the conventional model of developmental T&E and operational T&E from discrete segments of the acquisition cycle to an ongoing evaluation and evolution of the technology and concepts within the operational community. Military operators at all levels must become more familiar with and begin to employ T&E concepts and techniques within routine training operations. In addition, for any given development cycle, autonomous systems that learn can take advantage of datasets generated during previous development and training cycles to build on past successes—and learn from failures.

The outer circle for red teaming (red arrows) depicts the need for this in all stages of the cycle, from early concept development, to developmental and operational test and engineering, fielding,

and training. The rapidly advancing state-of-the-art in autonomy, together with the software-intensive and embedded processing characteristics of any autonomous system, whose functionality may evolve over time, introduce many more attack surfaces that will not all be discovered prior to deployment.

In conclusion, assuring appropriate calibration of trust assessments in the implementation of autonomy in DoD systems is clearly central to their adoption. The decision for DoD to deploy autonomous systems must be based both on trust that they will perform effectively in their intended use and that such use will not result in high-regret, unintended consequences. Without such trust, autonomous systems will not be adopted except in extreme cases such as missions that cannot otherwise be performed. Further, inappropriate calibration of trust assessments—whether over-trust or under-trust—during design, development, or operations will lead to misapplication of these systems. It is therefore important for DoD to focus on critical trust issues and the assurance of appropriate levels of trust.

3 Accelerating the Adoption of Autonomous Capabilities

Enabling transition to accelerate innovation encompasses many cross-cutting issues for the entire DoD enterprise. The study proposes an interdependent set of enterprise-wide recommendations that focus on critical enablers essential to accelerating innovation in the use of autonomy.

Tackling engineering, design, and acquisition challenges

Autonomous systems can be cyber-physical or totally cyber-dominated. In any case, these systems will be dominated by a software architecture and integrated software modules. The DoD historically has had difficulty in specifying, developing, testing, and evaluating software-dominated systems.

Autonomous systems present a new level of complexity for the acquisition community because of the potential for an evolving role of the human operator as part of a human-machine system. Additionally, the fact that some autonomous systems will be adaptive and have the capacity to learn from the environment in which they operate will require much higher levels of trust demanded by operational users, policy makers, and the public.

Prior to and during the requirements development process, a set of key decisions should be made as early as possible that will be critical to control cost, schedule, risk and vulnerabilities for the proposed autonomous system. Those decisions will require early definition of the functional modularity of the system that will allow for subsystem iteration based on technology advances and potential system learning. The functional architecture of the system must define the role of the human but also allow for an evolution of the functions performed by the human and those performed autonomously. Establishing the functional architecture early while allowing for system evolution need not be inconsistent objectives if the functional architecture is established at the appropriate level and the development process is flexible and can allow for the rapid insertion of new modules that allow the various subsystems of the overall autonomous system to be iterated independently.

Because autonomous systems are dominated by software and adaptive software architectures, cyber-security becomes even more imperative. Technologies that support cyber-security must be architected within the system from the outset, and they must be flexibly capable of being updated during the development process as new cyber-attack surfaces are detected. Cyber-security must also be integrated into the system as new capabilities are added since every new capability brings with it a new vulnerability.

Good systems performance seldom occurs in the absence of good, informed design. For autonomous systems, designers must anticipate a substantially broader range of operational circumstances. They must reconcile the fact that the system may be required to learn and adapt, so its specification is, to a degree, non-stationary. Experience has shown that the more constrained the design space, the more successful the acquisition will be (but the operational utility and system lifecycle may be limited). Conversely, the more fluid the requirements, the more likely it is that the

acquisition will be both less successful and costlier; poorly designed systems will be more expensive in a changing environment.

A consequential decision is how the requirement is parsed into subsystems. A modular system holds the promise that the overall system performance can be improved by changing out a module or two with little adverse effect on the other component modules, and without the need for full regression testing of the entire system. Ideally, the functional modules are sufficiently granular to allow system upgrades that represent significant advances in performance. A modular architecture can allow for slower-changing core functions and faster-changing peripheral functions. This describes good engineering systems design processes.

Autonomous systems allocate and share functions among humans and machines; this will influence functional modularity and decisions about system design. Users who generate requirements will often want functional allocation between humans and machines to be fluid with the expectation of enhanced autonomous performance as the system evolves. At the early stages of system requirements definition, the fundamental allocation and sharing of human-machine functions must be established, and the modular system design must allow for increased autonomous performance. These are critical design issues, which may increase initial system schedule and cost—while at the same time enhancing lifecycle effectiveness and reducing lifecycle costs. [An extreme—and costly—outcome might be to design the system to have the capability to be either fully human controlled or completely autonomous.] The use of modeling and simulation during the requirements development process will allow system designers to establish the functional modularity of the system and the evolution towards extended autonomous functionality while staying within cost and schedule.

Additionally, as early as possible, a comprehensive modeling and simulation capability, system use cases and the necessary data sets to allow for system design, modular architecture development, validation and verification of the system, and a continuous test and evaluation process must be defined, budgeted for, and implemented. Throughout the definition, requirements development, development, validation and verification, testing and evaluation, and operational fielding of autonomous systems, the capability for red teaming must be provided to address system vulnerabilities and establish system performance boundaries.

Engineering effective human-machine collaboration

Human-machine interaction is a highly specialized discipline with major branches for interacting with software (“autonomy at rest”) and with software-augmented hardware (“autonomy in motion”). It provides a framework for analyzing and assigning functions and tasks across teams of humans and machines, as well as providing guidelines for designing future systems with complex, adaptive, and dynamic modes of interaction with individual or human teams. The community includes theorists and practitioners from a number of areas, including human factors engineers, cognitive scientists, artificial intelligence researchers, and controls theorists, among others.

A major problem throughout the history of automation has been the human out-of-the-loop control problem, where a task or process encounters a problem and a human who is doing some

other task has to suddenly take control. The inability of a human to suddenly change mental gears and successfully diagnose a complex problem currently being addressed by the automation is well documented.

Alternatives to this are being developed, especially when exerting human control on an autonomous system is not an option. These include providing an explanation capability, both for the system to “explain” what it is doing and why, as well as a means for the human to “explain” what it means when a command may be vague or ambiguous, or when a situation may have changed from what had been anticipated. Complementing this explicit explanation capability is an implicit understanding of each other’s mental models (of the system by the human, and vice versa), in terms of goals held, situational awareness, decision plans, and so forth.

In addition, system self-awareness can provide for additional system robustness, through awareness of self-health and of the environment, and an understanding of where the system is operating with respect to its “envelope of competence” (or operating envelope). If the system includes learning, interactions with human operators would be facilitated if the system came with a design and training pedigree to help human teammates and supervisors anticipate novel system behaviors as the system evolves in reaction to past experiences and training.

Defending blue autonomy by red teaming

In the current context, red teaming is the relentless search and extirpation of vulnerabilities in one’s own systems. This single-mindedness is essential to the resilience of our autonomous systems. As an added benefit, red teaming may provide significant insight into vulnerabilities in the opposition’s autonomous systems, possibly unlocking an adversary’s autonomous system without direct access to the article.

The ingredients for successful red teaming include domain knowledge—an understanding of the mission and functions of the system, as well as the larger military purpose it was designed to serve. This means that a red team is intimately related, though still independent in thought, to the acquisition program or the operational element. It also means that red teaming expertise is distributed and often compartmented.

Because red teaming, in practice, is necessarily distributed, there is a benefit in establishing a community of interest (COI) to grow the art and practice while still acknowledging the need to know specific topics. Red teaming will be essential to all stages of the development of autonomous systems.

Recommendation 1.

USD(AT&L) should require that the following practices be developed and applied to all software dominated systems and, in particular, autonomous systems:

- Software designed to best engineering practices, and for incremental upgrades that can be implemented without full system regression testing
- Iterative development of subsystems

- Early decision on functional modularity to allow the system and the role of the human to evolve
 - A framework for analyzing and assigning functions and tasks across teams of humans and machines
 - Planning and budgeting for reliable datasets, use cases, modeling and simulation, validation and verification, testing and evaluation, user engagement, and self-monitoring at the earliest stages of autonomous system development
 - Red teaming at all stages of autonomous system development
-

Mitigating cyber issues

Autonomous functionality—basically decision-making—in a system resides in software replete with branching logic and tables of variables and parameters which, together, model the mission to be accomplished, the environment in which it must be executed, and the conditions that pertain. The more complex the mission and the more diverse the environment, the more extensive and complex is the software. Typically, too, autonomous systems will have organic sensors, a considerable corpus of stored information, and optional communication for some supervisory functions, along with a capability to receive and implement over-the-air updates. Insofar as they are mobile, they will implement precision, navigation, and timing (PNT) and collision avoidance. Additionally, there may be self-diagnostics and contingency fail-safe provisions.

Cyber intruders view such systems as target rich, rife with capabilities that also introduce vulnerabilities. Cyber defenders will worry about the extent of the cyber-attack surface, which might include “one touch” access at a number of stages in the lifecycle from the drawing board to the battlefield, remote network access, or entry points via the onboard sensor suite.

Exacerbating concerns about cyber complexity and vulnerability is the premier characteristic of autonomous systems: they have a wide range of consequential actions up to and, perhaps, including lethality, and may decide upon these actions with little or no human supervision. There may be provisions made for the machine to wrest control from the human when, as designed, it believes it knows best; is acting to protect the health and safety of humans or machines; or is given ambiguous, contradictory, or paradoxical instructions by its human supervisor. Curiously, when the machine assumes control under such circumstances, we say it is “out of control”—literally correct, but emotionally loaded.

In general, adversaries seek, and defenders worry about three rather different attack objectives—confidentiality, integrity, and availability. The first objective, breaching confidentiality or stealing secrets is a lesser concern here, although the autonomous system will admittedly be loaded with items of intelligence value: commander’s intent, mission orders, and target parameters, as well as elements of doctrine, CONOPs, and rules of engagement (ROE). The second and third objectives, availability and integrity, while less revealing, are nevertheless the essence of mission assurance—will the system be available when you need it and perform as designed and directed?

Corruption of these aspects can result in bad judgments, missed opportunities, inappropriate targeting, self-destruction, system abort and return to base, and so on.

The vulnerabilities that an adversary might exploit could result from poor design or implementation. They could have been introduced by an adversary who had even momentary access— physical or virtual—to the system. They also could stem from the design requirements: for example, the desire for fail-safe and return-to-base modes subject to compromise by an adversary.

Autonomous systems of the class under consideration may also have the capability to adapt and learn from experience—that is, adjust certain decision parameters like the estimated prior probabilities and presumed costs and values of potential outcomes. And, of course, machines will be expected to modify their decisions and actions in accordance with dynamic behavioral guidance. Such capabilities open the door to adversary manipulation. By presenting a misleading set of circumstances they might “mis-train” the system.

The bottom line is that capabilities almost always entail vulnerabilities, and autonomous systems will be quite capable. In addition, their software—indeed, their logic—will be quite complex, making it difficult to validate and verify actions and to root out induced vulnerabilities should the adversary have gained access.

The Department of Defense is working hard to respond to incessant attacks on its enterprise systems and those of its contractors (the defense industrial base), as well as to moderate the attendant publicity. Now it must consider how to ensure the viability of the autonomous systems it will field, both to guarantee mission success and to avoid further loss of public confidence.

There are important differences, to be sure, between the two situations. First, the enterprise systems attacks have breached confidentiality and secrets have been stolen. While regrettable, the damage seldom rises to that which would ensue if key autonomous military systems did not operate when and as directed. Second, the enterprise systems are heavily indebted to commercial off-the-shelf software and, to date, the marketplace has rewarded capabilities at the expense of vulnerabilities. Autonomous military systems are more likely to be under DoD control, even if commercial components are included, which provides a better chance to build in more cyber resiliency. Third, erosion of public confidence in autonomous systems—all the more critical if lethality could be involved—may seriously derail an otherwise desirable move toward autonomous systems.

These differences provide a challenge and an opportunity for DoD and the defense industrial base. A number of steps have been proposed to protect the integrity of autonomous systems throughout the entire lifecycle—drawing board to battlefield—including:

- Examining the requirements to eliminate unnecessary vulnerabilities
- Designing and building to the requirements using best practices
- Running an aggressive counterintelligence program to counter insider threats
- Ensuring adequate time, resources, and incentives for validation and verification
- Red teaming early and often
- Securing the supply chain

Operational considerations must anticipate the adversary and include:

- Prioritizing critical communication channels and integrity for C2 and over-the-air updates
- Planning for changing sources of reliable PNT
- Including sensors and processing to overcome spoofing
- Examining sensor inputs and providing reality checks

Both the National Security Agency (NSA) and the U.S. Cyber Command (USCYBERCOM) can play a pivotal role in finding and eliminating vulnerabilities, largely cyber, in the design and application of autonomous military systems. NSA has the preponderance of experience and firsthand intelligence access. The NSA Director adjudicates cyber offense-defense and, as Commander of USCYBERCOM, executes cyber operations. This experience base provides the necessary fundamentals for cyber red teaming.

Recommendation 2.

USD(AT&L) should address the special issues associated with cyber resiliency in autonomous systems, which include:

- Ensuring DoD-wide practices for cyber hygiene throughout the development process
- Mitigating initial introduction of vulnerabilities through requirements, design, and supply chain best practices
- Requiring system self-monitoring, redundancy, and baseline comparatives to identify spoofing, unauthorized access, or intrusion that alerts the operator
- Incorporating cyber issues into overall red teaming efforts
- Working with NSA and USCYBERCOM to harvest, integrate, and disseminate the science, distributed knowledge, expertise, and best practices of red teaming in the cyber domain

Creating new test, evaluation, modeling, and simulation paradigms

Autonomous systems present a number of challenges with regard to T&E, M&S, and related analysis. Most systems are continuously monitored by humans who can note deviations from desired performance and correct the behavior of the system. Autonomous systems may have periods of time with limited or no communication capability; during those periods the system must reliably behave in known ways to the full range of stimuli that the system is designed for. Insuring that the system will respond appropriately to all of the possible inputs will exceed the capability of conventional testing. It will require using a combination of modelling and simulation to explore thousands of test cases, statistically measuring system performance against the desired standard, then doing real world testing of the system to ensure that the modelled and real world behavior match for corner cases that span the range of system performance. This approach has been successfully used in commercial systems; the fidelity of the modelling and simulation must meet the requirements of the program and the modelling and simulation parameters must match the actual parameters of the system.

Maintaining accurate model parameters requires that attention be given to the process by which the parameters are chosen and changed. The Department also needs to ensure that it will have the full sources of all of the models and data available for its use. In some current DoD programs, models are generated in order to size the system and explore alternatives, but these models are not kept current with the system as design changes are made during development and are also not kept current throughout the life of the system. Using M&S and live testing early to evaluate the performance of the system with respect to the requirements allows for rapid cycles of model, build, test, and modify, to advance the design and performance of the system. The testing should include the expected and predicted adversary capability that may be used to thwart the system.

The act of modelling the system and field-testing will generate significant databases. This data provides deep insight into the capabilities of the system, and must be appropriately safeguarded. The data will be important in enhancing and upgrading the system, as well as in designing new systems. This valuable data must not be allowed to fall into adversary hands.

Test and evaluation for software that learns and adapts

DoD's current testing methods and processes are inadequate for testing software that learns and adapts. Because such software exhibits different behavior as it incorporates more data about its task, and learns to provide better results partly based on experience, such software cannot be exhaustively tested. There is no single, known correct answer. With experience, the software may discover unexpected, useful patterns that lead to better, different answers. Such software should be routinely evaluated as its behavior changes throughout its lifecycle. Consequently, testing should be conducted throughout the development and deployment processes. See Box 2 on page 31 for an example of this concept.

Commercial software developers have developed methods that permit software to be developed and tested in increments. After core functionality is developed, the software can be tested and put in the hands of users with increments of function tested and released incrementally. The lifecycle encompasses a long-lived sequence of incremental upgrades interspersed with test. This development approach delivers the software very early to users to gain early user feedback. As discussed earlier, the DoD has a rigid testing process—both for development test and operational test—which is in direct conflict with more advanced commercial development practices that are better suited for testing adaptive software. DoD's strong separation between developmental testing and operational testing is in conflict with the best known methods for managing the development of such software. Operators will have to change their mindset from expecting weapon systems that “just work” out of the box to systems that require their time and effort into shaping their ever-evolving instantiation, but will ultimately be better customized to their mission, style, and behaviors.

Box 2. Logistics as a Testbed for Software that Learns

Logistics planning and execution is a particularly good candidate application for T&E to experiment with new test methods for learning, adaptive software because the behavior of logistic software can be evaluated against crisply known metrics. Given inventory available at multiple locations to fill a customer requisition, there may be multiple logistics plans that are acceptable. Each of those plans may be more or less attractive considering dimensions beyond simply filling the customer's requisition, e.g., pallet packaging, container packing plans, urgency, and transportation constraints.

Core functionality of logistics software can be delivered as soon as it is developed and additional function can be delivered in later releases. An example later function would be anticipatory order filling based on knowledge of the customer's mission goals, priorities, and preferences or incorporating near real-time knowledge about the customer's situation, e.g., impending bad weather a week hence. A logistics system could then anticipate customer needs rather than simply filling requisitions.

A useful testbed will allow continual testing during development and then through the software system's lifecycle. Operators can be involved early to allow their evaluation of the logistics software, even with only core functionality, made known to the developers. The user interface is particularly important with adaptive software because the user needs to not only interact with the software, but to understand what it is doing, and occasionally to guide it or select among options that it proffers. Early user involvement will yield better interfaces, and inevitably, better functionality.

This method of testing that intersperses testing and code releases requires disciplined configuration control to permit frequent software updates, with routine, controlled rollback as needed when system behavior is not acceptable to the users. Testing can be aligned to match upgrade cycles to enable frequent incremental releases.

Such testing requires a realistic, synthetic environment that incorporates models of the environment, and even of human frailty when planning and executing logistics missions. The environment needs to be able to present scenarios to the software under test, and to evaluate the output of that software for each scenario. Classic regression testing might be used early, but the adaptive nature of the software under test requires that the testing infrastructure become increasingly knowledgeable as more functionality is added. An effective synthetic testing environment will collect accuracy and quality metrics and automatically track improvement (or lack thereof) over time.

Logistics is a good application for early experimentation because the behavior of the logistics system can, to a great extent, be evaluated rapidly using known and straightforward metrics. Likewise, given knowledge of inventory at hand, a requisition, and the plan produced by the software, expert logicians can readily judge the quality of that plan produced by the software.

With a successful logistics demonstration, lessons learned can be evaluated so these new test methods for software that learns can be applied to other applications that are more difficult to test and evaluate.

DoD will increasingly utilize software that learns and adapts for diverse applications. Such software incrementally enriches its database to describe relevant context, environment, threats, user inputs, and mission objectives. It records new input, then integrates and generalizes past experience to make decisions partly based on the accumulated data and experience. The software may record its reasoning and can report out how it derived intermediate decisions. With machine learning algorithms, more data is usually better, and the learning algorithms are expert at finding that data which is useful, and ignoring the irrelevant. For these reasons, it can be useful to transfer data from one instantiation of software application to another so that there is a broader base of data from which the software can learn. Additionally, reusing building blocks of previously evaluated and validated software modules will reduce the required time and resources for follow-on implementations of autonomous solutions.

Recommendation 3.

The Director of Operational Test and Evaluation (DOT&E), in conjunction with the Office of Developmental Test and Evaluation (DT&E), should establish a new T&E paradigm for testing software that learns and adapts. Considerations include:

- Opportunities to adopt or adapt commercial best practices in T&E of learning systems
- Experimentation and development of new methods and means for testing software that learns, such that the “correct answer” changes with context, experience, and new data

A more continuous and iterative approach to address validation and verification

As noted in previous sections, systems that learn or adapt present a significant challenge to the current T&E process because system behavior will depend upon the sequence of stimuli that the system receives. Research is required to develop approaches to test and evaluate the readiness of learning systems to be used by the warfighter. Humans serve as good models of learning systems that we have trust in. A similar approach to qualifying non-human autonomous systems may be used as a starting point for the research.

In non-learning autonomous systems, emergent behavior complicates the test and evaluation of the system. The emergent behavior will be manifested when groups of autonomous devices are presented with similar stimuli, causing them to act in a desired coordinated manner. Insuring that emergent behavior occurs in the desired ways and cannot be induced to occur in ways that would be detrimental to the system performance presents a significant T&E challenge. Research is required to understand the best ways to test for emergent behavior characteristics.

Recommendation 4.

The DoD test and evaluation community should establish a new paradigm for T&E of autonomous systems that encompasses the entire system lifecycle. Considerations include:

- Make extensive use of live and synthetic environments for evaluating and qualifying transition from development to fielded systems
- Establish standards and guidelines for continuous verification and validation (V&V) for autonomous systems
- Start testing early and iterate in development, operational testing, and fielding: “build, test, change, modify, test, change....”
- Develop datasets for assessing autonomous functionality, and expand based on live test validation results (experimentation can add to the datasets)
- Include expected adversary-induced environments, *e.g.*, cyber, electronic warfare, etc.
- Plan for involvement throughout the lifecycle of the system—the system will learn and self-modify, or be intentionally modified

Expanding the roles for modeling and simulation

Modeling and simulation are frequently used to explore the design space of a system, to predict the performance of the system, and to understand the limitations of design alternatives. Models typically are updated during the development process to better represent the system that will be built. These digital experiments are useful in understanding the design trade space of the system.

M&S is also used for experimentation and refinement of system requirements. For more sophisticated autonomous systems, M&S is needed to augment test and evaluation by simulating the testing environment and creating, running, and measuring thousands of random but realistic scenarios. By providing statistical evaluation criteria and behavior measurement, M&S has been shown to be effective in exploring the large scenario space in which a system is designed to operate. Thousands of test cases can be generated and simulated to increase confidence that the system will satisfy design requirements. Israeli Aerospace Industries has used the Gazebo simulation tool, which is built on the robot operating system (ROS), to perform such simulations and has shown that it is an effective means of performing a large number of digital experiments to gain a statistical measure of system behavior.

During the recent Robotics Challenge organized by the Defense Advanced Research and Projects Agency (DARPA), the use of M&S was further extended. By running the Gazebo simulator at the user workstation, the operator could watch a simulation of the robot’s performance that would predict the robot’s action when the communication channel was not able to provide real time updates, thus enabling operators to more effectively intervene when needed.

A simulator can also be run on the robot itself. This allows the robot to have an on-board monitor that can ensure that the robot stays within a set of operating rules. Incoming command sequences can first be run on the simulator to ensure that the outcome meets an established rule-set

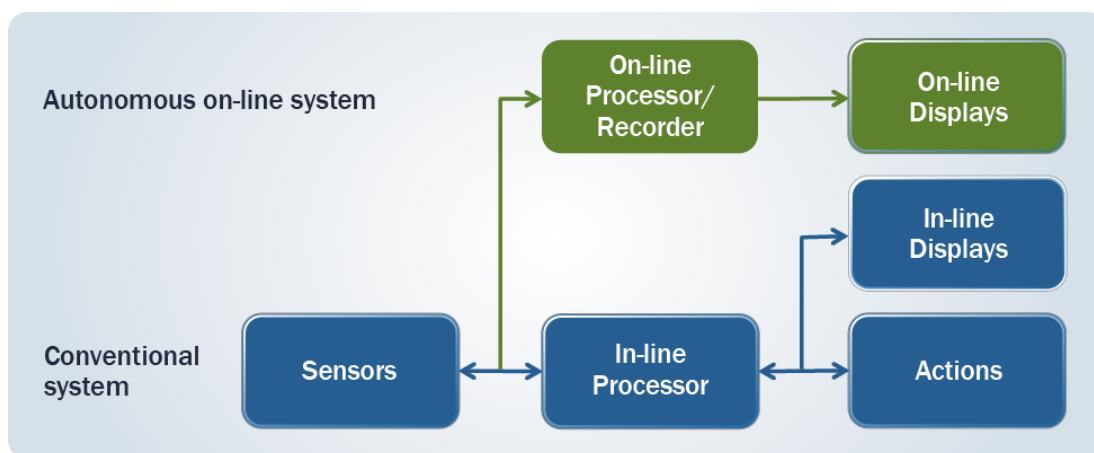


Figure 8 On-line processors implement new functions for system V&V and performance enhancement.

prior to execution by the robot. In addition to making sure that the robot does not perform an action that would cause undesired damage, this can help detect malicious commands. Finally, if the simulator determines that the robot will be forced into a situation where the machine will be compromised, it can force the erasure of key data and programs or take other actions to prevent the system from being reverse-engineered.

In a conventional system, the in-line control processor provides input to the in-line displays and system actions. For some systems, such as the Aegis Combat system, an additional processor has been added, termed an “on-line processor or recorder,” as shown in Figure 8. The on-line processors implement new functions and allow the operator to see the effect of the new function prior to it being implemented in the in-line system.

The approach of adding a second processor allows implementation of a number of new functions that are useful over the life of the system. In addition to demonstrating new functions, it can enable command validity checking and facilitate mission rehearsal. It can also provide a black box recording function to capture data for subsequent analysis of the inputs and decision sequence that the system implemented.

The secondary processor may also be useful in allowing explanatory dialogue of the robot actions and allow for new collaborative and cooperative behavior. Additional on-line functions may include advanced autonomy algorithms to provide local diagnostics, alternative courses of action, and explanatory dialogue.

Recommendation 5.

USD(AT&L) should require the acquisition community to establish and implement a consistent and comprehensive M&S strategy throughout the lifecycle of the system, including:

- Alternative system approaches and implementations
 - Design and prediction of system performance
 - Performance of test and evaluation
 - Support for manufacturing
 - Support for training
 - Support for operations and sustainment
 - Exposure of vulnerabilities at all stages from requirements development
 - Ensure the harmonization and coordination of existing models and simulations across the acquisition and training communities
-

Integrating technology insertion, doctrine, and concepts of operation

During recent conflicts, unexpected threats on the battlefield spurred efforts to pursue rapid-equipping initiatives. The goal was to deliver relevant technology to the battlefield as soon as practical. Among initiatives launched to expedite solutions was the creation of Joint IED Defeat Organization (JIEDDO), charged with finding solutions to serious improvised weapon threats.²⁷ These accelerated acquisitions often followed the “try a little, buy a little” development and proof-testing model. Some efforts were more fully integrated, with developers and users working together during all phases of the expedited procurement—from design and modeling, through simulation, prototyping, testing, CONOPs and TTPs (tactics, techniques, and procedures) development, and issuance to the field—all performed in an overlapping, parallel joint atmosphere. Ongoing operational feedback improved effectiveness as adversaries adjusted their own attack TTPs.

The Rapid Equipping Force Initiative provides solutions to unanticipated needs and is a useful example for expediting the development and fielding of military semi-autonomous and autonomous systems. As the DoD moves to peacetime acquisition processes, there is a need to create and institutionalize an accelerated acquisition system to cope with threats that are developing very quickly using globally available technology to conceive and build capabilities that are often inexpensive, broadly available, and potentially lethal.

Today, many more tools are available to develop, test, and assess progress, delivering constantly improving capability in modeling and simulation. The goal is to create an environment where all the developmental and operational players are involved from the start rather than sequentially scheduled into a linear program evaluation and review process. With this change, the developer sees better ways and means to deliver the stated requirement, and the user develops early CONOPs and TTPs, while seeing new potential capabilities and uses not previously conceived in

²⁷ JIEDDO is currently called the Joint Improvised-Threat Defeat Agency (JIDA).

the requirements document. This synergistic cooperation continues through prototyping and field test where new possible use cases will continue to emerge.

This presents opportunities throughout the conceptual design and development process to make tradeoffs in capability, performance, cost, and schedule. In the standard linear model, such opportunities are not generally visible until much later in the acquisition process. This parallel approach allows the operator or user to develop draft CONOPs and draft TTPs alongside product development processes.

This approach—joint development and operational conceptualization—argues that time from concept to initial fielding should be the primary performance metric, recognizing that the global market is rapidly introducing new dual use capability with countless applications, recreational as well as potentially nefarious.

Recommendation 6.

Military Service Chiefs should integrate technology insertion, doctrine, and CONOPs by:

- Ensuring early experimentation for autonomous systems includes multiple methods:
 - Modeling, simulation, and operational gaming
 - Prototype testing employing hardware and software modifications to existing systems to allow autonomy investigations
 - Developing unique prototypes that explore new capabilities not possible in systems designed to be controlled by humans
- Employing alternative sources throughout the experimentation process:
 - Use available commercial systems to allow early experimentation
 - Strongly consider Military Service Laboratories, federally funded research and development centers' (FFRDCs), and university affiliated research centers' (UARCs) participation in prototype development and red teaming
- Ensuring that field experimentation with developmental autonomous hardware and software informs employment doctrine and CONOPs:
 - Designing and fielding early test hardware examines difficult trades early in the systems effort
 - Accelerating system maturity validates technical and operational expectations

Developing an autonomy-literate workforce

American military forces, formerly equipped with largely electro-mechanical platforms, are now fielding systems that are dependent on software for combat effectiveness. This technology shift has placed a huge demand on education and training to provide qualified people across all aspects of the economy—a demand that is far from satisfied and is growing.

Skilled technical people are essential to develop and measure accelerated joint technical and operational development and later fielding. Such an undertaking is data-heavy in all phases, from design, through modeling, simulation, validation, verification, tech insertion, and operational concepts and tactics, techniques, and procedures.

With commercial development the technology leader, it is also a more effective competitor for talent. The military can train its own but is at a serious disadvantage to retain experience—talented operators, maintainers, supervisors, and technology leaders.

Programs need to be developed that formalize broad exchanges between government, military, and commercial enterprises for extended periods—closer to months rather than days—so that both government and commercial personnel can learn and understand emerging technologies and capabilities as well as the range of user concepts and applications.

Sustaining an autonomy literate workforce challenges both the commercial enterprise and the government. The demand for such talent spans the U.S. economy, across a growing number of industries. In addition to finding new and innovative ways to establish broad personnel exchanges across the commercial, government, and military divides, the military needs to identify and functionally manage the currently small but essential cohort of personnel who are experienced and knowledgeable about autonomy, both technical and operational.

Such necessary measures may include creating a military service career identifier, insuring their continued assignment in the autonomy field, categorizing autonomy trained personnel in the highest pro pay category, and offering significant re-enlistment bonuses and officer retention bonuses. These steps are necessary but insufficient measures that need early implementation.

It is also worth assessing all four Military Service recruiting screening processes, testing, and selection for identifying the right inductees with the requisite talent and capability for assignment to the autonomy function or related cyber career fields.

New approaches also need to be considered for the training of reserve forces, both National Guard and the Reserve, who are currently assigned to units accomplishing autonomy-related missions. Some already work in commercially related enterprises as their day job. Why not assign them annual training to commercial enterprises in the autonomy domain? Or, take advantage of Reserve personnel who work in the intelligent system industry? Such measures could create a sustaining link that benefits both parties—the developer and the user—and provides focused quality training to the Reserve component.

Finally, the DoD should develop a formal program to regularly draw key members of relevant commercial autonomy enterprises to serve in Defense Agency and Military Service postings, especially at the operational level where they should be maximally exposed to field operations and exercises in order to understand military employment concepts and future needs. All available incentives should be used to attract such personnel, to include broad use of the Intergovernmental Personnel Act (IPA) process to compensate at rates as competitive as possible.

Recommendation 7.

The Undersecretary of Defense for Personnel and Readiness (USD(P&R)), working with USD(AT&L) and Military Service Chiefs, should develop an autonomy-literate workforce by:

- Establishing standing relationships with the commercial sector to be technologically aware:
 - Assign DoD civilian and military personnel to work within commercial sector on temporary rotating assignments
 - Ensure some National Guard and Reserve personnel do their training within commercial industry
 - Ensure National Guard and Reserve personnel who work in the intelligent systems industries can use their experience during military assignments
 - Immerse commercial technical personnel in operational field environments and exercises
- Establishing Military Service career identifiers for key specialists and operational employment experts:
 - Align recruiting tests and qualifications with finding the best candidates to employ autonomous systems
 - Implement performance excellence recognition (pro pay, etc.) programs with persuasive retention incentives

Improving technology discovery

As covered in the introduction to this study, autonomous systems and enabling technologies comprise a growing, global commercial enterprise. Furthermore, R&D to advance the state of the art is also occurring on a global scale. As a result, DoD will need to be exquisitely aware of capabilities in the commercial sector and the possibilities of how they might be used outright, or modified or adapted for military use.

Part of the approach for achieving and maintaining global awareness is already well established in DoD. Traditional “tech watch” programs, such as the Office of Technical Intelligence (OTI) within the Office of the Assistant Secretary of Defense for Research and Engineering (ASD(R&E)), and the Military Services’ 6.1 programs, such as the Office of Naval Research (ONR) Global and Air Force Office of Scientific Research (AFOSR) International, have been part of the Department’s R&D community for many years. In addition, improved techniques for horizon scanning to improve the ability to anticipate advances is being pursued, such as efforts in data analytics and cluster analysis. The current level of effort in these organizations related to autonomy may or may not be sufficient to support a Department push to introduce new capabilities into the forces.

While tech watch programs are necessary, they will not be sufficient for timely understanding of the potential impact of advances, whether to more rapidly field capabilities or that might pose a serious threat to our mission success. The fast pace at which autonomy is moving, both technically

and commercially, coupled with limited governance on the use of autonomous systems, is leading to many surprising and novel applications, as well as honing operator skills outside of DoD.

One effective way to compete in this rapidly evolving environment is for DoD to create and nurture a community charged to build on the findings of technology watch programs to prototype and experiment with commercially available technologies and systems in a military operational context. Such a community would not only explore new uses for autonomy, counter-autonomy, and countering potential adversary autonomy, but also more realistically inform what the tactical advantages and vulnerabilities would be to both the U.S. and adversaries in adopting or adapting commercially available technology. Such a program could also create options for insertion into current capabilities that might initially be too risky or too disruptive. Some of the best sources for participation in this effort are the government laboratories, and independent, not-for-profit laboratories, including the FFRDCs and UARCs, because of both the technical proficiencies of their workforces and their working knowledge of national security missions. Some of these individuals and organizations may have bridges to the commercial and academic sectors that can aid government programs.

Another way to expand beyond tech watch, and more importantly, to advance the quality and skills of the workforce is to arrange for sabbaticals for both civilian and uniformed scientists and engineers in commercial and laboratory organizations worldwide that are engaged in autonomy-related R&D. An important corollary effort should be the aggressive recruitment of IPAs from a wide array of outside R&D organizations to perform R&D in DoD.

A final ingredient for success in harvesting the fruits of a robust awareness program is a partnership with the intelligence community (IC). The team's results can provide cues for the IC on what to look for, while enhanced technical collection by the IC can better focus efforts.

Recommendation 8.

The Assistant Secretary of Defense for Research and Engineering (ASD(R&E)) should improve global autonomy technology discovery by:

- Enabling civilian and uniformed scientists and technologists to do sabbaticals in R&D organizations in the U.S. and abroad to understand and practice the state of the art in autonomy
- Attracting more IPAs to perform autonomy R&D in DoD
- Engaging government lab, FFRDC, and UARC personnel to establish a robust program of autonomy technology demonstrators using globally available technologies to:
 - Explore new uses for autonomy, counter-autonomy, and countering potential adversary autonomy
 - Investigate the tactical advantage and potential vulnerabilities of both blue and red
- Coordinating with the IC to expand technical intelligence regarding global advances in autonomy with military relevance.

Improving DoD governance for autonomous systems

This section explores the techniques for the management of autonomous systems with the desire to establish whether management changes would assist with the acceptance of systems employing this rapidly changing technology. Explored here are different and nonstandard techniques successfully employed in analogous situations.

Experience indicates that nonstandard techniques should be employed only when the imperative of operational exploitation outweighs the disadvantages of establishing additional organizational structures. Exemplars such as the exploitation of low observables technology, cruise missiles, remotely piloted vehicles, and, more recently, electronic warfare demonstrate the imperative for use of nonstandard techniques. The discussion of imperatives for nonstandard governance for autonomous systems includes:

- Acceleration of development and fielding of autonomous systems—including complex systems that include embedded autonomous functionality
- Stay current with the speed, diversity, and global nature of relevant technology development
- Countering potential adversary use of autonomy

Over recent history alternative methods of acquisition have been employed that could be helpful in accelerating deployment of autonomous systems to the operating force. One management strategy was employed by JIEDDO to counter enemy use of improvised weapons in Iraq and Afghanistan. This approach employed a large permanent workforce with intensive management to ensure dedicated focus on the required developments. The centralized office controls finance and directs development and fielding.

This approach has also ensured the rapid introduction of cruise missiles and remotely piloted vehicles. The management of autonomous systems could similarly be assigned to an existing DoD organization, such as the JIEDDO, or could serve as the mission for a new organization. Generally, such direction is exercised through Military Service mechanisms to ensure transition to requisite Military Services at the appropriate time in the development process. While this process provides the best assurance of acceleration, successful transition to Military Service management has proven difficult.

The recommended approach is to establish an executive committee (EXCOM). This is a coordination mechanism comprised of Defense Agency and Military Service principals acting in their official capacities. The advantage of this approach is that the standard acquisition approach is accelerated by the intensive focus provided by the principles working through the EXCOM. Concentration of senior management on resolving financial, acquisition, and policy issues is an important benefit of this approach. This approach was successfully applied to the development and exploration of stealth and counter stealth, and is currently being employed for the electronic warfare community.

Recommendation 9.

The Deputy Secretary of Defense should establish departmental governance of autonomy by:

- Creating an EXCOM with the responsibility to oversee and ensure the development and fielding of autonomous systems:
 - Ensure funding is available to Military Services for near-term autonomous systems demonstrations and system acquisitions
 - Ensure an integrated approach to building trust (recommended by this study)
- Tasking Military Services to:
 - Establish advocates (*e.g.*, Deputy Under Secretary of the Navy for Unmanned Systems) for resourcing autonomous programs and move them into development and acquisition cycles
 - Charter program managers for each new system to implement all aspects of the recommended processes, *e.g.*, coordinated M&S, red teaming, use cases
 - Ensure the operators cover all of doctrine, organization, training, materiel, leadership and education, personnel, and facilities (DOTMLPF) for rapid adoption of autonomous systems

One of the most important aspects of accelerating the adoption of autonomy is shifting the underlying policy, legal, and cultural framework. An effective governance structure must coordinate efforts to implement these changes.

Recommendation 10.

USD(AT&L), the Under Secretary of Defense for Policy, and the Assistant Secretary of Defense for Public Affairs should take a proactive, two-pronged approach to anticipate cultural objections to the use of autonomy. In particular, they should instruct and require that each autonomous program:

- Establish and refine a communications plan that provides transparency into the trust building measures (*e.g.*, safety and security systems, information assurance, anti-tamper, audit trails) undertaken from the start of the development of every autonomous system
- Routinely engage the public to build confidence that the Department is acting in accordance with applicable treaties and the Department's policies

As a starting point for effective public engagement, promulgate the lessons learned from the Army's programs on chemical demilitarization and assembled chemical weapons alternatives. Together the two programs offer numerous guidelines for both what works and what does not in engaging the public interest and understanding.

Countering adversary use of autonomy

As has become clear in the course of the study, the technology to enable autonomy is largely available anywhere in the world and can—both at rest and in motion—provide significant advantage in many areas of military operations. Thus, it should not be a surprise when adversaries employ autonomy against U.S. forces. Preparing now for this inevitable adversary use of autonomy is imperative.

This situation is similar to the potential adversary use of cyber and electronic warfare. For years, it has been clear that certain countries could, and most likely would, develop the technology and expertise to use cyber and electronic warfare against U.S. forces. Yet most of the U.S. effort focused on developing offensive cyber capabilities without commensurate attention to hardening U.S. systems against attacks from others.²⁸ Unfortunately, in both domains, that neglect has resulted in DoD spending large sums of money today to “patch” systems against potential attacks. The U.S. must heed the lessons from these two experiences and deal with adversary use of autonomy now.

While many policy and political issues surround U.S. use of autonomy, it is certainly likely that many potential adversaries will have less restrictive policies and CONOPs governing their own use of autonomy, particularly in the employment of lethal autonomy. Thus, expecting a mirror image of U.S. employment of autonomy will not fully capture the adversary potential.

The potential exploitations the U.S. could face include low observability throughout the entire spectrum from sound to visual light, the ability to swarm with large numbers of low-cost vehicles to overwhelm sensors and exhaust the supply of effectors, and maintaining both endurance and persistence through autonomous or remotely piloted vehicles.

Autonomy also inherently provides a greater surface of vulnerabilities and opportunities that may enable countering these advantages. Using deception to confound rules-based logic, overloading the processing capabilities embedded in a vehicle swarm, or disrupting the adversary’s supply chain all provide opportunities to limit or defeat the use of autonomy against U.S. forces.

Despite understanding that autonomy used against U.S. forces provides both a threat and an opportunity, DoD capabilities and knowledge in this area are fragmented, often compartmented, and provide little opportunity to benefit from both offensive and defensive technologies, techniques and programs. What needs to be done is better integrate these activities; share the knowledge gained from both sides of the offense–defense paradigm; and build a “ladder” for red teaming, with each rung informed by what has been learned on one side prior to the exercise, thus providing new knowledge and capabilities for the other side on the next rung of learning. Integration of both red and blue use of autonomy will thus help shape both U.S. offensive and defensive initiatives and responses.

²⁸ Defense Science Board, *21st Century Military Operations in a Complex Electromagnetic Spectrum* [2014]. Available at www.acq.osd.mil/dsb/reports/DSB_SS13--EW_Study.pdf (Accessed June 2016.)



Figure 9 Both inexpensive systems, such as the Flight Red Dragon Quadcopter (left), and more expensive ones, such as the Haiyan UUV (right), are becoming more capable and more available.

The U. S. will face a wide spectrum of threats with varying kinds of autonomous capabilities across every physical domain—land, sea, undersea, air, and space—and in the virtual domain of cyberspace as well.

Figure 9 (photo on left) is a small rotary-wing drone sold on the Alibaba web site for \$400.²⁹ The drone is made of carbon fiber; uses both GPS and inertial navigation; has autonomous flight control; and provides full motion video, a thermal sensor, and sonar ranging. It is advertised to carry a 1 kg payload with 18 minutes endurance.

Figure 9 (photo on right) shows a much higher end application of autonomy, a UUV currently being used by China. Named the Haiyan, in its current configuration it can carry a multiple sensor payload, cruise up to 7 kilometers per hour (4 knots), range to 1,000 kilometers, reach a depth of 1,000 meters, and endure for 30 days.³⁰ Undersea testing was initiated in mid-2014. The unit can carry multiple sensors and be outfitted to serve a wide variety of missions, from anti-submarine surveillance, to anti-surface warfare, underwater patrol, and mine sweeping. The combat potential and applications are clear.

Figure 10 shows a variety of small UA characterized by their gross takeoff weight and the weight of the payloads they can carry. They lie on a line close to a 45 degree slope, meaning that a vehicle of x pounds can carry a payload of an equal weight. The Airborg H6-1500 (pictured) follows this trend, shown in the highlighted triangle. This is a more robust 1500mm hex rotor UA.³¹ This vehicle has 6 26” carbon fiber propellers; an estimated flying time of 2 hours, at a maximum velocity of 40 mph, with a maximum payload of 9 kg (20 lbs); a maximum range of 160 km (100 miles); and can operate in wind/gust conditions up to 35 mph.

²⁹ iFlight Red DragonFly Quadcopter. Available at www.alibaba.com/product-detail/iFlight-Red-DragonFly-FPV-Quadcopter-Quadcopter_60020379201.html (Accessed January 2016.)

³⁰ *China tests long-range unmanned mini sub* [June 29, 2014]. Available at www.china.org.cn/china/2014-06/29/content_32804788.htm (Accessed June 2016.)

³¹ Top Flight Technologies, Malden MA. Available at www.tflighttech.com/products.html (Accessed June 2016.)

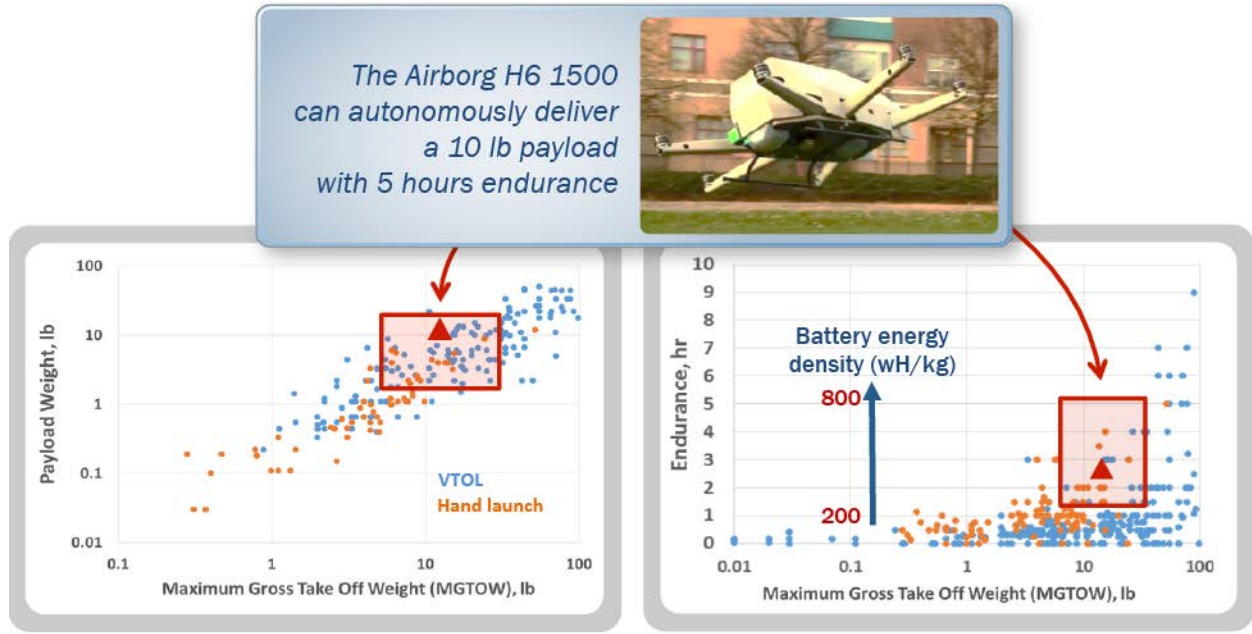


Figure 10 The Airborg (center top) capabilities are shown in the red boxes. The maximum gross take off weight of UA is compared with payload (left) and endurance (right).

Recommendation 11.

The Deputy Secretary of Defense should take immediate action to counter adversary autonomy, including:

- Direct USD(AT&L) to establish a counter autonomous systems community to develop and test counter-autonomy technologies, surrogates, and solutions:
 - Conceive, fund, develop, experiment, and demonstrate approaches
 - Develop adversary autonomous systems as “test targets”
 - Include full-spectrum expertise from cyber to directed energy to electronic warfare, cross-Service participation, and standing participation from the intelligence community
 - Create a standing countering adversary autonomy red team³²
- Direct the Undersecretary of Defense for Intelligence (USD(I)) to raise the priority of collection and analysis of foreign autonomous systems:
 - Technologies important to advancing autonomy capability are global in nature and commercially available
 - Actively maintaining global awareness of emerging advances in technical capability and field application must be understood and taken into account
- Direct the Military Service Chiefs to:
 - Equip Military Service opposition forces in training exercises with autonomous systems and counter-autonomy capabilities
 - Use lessons learned to support CONOPs, doctrine, and training

³² A useful model for this is the Low Observable and Counter Low Observable (LO/CLO) Executive Committee.

4 Strengthening Operational Pull for Autonomy

Commercial interest is exploding for autonomy, from widespread commercial development of UA for civilian applications, to self-driving and driver-assist applications for automotive applications, to dynamic spectrum management for cell phones, to IBM's Watson technology providing decision support to human operators in a wide range of big data applications.

Autonomy delivers operational value across a diverse array of vital DoD missions. The study established a categorization for the ways that autonomy can benefit DoD missions:

- Required decision speed – More autonomy is valuable when decisions must be made quickly (*e.g.*, cyber operations and missile defense).
- Heterogeneity and volume of data – More autonomy is valuable with high volume data and variety of data types (*e.g.*, imagery; intelligence data analysis; intelligence, surveillance, reconnaissance (ISR) data integration).
- Quality of data links – More autonomy is valuable when communication is intermittent (*e.g.*, times of contested communications, unmanned undersea operations).
- Complexity of action – More autonomy is valuable when activity is multimodal (*e.g.*, an air operations center, multi-mission operations).
- Danger of mission – More autonomy can reduce the number of warfighters in harm's way (*e.g.*, in contested operations; chemical, biological, radiological, or nuclear attack cleanup).
- Persistence and endurance – More autonomy can increase mission duration (*e.g.*, enabling unmanned vehicles, persistent surveillance).

As has been the case in a number of other technologies, most notably information technology and the Internet, where the DoD was at one point the driving force behind technology development, much of the leading research in autonomy is happening outside the Department and, in some cases, outside the United States. A key objective of this study was to identify opportunities for DoD to more rapidly exploit ongoing advances. By selecting several demonstrations of autonomous systems with near-term benefits, the study hopes to both demonstrate the operational value of autonomy, **while simultaneously strengthening the enterprise business practices recommended in this report that will make or break the transition of new technologies from the lab to the battlefield.**

Because the DoD mission is so broad, it was beyond the scope of this study to conduct an exhaustive review of, and search for, all of the beneficial roles for autonomy. Rather, the study chose to select representative system and mission applications to illustrate the potential value of autonomy. The study investigated four areas in depth—protection, battlespace awareness, force application, and logistics. These are joint capability areas that could immediately adopt existing autonomous technologies.

Within these capability areas, the study evaluated potential experiments against the relative benefits of autonomy. Descriptions of ten representative demonstration projects are interspersed in the following sections to show where focused prototyping and experimentation can validate advanced CONOPs and demonstrate the benefits associated with more aggressively adopting autonomy technology. In Figure 11, these 10 projects are shown plotted against the six benefit areas listed above. These specific efforts have the potential to either present significant challenges to an adversary—by costing more to counter than it costs the U.S. to deploy—or to negate an adversary’s transformative capability.

Each of the following 10 projects could be started immediately and is predicted to yield wide-ranging impacts. While these specific demonstrations are strongly recommended, the list is not intended to be exhaustive. Other demonstrations could deliver the same—or additional—benefits.

Autonomy for battlespace awareness

The evolving national security landscape places an increasing premium on the Department’s ability to develop and sustain situational awareness. Across the globe, threats are increasingly diffuse and growing, decision cycles—especially for cyber-threats—are dramatically shortened, and resource constraints will continue to limit the ability to cover every scenario with equal vigor.

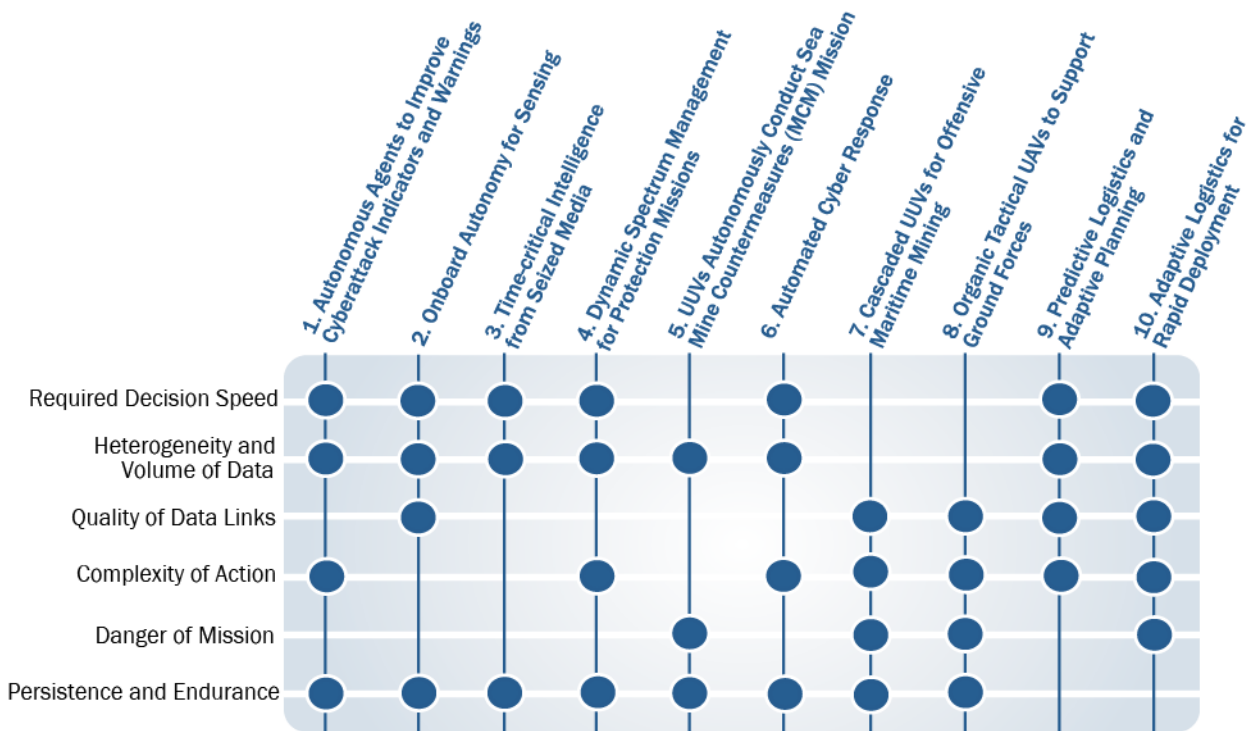


Figure 11 The study evaluated many candidate projects and selected those that encompassed the range of benefits of autonomy.

Providing battlespace awareness to the warfighter in potentially far-flung and congested battlefields is an increasingly complex problem, with a solution that is increasingly multifaceted. It will require better sensors and better organization, as well as more robust communications capabilities and data security. It will perhaps require a new tool: that of the capability for increased autonomy. Autonomy has the potential to enhance the ability to do more, at speed, at the sensor platform, which would then facilitate rapid re-direction of sensors; reduced transmission of data; enhanced sensor functionality; and continuing mission relevance in the absence of ground communications or a communications denied environment.

Autonomy can allow the commander to understand the real-time battlespace using data and techniques in ways that humans cannot by addressing volume, complexity, speed, and continuity. It can also assist in the merging of open source data with classified data sources.

Autonomous tools can help commanders integrate sparse, disparate, and unformatted data sources. Autonomy allows the human to detect and resolve inconsistencies deriving from any source by making information harder to hide and harder to spoof. At the same time, a typical adversary will be motivated to modify situational awareness capability to their advantage. Autonomy will enable intelligent sensing by helping sensors utilize their own data to collect better data through reduction of uncertainty, energy use, and communication bandwidth, as well as to address dynamic sensor configuration to support mission-relevant collection that is tied to commander's critical information requirements.

Battlespace awareness has typically been associated with the sensing part of the sense-think-decide-act process, but the rapid advance of technology may begin to blur these distinctions as battlespace awareness platforms must necessarily become more intelligent and more mobile. Over centuries of warfare, technology has inexorably expanded to permeate the various steps of sensing-thinking and deciding-acting, with emphasis first on the actuation of systems in battle. With the rise of electronics, sensing has extended far beyond the capability of human operators. Today, computers and networks are increasingly integrated with sensors, creating complex systems of systems of battlespace awareness capabilities.

Evolving support for conflict

Over the past decade, a generation of new sensors and decision support tools has been added to the Department's arsenal to address the growing complexity, ambiguity, and velocity of conflict. The ISR and related counter-IED systems employed in Iraq and Afghanistan, such as Constant Hawk, Gorgon Stare, and Argus provide wide area persistent surveillance. Similarly, a new class of sensors operates in cyber-space to enable the Department to assimilate threat data from forward deployed intelligence systems. These feed boundary defense devices that detect and counter cyber-threats at scope and scale in real- to near-real-time.

While these advances are impressive, these systems continue to require a vast logistics pipeline and a significant footprint of human talent to collate, cross-walk, and groom the sensor data in order to feed the downstream systems that generate responsive courses of action. Current systems

continue to strain to fully leverage the data volume, data velocity, data variety, and data veracity, which continue to outstrip human capacity to process and understand. The ability to extend the autonomy of these tools and techniques for the DoD could have a significant impact when applied to battlespace awareness in physical space as well as cyber-space.

The study concluded that appropriate employment of autonomy across the DoD military enterprise will yield a wide range of improvements in performance. Many commercial companies have developed, either for their own use or for sale, products that provide situational awareness and that incorporate various levels of autonomy. The Department is capitalizing on these where its operations are similar to those of commercial enterprises, but additional opportunities exist.

Project #1: Autonomous agents to improve cyber-attack indicators and warnings

Cyber-threat actors still enjoy a significant advantage over defenders in their ability to mount and sustain attacks using the natural camouflage that derives from the fact that any one defender can only observe a small component of the overall stream of adversary actions. This is especially acute as defenders attempt to detect threat streams that cross multiple networks, jurisdictions, and areas of responsibility.

The DoD very often cannot share information in a timely way because doing so would compromise classified sources and methods. At the same time, DoD frequently fails to capture insights available in open sources. While there is broad recognition of the imperative to share information across sectors and organizations to “connect the dots” and more easily reveal these threat actors, a particularly difficult challenge remains in the marriage of classified and unclassified data sources to feed a common operational picture. Much threat relevant data is held by the private sector, while the U.S. government generates and holds unique and valuable information in the form of insights gained from classified sources on both threat actors and activities.

Recent successes with cutting-edge tools for data analytics, such as conditional random fields for scene classification, and the ability of data analytics to synthesize and derive near real-time insight from large and complex volumes of data, offers the opportunity to employ these same techniques to the extraction of insights, tips, and queues from disparate data across classified and unclassified sources.

DoD alone enjoys permissive access to classified sources while at the same time enjoying access to increasingly robust open source data obtained through its own experience and authorities, or through collaborative or customer-provider relationships with the private sector. A demonstration is proposed to show the usefulness of this approach without compromising classified material.

An autonomous agent could examine open source data in multiple formats based on a cue from a classified product. Autonomy will allow the agent to aggregate information from multiple sources while obfuscating the search. The result will be sharable information that will not point back to classified sources and methods.

A successful program would add value to existing data from all sources. Most importantly, this process could enable cooperation among groups with different classification accesses.

Recommendation 12.

NSA, in partnership with DARPA and the Intelligence Advanced Research and Projects Agency (IARPA), should fully develop the means to tip and cue the Defense Information Systems Agency (DISA) and the defense industrial base to defend the DoD information infrastructure, extending to U.S. government and private sector support as appropriate. The study recommends an allocation of \$50 million over three years to pursue an aggressive goal to develop a working prototype comprised of the following:

- Take lessons learned from nascent efforts underway since 2007 and consider the threat-paced time urgency for fully mature defensive systems³³
- Leverage commercial and intelligence community tools and datasets to develop a continuous assessment of network conditions, threat vectors, and anomalous behaviors, and a rapidly configurable toolkit to provide fine-grained intelligence in support of threat identification, attribution, and tipping of operational cyber-defense forces
- Extract and integrate information from multiple and dynamic sources, and obfuscate search using dummy queries and meta-queries
- Rapidly sanitize information for sharing and dissemination to supported customers, the private sector, and allies

Project #2: Onboard autonomy for sensing

Counter-terrorism, time-critical targeting, and urban operations, among other missions, require wider field-of-view sensing with higher resolution and frame rates. A new class of such sensors is reaching the battlefield, including Constant Hawk, Gorgon Stare, and Argus. As users gain access to high-definition full-motion video, they are becoming dissatisfied with low-resolution images. However, comprehensive transmission of complete high-resolution collection, even in uncontested environments, is not feasible on the foreseeable communications infrastructure. Communication networks are already overburdened and are a key vulnerability in contested environments.

Sensor technology for ISR is rapidly expanding in terms of both resolution and coverage area. The pace of this growth has greatly outpaced our ability to communicate raw sensor data back to ground stations for processing and analysis, and even is outpacing processor capability growth as widely characterized by Moore's Law (Figure 12). This is in part due to enhancements in focal plane array technology.

Operationally, video analysts still manually review ISR video where the relevant information content to data ratio is quite low. Data fusion and analysis software is maturing. Dynamic time

³³ Such a system is Tutelage, which tips and queues SIGINT and collateral information on cyber-threats to sensors protecting the Defense Industrial base and DoD-managed networks.

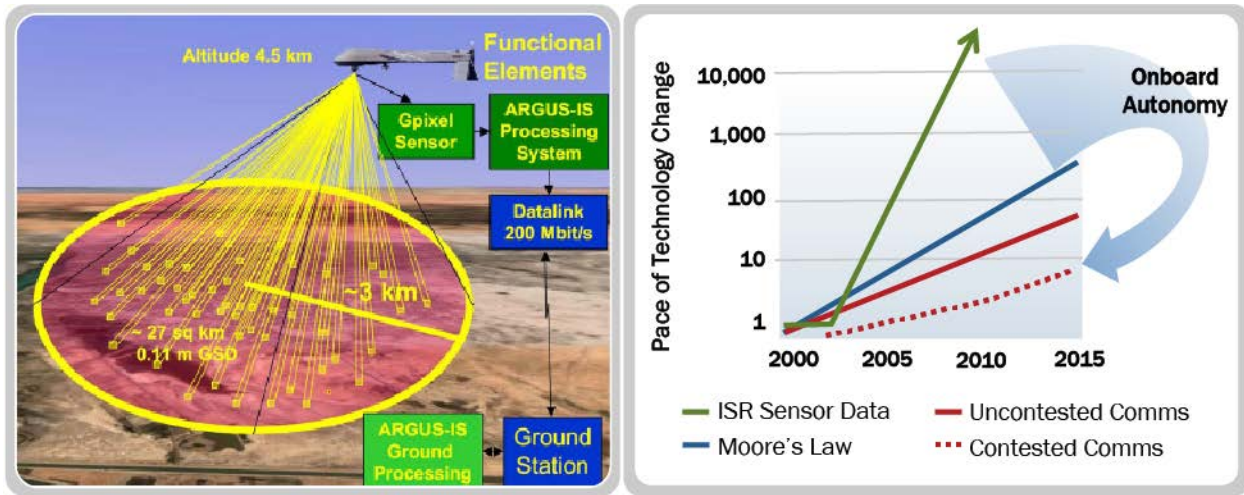


Figure 12 Elements of the ARGUS-IS Wide Area sensor are shown (left), along with the pace of technology change in sensor capabilities that can enable onboard autonomy (right).

critical warfighting capability has been developed by the Air Force and reached technology readiness level (TRL) 7 in 2009, with a focus on ground-based multi-source fusion for time critical targets. This capability was able to demonstrate greater than 90 percent success in automatically detecting moving humans and vehicles.³⁴ The Air Force Research Laboratory (AFRL) has an ongoing program called Planning & Direction, Collection, Processing & Exploitation, Analysis & Production, and Dissemination Experimentation (PCPAD-X) for ground-based fusion as well. DARPA is finishing the Insight program with a similar focus on ground-based multi-source ISR sensor fusion.

Some unattended ground sensors and undersea systems already use onboard, autonomous sensor processing. Technology advances in high-throughput, embedded processing, and machine learning offer the promise of onboard processing of high-resolution, multi-source airborne ISR sensor data. Autonomous sensor processing and high-level information generation would greatly reduce the required communications bandwidth and reduce the burden on human analysts. It also provides higher quality, improved persistence, better resilience, better tasking, and higher reliability. Perhaps the greatest operational benefit is agility: a single platform can adjust collection in real time based on observation, and many such platforms can coordinate to achieve better theater coverage. This would dramatically improve targeting information.

The application of machine learning enables much of this autonomous behavior. While the data needed for real-time analysis and cueing may be greatly reduced in this scenario, it is important to capture and retain the datasets to calibrate and retrain autonomous systems. Autonomous sensor data screening and fusion software will have increasing levels of complexity and will require different models for learning. They can be categorized as follows:

³⁴ A. J. Newman and G.E. Mitzel, *Upstream Data Fusion: History, Technical Overview, and Applications to Critical Challenges* [DARPA MTO Industry Day, 2015].

- Level 1—Object Refinement: A consolidated estimate of the observed and observable objects (*e.g.*, vehicles, facilities, persons) in the battlespace, and their kinematic state, representing the combined information from all sensors and information sources
- Level 2—Situation Refinement: An estimate of the militarily relevant entities (*e.g.*, units, groups, events) in the battlespace, their status, and relationships among the entities and between entities and observable objects
- Level 3—Threat Refinement: An estimate of the threat posed by entities in the estimated situation, including intent relevant to blue force plans, projected range of actions, and potential impact on those plans

Recommendation 13.

DARPA, working with AFRL and the 711th Human Performance Wing, should initiate a new program to adapt existing ISR data screening and fusion tools, such as the Air Force’s Dynamic Time Critical Warfighting Capability (DTCWC) or PCPAD-X, or DARPA’s Insight for autonomous, real-time use. The estimated cost for this effort is \$80 million over three years. Some suggested implementation steps include:

- Integrate software into an embedded processing payload on a BizJet-class UA platform to autonomously prioritize and process extracted sensor data and transmit mission-relevant information over communication channel capacities that would be available in a contested military environment
- Demonstrate a likelihood greater than 90 percent of including target information in the real-time communications stream for the autonomous screening and processing algorithms as needed to achieve operational relevance
- Participate for demonstration purposes in a Red Flag exercise of time-critical detection and tracking of surface-to-air missiles in a denied environment

Project #3: Time-critical intelligence from seized media

Special operations forces and Military Service tactical document and media exploitation (TACDOMEX) teams routinely seize massive quantities (terabytes) of diverse data types on digital media (*e.g.*, computers, tablets, smart phones) from adversaries. For time-critical, counter-terrorism operations, these media can provide valuable intelligence on people, places, and organizational structures— if exploited on operationally relevant timelines, meaning hours or days rather than weeks or months. Current document and media exploitation (DOMEX) operations focus on extracting information only from human-searchable files such as text and metadata. However, images, video, and audio can provide additional valuable information that can automatically be extracted by commercial tools like image analysis, translation, summarization systems, email network analysis, scanning and word recognition, and speech analysis. This in turn can be used as input to a new tool that constructs, for example, a social network graph and with node annotations, *e.g.*,

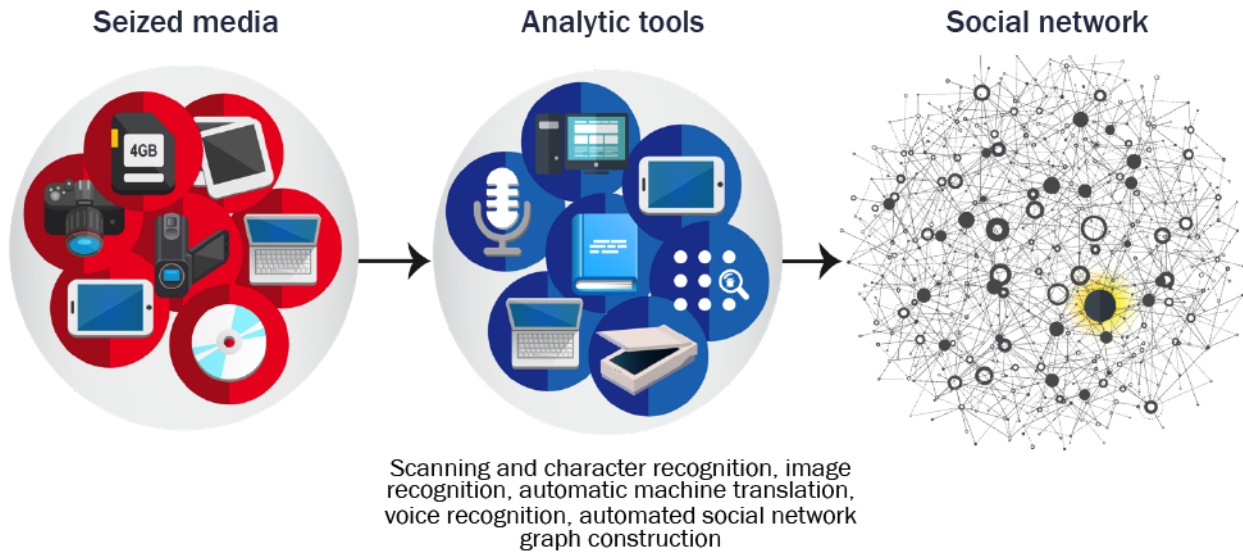


Figure 13 Examples of seized media are shown (left), along with tools that can make sense of the stored information in real time (center). The resulting social network can reveal a real-time threat (right).

approximating a human terrain map that unifies disparate data sets to reveal adversary actions, aspirations, capabilities and future plans (Figure 13).

High fidelity, real-time, situational awareness derived from this “digested information” can drive ongoing operations intended to optimize blue force defenses and sustain tactical, offensive advantage over adversaries already disrupted by the initial phase of an operation. When appropriate, significant leverage will be derived from the real-time marriage of these time-sensitive assessments of captured information, and the on-going restrike of adversary capabilities and initiatives.

Commercially available computing hardware and open source deep learning software will be an important foundation of this work, assisting in the production of finished intelligence in real time, revealing adversary networks, communications, plans, operational methods, and insight. Outputs from these tools can be opportunistically communicated to human analysts for further insight so that benefits accrue to both ongoing operations and the intelligence operations intended to inform them and future plans. The result will be more timely and actionable information in the field, significant improvements in ongoing operational leverage, faster and higher fidelity cues to human analysis, and deeper analysis derived from holistic assessments beyond the scope of current tools.

Recommendation 14.

The Defense Intelligence Agency (DIA) and the Special Operations Command (USSOCOM) should integrate commercial components and build a new machine-learning analysis tool, and prototype the resulting system using existing historical data, seized media, and commercial (collateral) sources. The cost is estimated at \$20 to 30 million over 2 to 3 years:

- USSOCOM should develop the capability to digitize any non-digital information and transmit this and all collected digital material in real time to analysis cells

- Rapidly sanitize information for sharing and dissemination to supported customers, the private sector, and allies
-

Autonomy for protection

The ability to prevent or mitigate the adverse effects of attacks on personnel, both combatants and non-combatants, and the physical assets of the United States, allies, and friends is an important DoD mission.³⁵ It includes both kinetic and non-kinetic attacks and other adversary offensive capabilities from all domains—surface, air, undersea, space, cyber and electronic warfare. It also addresses fixed facilities and locations, such as bases, borders, and air and missile defense sites, as well as mobile assets, including aircraft, ships, satellites, and personnel.

As adversaries make use of more sophisticated technology and weaponry, maintaining vigilance against potential attacks and responding rapidly to threats will require increasing use of autonomy-enabled capabilities. Because missions are largely defensive, fewer policy issues arise in comparison to those emerging around the offensive use of lethal autonomous systems.

Current and potential uses of autonomy in protection

Two of the six major drivers to stimulate the adoption of autonomy technology are behind many of the current uses in the protection missions. The two drivers are the required decision speed and the need for persistence and endurance. When both drivers are in play, there may be no alternative to autonomy for missions such as missile, space, and cyber-defense.

For these applications, collaboration with the human operator is often limited to system initiation and oversight. For example, in many air and missile defense systems there are two modes—manual and automatic. In the manual mode, the operator collaborates with the system to identify and validate adversary targets, and then launches an autonomous interceptor. The system then switches to automatic mode, where the operator monitors the system as it detects and engages targets autonomously unless the operator chooses to call off the engagement.

Protection missions, such as defending against an incoming missile salvo or cyber-attack, are also driven by both the need for persistence and rapid decisions. They are enabled by extensive signal processing of large volumes of sensor data coupled with the use of autonomous agents for rapid decision-making and actuation of control systems. The cyber-defense mission is almost completely analogous to the protection of civil and commercial information technology systems, and, consequently, DoD can benefit greatly from collaboration with the commercial sector to benefit from developments in that market.

When persistence is the key driver for autonomy, the autonomous system of choice has often been a UA that can protect large geographic areas over long times. These UA are generally remotely piloted today, with only a few functions delegated to autonomy, such as station-keeping or sensor

³⁵ Joint Capabilities Assessment, 2010 [Refinement approved April 8, 2011].

management. Systems may also have automatic target recognition, either alone or as a cueing aid for a human operator.

To protect limited geographic areas with reasonably predictable terrain features, unmanned ground vehicles (UGVs) have often been selected. The Mobile Detection Assessment and Response System (MDARS) is such an application that has numerous commercial analogs. Unmanned ground vehicles have also been deployed to minimize human exposure to hazards in the IED land mine clearance mission. The QnetiQ Talon (Mobile Tactical Robotic System, MTRS Mk1) and iRobot Packbot (MTRS Mk2) systems were deployed extensively in this role during the recent wars in Iraq and Afghanistan. The next generation of this capability is under development in the Advanced EOD Robotics System (AEODRS) program. As is true for UA, most UGVs are remotely operated and use autonomy technology almost exclusively for navigation and obstacle avoidance.

The study concluded more opportunities will emerge to delegate cognitive functions to an intelligent system in such areas as vehicle health monitoring or situational awareness. Such applications would exploit mature capabilities already in commercial use.

Commercial technologies are also outpacing military technology in autonomy in unmanned undersea applications. Developing and employing autonomous undersea systems has long been the purview of the U.S. Navy, but in recent years the commercial undersea survey and oil exploration industry and the scientific oceanography sector have taken the lead in deploying autonomous, often low-cost, platforms. While the Navy has kept pace in conducting foundational research and developing prototype systems in this area, there is significant value yet to be realized in operationalizing military systems. Currently deployed counter-mine applications use UUVs for persistence and protecting humans from danger, but rely on human operators at a command center to process data for target classification. This is followed by a separate mission to neutralize any mines detected. Autonomy can reduce both the time to neutralize the threat and the danger to the personnel assigned to the task.

Finally, the use of autonomy in the mission to counter chemical, biological, radiological, and nuclear threats is motivated by both the need for persistence and endurance and to protect humans from danger. Several current and potential programs aimed at using autonomy to reduce risk and improve protection of U.S. assets were identified.

Protection is an area where the benefits of autonomy have been well demonstrated, because it requires persistence and endurance that are often limited by human factors. Further, in many situations, protection requires speed of response or exposure to hazardous environments that may be better addressed by an autonomous system with an appropriate level of supervision by human operators.

The study reviewed the impact of a broad sampling of current and developmental uses of autonomy to improve protection capabilities. Many opportunities were identified to exploit the benefits of autonomy technology. Demonstration and early successes with deployed systems could accelerate adoption across the protection capability area.

Project #4: Dynamic spectrum management for protection missions

Today, the military use of the radio frequency (RF) spectrum is manual and largely pre-planned. Because of the complexity and dynamic nature of the environment, this approach can neither maximize use by U.S. forces nor deny adversary use. If not addressed, this situation will only worsen.

The opportunity presented by automating sensor, communications, and jamming coordination within the environment is to protect the ability to achieve information dominance while imposing high “cost” and disruption on adversaries.

Recommendation 15.

The study recommends two simultaneous and complementary programs that inform each other to achieve dynamic spectrum management:

- The U.S. Army Communications-Electronics Research, Development and Engineering Center (CERDEC), AFRL (Rome, New York), and the Space and Naval Warfare Systems Command (SPAWAR) should develop Military Service prototypes for local, agile spectrum deconfliction and control among a few systems. One demonstration per Military Service should be coordinated through semi-annual collaboration conferences. The estimated cost for this program is \$400 million over 5 years.
 - Each prototype demonstration should involve at least two non-collocated systems carrying out different but contemporaneous missions. A first demonstration is suggested for a locally shared electromagnetic spectrum common operating picture to demonstrate negotiation techniques that include agile adaptation to the electromagnetic environment. It should demonstrate the first implementation of centralized rules and policies, rather than centralized spectrum assignments and should be carried out in coordination with the single, joint program on spectrum management. The results should inform the evolution of that program.
- DARPA should develop an architectural framework and algorithms for near-real time, theater-level spectrum deconfliction and control for a full ensemble of joint, coalition systems. The estimated cost for this program is \$180 million over 3 years.
 - The recommended program would identify and develop the enabling framework and technologies for dynamic spectrum management over a large area involving thousands of systems, both friendly and hostile. The program would identify the protocols and algorithms required for distributed negotiation for mitigating inference among blue forces and maximizing electronic attack on red forces.

Project #5: UUVs autonomously conduct sea mine countermeasures mission

Sea mines have been laid during various conflicts since the Civil War and threaten both military and civilian maritime operations. The sea mine countermeasures (MCM) mission, localizing



Figure 14 Current Mine countermeasure capabilities utilize two separate vehicles – an autonomous UUV for search and detection (left) and a vehicle remotely operated by a manned ship in the mine field for disposal (right).

and neutralizing mines, is critical to ensure the safety of waterways for civilian traffic and to ensure military access to areas of importance.

Current manned and unmanned MCM platforms all require personnel in the minefield. MCM-1 class ships can detect, classify, and neutralize all known types of mines, but are manned by a crew of over 80 individuals. Increased utilization of autonomy-enabled UUVs can significantly reduce personnel risk during MCM operations by allowing uniformed personnel to supervise MCM operations remotely rather than entering the minefield. The MK-18 Mod 2, shown in Figure 14, program has demonstrated significant progress in utilizing Remote Environmental Monitoring UnitS (REMUS) UUVs for MCM. Further gains are possible. Current MK-18 operations, for example, continue to require long tactical timelines with intensive operator involvement, including a manned platform entering the minefield during the neutralization stage. Increased autonomy could reduce the demand for manning and personnel risk, and decrease the tactical timeline.

The study acknowledges that some of these recommended actions are already being explored and stresses the importance of continued progress in this area to further reduce personnel risk and decrease the time needed to conduct mission-critical MCM operations. For example, in the Single Sortie Detect to Engage program, ONR is developing enabling technologies to support integrating these two elements of the mine countermeasure mission into a single activity.

Additionally, while significant progress has been made in the basic research domain, there is a need for further research and development to harden and make robust the aforementioned capabilities. Continued focus in parallel R&D, particularly in improved automatic target recognition (ATR) capability and autonomous launch and recovery of UUVs, will ultimately enable delivery of an MCM mission package by larger unmanned platforms.

Recommendation 16.

The Navy Program Executive Office Littoral Combat Ships (PEO-LCS) should conduct a user operational evaluation system (UOES) program run by PEO-LCS in partnership with ONR. The estimated cost is \$60 million per year for three years. Suggested implementation actions for this demonstration project are the following:

- Equip an existing UUV platform, such as the MK 18 Mod 2, with embedded ATR capability to enable autonomous detection, classification, localization, and identification. Embedded ATR algorithms will utilize sonar and optical sensors to locate and identify mine-like contacts utilizing viable communication channels to query a remote operator to confirm identification. This will reduce the time for mine detection, classification, and identification, currently conducted over two sorties with intensive operator interaction, to a single-sortie task with streamlined operator interaction.³⁶
- Update an existing mine disposal platform, such as Seafox, with contact reacquisition and neutralization capability. Seafox is a wire-guided mine neutralizer that uses a shaped kinetic charge. Providing communications from the mine disposal platform to operators, along with higher-level autonomous control functions, will retain operator control of neutralization and will remove the need for personnel to enter the minefield to execute fly-by-wire operations.
- Expand delivery of the MCM sensing and neutralization nodes to an unmanned surface vehicle (USV) or large UUV to enable the detection, identification, and neutralization to be accomplished in a single sortie. As a stretch goal, delivery of this mission package by a large UUV could enhance mission covertness when required, for example, for MCM in support of an amphibious assault operation. The delivery platform would facilitate communications and command and control functionality as required.

Utilizing this acquisition model, fleet operators would work with developers during the course of the program to experiment with the system to rapidly evolve CONOPs, and design and characterize system strengths and limitations. After four years, the program would transition the enhanced MCM package to the Navy's 5th Fleet.

Project #6: Automated cyber-response

Despite increased awareness of cyber-threats created by the daily revelation of yet another audacious hack into systems containing sensitive and personally identifiable information, threat actors still seem to be winning in the daily contest with defenders across both the public and private sector. DoD classified and unclassified networks, the financial sector, industrial control systems, key elements of the DoD industrial base, elements of the power grid, cryptography firms, commercial companies, gaming industry, and even the entertainment industry have all been affected. The unfortunate truth is that even a greater number of incidents are not made public, and in a still greater

³⁶ D. Scheidt and G. Pollitt, *Hybrid Control Algorithms for Cooperating Vehicles Final Report, Appendix B: Benefits of Autonomous Operations (AO) Beyond Undersea Cooperative Cueing and Intervention* [September 2009].

number of incidents, the victim is unaware of the compromise.³⁷ In the known incidents, the security posture of the target was clearly not up to the challenge.

In many cases, the target failed to apply patches needed to address published vulnerabilities or inadequately trained staff made elementary mistakes. The aftermath of these incidents is unpleasant, embarrassing, costly to fix, and damaging to reputation and attendant investor and customer confidence. Worse still, the *modus operandi* for defending these systems is still largely focused on solutions that address adversary attacks based on previously observed behavior and subsequent mitigation based on detecting a recurrence. The systems employed to address individual adversary tactics are often not networked to achieve a comprehensive, let alone real-time, sense of adversary behaviors that cut across the otherwise stove-piped and unconnected defensive sensors. The present practice of bringing in security experts to remove the malicious software/hardware, restore the system, and determine attack attribution cedes initiative and advantage to adversaries who pick the time and place of their attacks and overwhelm defensive tactics with increasingly sophisticated campaigns.

Even after addressing the obvious shortfalls of not patching and inadequate training, systems will remain vulnerable to more sophisticated attacks that take place much faster than human decision cycles can address. The laudatory goal of perfectly secure systems is an impossible one, a reality deriving from their complexity, the constant changes they undergo as a result of system and infrastructure upgrades, and the reality of unpredictable human behavior on the part of both authorized and adversarial humans. As a result, security doctrine across all sectors is moving from one of creating and maintaining secure systems (impervious to attacks) to one of creating defensible systems that are well defended and supported by a diverse array of tools, authorities, and intelligence.

The emerging strategy represents a fundamental shift from a focus on inherent properties of systems, which remain important, to a dynamic understanding of the behavior of systems and actors and the active management or interdiction of them over time. Given the speed at which these changes take place and the complex nature of the systems themselves, we need greater levels of automation and autonomy to increase our ability to protect these systems within the timeframe of the attack. The foundation—Tutelage—for such a system already exists within the intelligence community. Today, this system provides real-time protection of the Non-classified Internet Protocol (IP) Router Network (NIPRNet), inspecting and analyzing more than 3 million packets per second for threats. Over the past five years Tutelage has prevented hundreds of millions of attacks. Additionally, this system is designed to provide actionable intelligence to DoD partners to help protect the nation. The proposed investments will significantly build upon these current capabilities.

A comprehensive network of sensors needs to be designed and deployed within blue space (systems we own), grey space (systems being used by the attacker that are owned by an unwitting third party), and even deeply buried within red space (systems being used to support the attack and owned or residing within the domain of the attacker). Clearly, there are technological, operational, and policy challenges with such an architecture.

³⁷ Verizon, *2015 Data Breach Investigation Report*. Available at www.verizonenterprise.com/DBIR/2015/ (Accessed January 2016.)

This sensor network that is fed into advanced autonomous systems will in real-time develop options to thwart the attack in the timeframe required to protect the target. Because many of these options could impact U.S. law, authorities, and policy, an option development engine must ingest these rule boundaries and then be able to rack and stack options that stop the attack while operating consistent with agreed-on rules of engagement. In the event that all options infringe upon at least one of the rules, the option and rule infringement should be highlighted, and a workaround identified or a waiver requested. While search results may not be able to stop an imminent attack, they could make systems ready for the next one.

In cases where the execution of the option is within the legal, policy, and authority bounds, additional infrastructure must be developed or existing infrastructure augmented to autonomously engage the cyber-weapon and block or stop the weapon's effect on the target. In different cases, the response option may target the weapon, the weapon infrastructure, the weapon command and control, or the effectiveness of the operator. Some examples of industry efforts for autonomous detection and mitigation include:

- Better analysis and detection can be facilitated by safe browsers with access to endpoint data. This has been used to analyze the Great Cannon attack on Github.³⁸
- Google Project Shield is an initiative to explore new ways of using Google's attack mitigation technology to offer news sites free protection from distributed denial of service attacks.³⁹
- Malware sharing forums, Internet protocol (IP) attack compilation statistics, industry cyber-attack analysis (IBM, Symantec, McAfee, Intel, FireEye, Verizon). Most sharing is done pairwise based on trust.
- The Structured Threat Information Expression (STIX) is a collaborative community-driven effort to define and develop a standardized language to represent structured cyber-threat information. The STIX language intends to convey the full range of potential cyber-threat information and strives to be fully expressive, flexible, extensible, automatable, and as human-readable as possible. All interested parties are welcome to participate in evolving STIX as part of its open, collaborative community. Trusted Automated eXchange of Indicator Information (TAXII) is the main transport mechanism for cyber-threat information represented as STIX. Through the use of TAXII services, organizations can share cyber-threat information in a secure and automated manner.⁴⁰

³⁸ Ars Technica, *Meet "Great Cannon," the man-in-the-middle weapon China used on GitHub* [April 10, 2015]. Available at arstechnica.com/security/2015/04/meet-great-cannon-the-man-in-the-middle-weapon-china-used-on-github (Accessed June 2016.)

³⁹ Google Project Shield. *Protecting free expression from DDoS*. Available at projectshield.withgoogle.com/public/about (Accessed June 2016.)

⁴⁰ STIX/TAXII Standards Transition – Frequently Asked Questions. Available at stixproject.github.io/oasis-faq.pdf (Accessed June 2016.)

Recommendation 17.

The U.S. Cyber Command should take the lead to develop an automated cyber-response, in partnership with the Central Intelligence Agency (CIA), the Federal Bureau of Investigation (FBI), NSA, DARPA (Plan X), key cyber-security industry players, and DISA. The team should implement a limited demonstration within one year and full operational capability within three years. The estimated cost is \$50 million per year for two to three years. Some suggested implementation actions include:

- Build upon existing success of the DoD’s current active defense system (Tutelage), leveraging private sector innovation and attendant rapid advances being achieved there:
 - Clarify authority and accountability within the Department for the overall design and implementation of the system that will connect and leverage sensors and analytics across DISA, the services, USCYBERCOM, and the combatant commands.
 - Design a family of sensors that can be introduced unilaterally and in partnership in blue, grey, and red domains.
 - Compile the legal authority and policy constraints in a form that can be ingested into the option development engine. Engage legal and policy experts early to identify obstacles.
 - Develop a global clandestine infrastructure that will enable the deployment of the defensive option to thwart an attack.
 - Ensure that the sensors, tools, options, and infrastructure used to support this enhanced defensive mission architecture do not compromise capabilities that support our various other missions that must remain much less visible.
 - Once designed, benchmark the effectiveness of the system against the plethora of historical attacks on U.S. targets.
 - Crawl, walk, and run in lock step with legal, policy, and partners to demonstrate system effectiveness with acceptable levels of unintended consequences—and in doing so, develop confidence and trust.

Autonomy for force application

Force application is the ability to integrate the use of maneuver and engagement in all environments to create the effects necessary to achieve mission objectives. Plans may include maneuver to insert, to influence, or to secure a location. Engagement may be through kinetic or non-kinetic means, using both lethal and non-lethal weapons. Force application planning assesses the most appropriate capability to achieve the objective.

Autonomy can improve the speed and accuracy, and by extension, the effectiveness of all aspects of force application. Anti-access and area denial (A2/AD) is a primary example of a mission that could be enhanced by autonomous systems. Autonomously operating UA could assume several functions now performed by manned aircraft in areas that are difficult to access (*e.g.*, aerial refueling, airborne early warning, intelligence, surveillance, reconnaissance, anti-ship warfare, and command). Additionally, large UA could be designed to dispense small UA that could operate autonomously to facilitate both offensive strike (via electronic warfare, communications jamming, or decoys), as well as defensive measures as decoys, sensors and emitters, target emulators, and so on—to confuse, deceive, and attrite adversary attacks.

These small swarms could be deployed as perimeter and close-in defensive actions with payloads tailored to the situation.

These concepts could be readily applied to other missions. For undersea missions, acoustic and RF decoy payloads would likely be much smaller than sea mines, and thus could be more easily deployed in quantity from existing commercial UUVs. While today's electromagnetic maneuver warfare

capabilities are limited, UUVs could provide a means to significantly extend capabilities and enable a covert option with a small observable footprint until electronic warfare (EW) operations are initiated. Typical communications systems could be emulated in the size, weight, and power available onboard 12-inch commercial vehicles.

Potential adversaries to the U.S. are creating systems (*e.g.*, very quiet submarines) and capabilities (*e.g.*, sophisticated sensors) that threaten U.S. forces as well as the undersea infrastructure. Moreover, with the current reliance on exquisite platforms, such as nuclear submarines, the U.S. runs the risk of being asymmetrically disrupted. To mitigate the risks, the U.S. must be more proactive and complement their submarine force with other capabilities, such as powerful new autonomous UUVs and sensor networks.

The Navy and DARPA have performed foundational work in many undersea areas, but there is no lack of additional possibilities to explore. Autonomous UUVs, in particular, hold great promise. Having been used by both DoD and the commercial sector (*e.g.*, in the oil industry), there are several UUV platforms that can provide a basis for rapid prototyping and experimentation.

Project #7: Cascaded UUVs for offensive maritime mining

One area, in particular, that could be leveraged more effectively is cascaded use of autonomous UUVs. With cascaded operations, extra-large autonomous UUVs (that may be close to 100 feet in length with large carrying capacities) would deploy smaller UUVs with targeted payloads, as illustrated in Figure 15. The extra-large UUV, acting as a submerged delivery vehicle, could be

In this study, an unmanned aircraft (UA) refers to a single asset. An unmanned aircraft system (UAS), however, refers to a system or systems of aircraft and payloads, command and control systems, communications architecture, ground stations, and CONOPs, which together comprise an entire capability that is greater than the sum of the individual parts. Because the individual UA are heterogeneous, the UAS provides an integrated capability far beyond that offered by a single UA.



Figure 15 The cascading unmanned undersea vehicle concept deploys smaller UUVs with targeted payloads.

launched from shores far from the area of operations or from surface ships, and then traverse autonomously to desired locations. Once close enough to the area of operations, the extra-large UUV would deploy UUVs with specialized payloads that could make their way to desired points for final action. An additional benefit of this model is that the extra-large UUVs would have the capacity to collect, process, shape, and then pass higher fidelity real-time contextual information to the smaller UUVs at the time of the smaller UUV's deployment.

Cascaded operations could enable many missions, including offensive mining, sea mine countermeasures, chokepoint monitoring and control, decoy delivery, and others in which it would be difficult to send surface platforms or submarines. While the concept is generalizable, there are many details to be worked and questions to be answered that would be influenced by the mission being executed, such as how critical recovery of significant deployed assets really is.

To illustrate the concept of cascaded UUV operations, consider offensive maritime mining. Today's offensive sea mining capabilities are limited, but UUVs could provide a means to significantly extend capabilities by increasing the influence range via mobility. Extra-large UUVs could be deployed from one or more shore sites or surface ships, and autonomously travel to an area of operations. Once they arrive, they could deploy a number of smaller UUVs or variants of modular torpedoes that have both automated target recognition capabilities and enough explosive material to disrupt or disable (or possibly even destroy) surface vessels. The UUV modular torpedoes would essentially serve as intelligent mines that are able to maneuver in an area and disrupt or disable adversary ships upon target verification. This would enable friendly forces to restrict adversarial freedom of movement and control access to key maritime areas, such as chokepoints in harbors. The UUVs also could be used to stop adversary ships from returning to their ports, thus precluding replenishment.

Recommendation 18.

U.S. Navy and DARPA should collaborate to conduct an experiment in which assets are deployed to create a minefield of autonomous lethal UUVs. The cost for this effort is estimated to be \$60 million over three years:

- The funding would cover the cost of leasing an extra-large UUV, purchasing or leasing USVs, purchasing or leasing UUV modular torpedoes, adapting the UUV modular torpedoes to work together, modifying and developing needed software, performing trade studies, establishing range and range requirements, and conducting the experiment.
- The experiment would culminate with an extra-large UUV traveling to the area of operations and deploying UUV modular torpedoes. The large UUV could also serve as an intermediate communications point for limited minefield control. The UUV modular torpedoes would be equipped with ATR capabilities and would draw upon a database of acoustic and other sensor signatures, including RF intercepts, to monitor ships in the area, some of which would be targets. Ideally, the UUV modular torpedoes would autonomously heal holes in the minefield caused by loss of mines due to, for example, vessel destruction. It should be noted, however, that this self-healing capability would be difficult to achieve, especially if being covert was critical and communications and control were highly constrained.

The experiment could be conducted in phases to address key technology areas in an incremental manner. The first phase of the experiment would use USVs to demonstrate the ability to detect and home on surface combatants using RF emissions and acoustic signatures. It would also demonstrate the cooperative behaviors needed to initiate a kill and maintain control of a chokepoint.

A second phase, if warranted, would be similar to the first, but would be performed with a field of UUV modular torpedoes. A key requirement of the second phase would be to demonstrate the needed underwater command and control. Before this phase of the experiment, a trade study should be performed to assess existing commercial UUVs as well as the state of development of the modular torpedo. Typically, the terminal approach to intercept a ship is difficult and requires a speed advantage. Existing commercial UUVs generally do not have the propulsion capability needed for effective terminal approach and may also be limited in payload capacity for a warhead capable of a hard kill. As a point of comparison, the MK62 Quickstrike mine carries 200 pounds of high explosives and the MK65 carries 1350 pounds of high explosives. Thus, use of existing commercial UUVs may result in reduced influence range, effectively requiring a ship to pass over the UUV in a tripwire rather than allowing the UUV to extend the influence range by introducing mobility to intercept the ship. Payload limitations (*i.e.*, less than a few hundred pounds of high explosives) may also require consideration of mission kill rather than hard kill. Thus, the modular torpedo that is currently under development may ultimately be a better option for addressing this mission, but it may not be ready in time for this experiment.

A final phase would use an extra-large commercial UUV to autonomously deploy the field of UUV modular torpedoes. The extra-large UUV could be launched from a surface ship at sea or from a shore. It would autonomously navigate to an undersea test range area. Once in the range area, the extra-large UUV would deploy UUV modular torpedoes. These UUV modular torpedoes would either loiter in the area or execute a search pattern to seek adversary targets. At some point in the deployment a message would be sent to either the extra-large UUV before releasing the UUVs or modular torpedoes or directly to the UUV modular torpedoes to let them know it was time to go active.

For covert communications, messages could potentially be hidden in acoustic masks designed to appear like sounds in the local area. These masks could, for example, be sounds of local maritime traffic or aquatic animals. In a sense, this can be seen as the use of steganographic techniques in the acoustic domain. If the message only needs to go to the extra-large UUV, then the UUV modular torpedoes would be active as soon as released. If, on the other hand, the UUV modular torpedoes could receive messages, then they could stay passive while loitering or searching and go active when they received direction to do so. The final act would be the navigation of the UUV modular torpedo towards the adversarial target, once one was recognized, and the detonation of surrogate explosive when the UUV modular torpedo was in close-enough proximity.

The combined set of experiment phases would require commercial unmanned surface vehicles, an extra-large commercial UUV, and several commercial UUV modular torpedoes. The extra-large UUV would need autonomous navigation capabilities to get to a particular area of operations. It would also need to be modified to carry and deploy smaller UUV modular torpedoes. These UUV modular torpedoes need to be equipped with sensors and ATR software. In addition, they would have to be equipped with surrogate charges that could be virtually detonated when the UUV modular torpedo was in close-enough proximity to a target ship.

Project #8: Organic tactical unmanned aircraft to support ground forces

To achieve the U.S. defense strategy's mandate to project power and win decisively, U.S. ground forces must be able to enter foreign territory in the presence of armed opposition as well as an advanced A2/AD environment (*e.g.*, air and missile defense, jammed communications, and so on).

Currently, tactical ground units engaged in asymmetric and near-peer conflicts are under constant threat, operating in an environment that is complex, constantly changing, and unpredictable. The speed at which ground units discover, assess, and react to battle-space change is vital to tactical success. A unit's agility, or the ability to rapidly respond to unexpected change, is known to be an important characteristic of highly capable units.⁴¹ Recognizing the relationship between decision-making speed and mission success, the U.S. Army promotes agility by instructing

⁴¹ D.S. Alberts and R.E. Hayes, *Power to the Edge* [CCRP Press, 2003]. Available at www.dodccrp.org/files/Alberts_Power.pdf (Accessed June 2016.)

soldiers to unilaterally take decisive action when necessary as long as that action falls within their commander's intent.

Unmanned aircraft system (UAS) support of an agile force requires that UAS capabilities be made available to platoon and squad-level units on an as needed basis. UASs carrying sensing, communications, jamming, and strike packages are capable of providing ISR, EW, overwatch, and tactical strike support that represent useful—potentially decisive—advantages for small tactical units. To be useful, UAS support must be timely. Since unit agility is often in response to unanticipated changes in conditions, the supporting UAS must be pre-positioned in advance of specific unit request. Because an operating theater may contain large numbers of small units operating over a large area, and the exact nature, timing, and location of unit need is unpredictable, a UAS may be required to provide simultaneous cover over a large area. In addition, the UAS manning requirement must be low so as to avoid placing an undue burden on front-line or supporting units.

Currently fielded UASs rarely factor into small unit ground combat because the systems and operating procedures used in their deployment cannot support the rapidly evolving needs of an agile force. Small organic UAS, such as the RQ-11 Raven, are difficult to use in unexpected engagements because the engaged unit is required to dedicate personnel to prepare and launch the Raven, which takes two soldiers to launch, distracting them from other tasks. Larger, centrally controlled vehicles, which include the Boeing Scan Eagle, RQ-7 Shadow, and MQ-9 Reaper, are difficult to use for small tactical forces. First, the cost and manpower requirements make dedicating a UAS for each small tactical unit infeasible; second, the time required to request for UAS support and subsequently vector support to an engagement is too long to be useful in many engagements. Thus, UAS capabilities are rarely used to support small unit actions.

The future battlefield environment will necessitate that ground forces achieve a high degree of autonomy within their indigenous weapons systems. This new reality will require blending of the attributes of autonomy with multiple UA—resulting in an effective force multiplier.

The value proposition of this blended approach is that autonomous systems:

- Are capable of operating in denied environments—supporting the rapid shaping of an engagement by denying adversaries effective communications and sensing while maintaining effective communications/sensing for blue forces
- Enable pervasive, persistent perimeters to be maintained—keeping the enemy at arm's length from capital assets. The “perimeter” can also be used to rapidly strike targets of opportunity (because they are pre-positioned).
- Permit persistent tripping of adversary forces' sensor networks—thereby complicating the adversary's ability to effectively determine and then target the critical elements of our attacks. This will produce new opportunities for assured access to previously denied areas.

Autonomously operated UA with various payloads for battlespace awareness, strike, or jamming, will provide small tactical ground units with the ability to protect themselves and facilitate offensive action. This operational concept substantially increases the operational tempo of unmanned vehicle capabilities supplied to front-line squads including:

- Supporting rapid, on-demand tactical strike
- Providing immediate alerts and battlefield intelligence (to threatened squads)
- Providing adaptive, continually reforming communications and navigation infrastructure
- Providing cover via coordinated deception and electronic warfare

Teams of heterogeneous autonomous UA could promote needed agility within small tactical units by providing an organic ability to use UAS support to anticipate threats, provide protection, and facilitate offensive action, as shown in Figure 16. Autonomous UAS support will provide immediate response to unit ISR, EW, and strike needs. In addition to these core capabilities, heterogeneous autonomous UA can further improve unit effectiveness by providing blue force communications, PNT, and blue force tracking.

Launched and recovered from a central base or ship, 10 to 40 medium-sized heterogeneous, autonomous UA could provide a persistent cover to small units operating over large areas. The UAS squadron would provide services to line units on patrol or located in forward outposts at speeds unattainable by human-piloted aircraft by accepting tasking from and providing services directly to the front-line user. Because the individual UA are heterogeneous, the UAS provides an integrated capability far beyond that offered by a single UA. Services provided by the UAS include electro-optic, infrared, and acoustic surveillance; jamming and EW spoofing; PNT, communications; and strike options. Presented with a diverse set of tactical needs from multiple units, the UAS squadron will self-organize, autonomously coordinating to satisfy the needs of all units.

The autonomous UAS squadron will be resilient and robust, and capable of operating in denied environments and accepting losses without compromising performance. Each UAS operates independently and is capable of using its own sensors, payloads, and autonomy software to perform mission objectives



Figure 16 The concept for organic tactical ground vehicle support with an unmanned aircraft system would enable small units to anticipate threats, provide protection, and facilitate offensive action.

provided to it by small units. Distributed artificial intelligence is used by *ad hoc* teams of one or more UA to self-organize, divide up tasking, and accomplish the mission objectives provided by small unit users. It should be equipped to support deception tactics and thwart deceptive tactics used by adversaries. Using EW payloads that include jammers, spoofing transmitters, and digital radio frequency memory (DRFM) transceivers, UAS teams will synchronize to provide misleading contacts, obfuscate blue force signatures, identify false targets employed by adversaries, maintain needed communications between small units and command, and limit adversary communications. Communications between a UAS and a peer UAS, end users, and senior command improve mission performance when available, but are required only when lethal force is a factor. UAS resilience is facilitated by the use of an *ad hoc*, delay and disruption tolerant network that allows each UAS to work with local cliques of peers and users when end-to-end connectivity is unavailable, and to use temporary local communications links to coordinate asynchronously. Resilient operations require that common operating pictures are also shared in an asynchronous, *ad hoc* manner that requires decentralized sensor fusion and delay tolerant information exchange. Heterogeneous autonomous UAS squadrons that are capable of performing missions without reliable communications will enable missions in denied environments or when stealth requires a communications blackout.

Small tactical units, particularly in an unexpected engagement, do not have the excess manpower or equipment required to direct UAS operations. For autonomous UASs to be useful, the interface between the frontline unit and the UAS support team must be minimal, consisting of an app on an existing handheld device or audio interactions over an ear bud and microphone, similar to the Apple Siri app. Use of the UAS team is managed by objective; users are not required to provide explicit waypoint directions or to monitor UAS performance.

Recommendation 19.

The U.S. Marine Corps, DARPA, ONR Code 30, and an FFRDC or UARC develop and experiment with a prototype heterogeneous, autonomous UAS support team that includes ten or more UA. The estimated cost for the effort over three years is \$40 million.

This experiment should be part of a Marine Corps training exercise conducted at 29 Palms or Camp Lejeune. The exercise should be conducted within three years with a stretch goal of quickly transitioning to a fieldable initial operating capability. In parallel with preparing for and conducting the exercise, the Marine Corps Combat Development Command (MCCDC) should conduct a design competition for a suitable production platform. The prototype UAS should be based upon commercially available hardware, including commercial airframes, payloads, radios, and ground stations. The baseline UAS should be capable of carrying between 2 and 12 kg of payload with a flight duration of at least twelve hours. The following tasks will be necessary in preparing for and conducting the exercise:

- Develop and assess CONOPs for local tactical employment of a small swarm UA fleet in a combat environment (MCCDC, FFRDC, or UARC)

- Integrate low cost UAS fleet of at least ten vehicles with distributed ISR, EW, communications, and strike payloads. The initial strike payload should be a non-lethal proxy that is suitable for experimentation purposes. (DARPA, ONR)
- Develop and integrate autonomy applique for UA fleet. The applique should support autonomous ISR, EW, communications, and strike missions. (FFRDC or UARC)
- Develop and integrate an *ad hoc*, delay or disruption tolerant communications infrastructure for UAS and small unit coordination (FFRDC or UARC)
- Develop and integrate an *ad hoc*, delay or disruption tolerant information management system capable of multi-sensor fusion and sharing common operating pictures (FFRDC or UARC)
- Develop and prototype a lightweight, mission-focused user interface that supports the CONOPs with audio and visual user-UAS dialogues on equipment with a mass of less than 1kg (FFRDC or UARC)
- Develop and refine a launch and recovery process in which a squad of no more than three is capable of recovering, refueling, and re-launching a UA in under 15 minutes (MCCDC)
- Conduct simulation-based testing to validate the CONOPs and prepare for hardware in-the-loop testing (MCCDC, FFRDC or UARC)
- Acquire a non-invasive test process capable of providing UAS safety assurance during experimentation in cooperation with the Test Resource Management Center (TRMC) and the Test Automation Center of Excellence (TACE)

To demonstrate that the heterogeneous, autonomous UAS concept is not only technically and logistically feasible, but economically feasible, MCCDC should conduct an unmanned air vehicle design competition. The competition should provide for an unmanned air vehicle design that includes:

- Non-proprietary avionics bus that supports modular payloads and a secondary processor suitable for hosting autonomy and data fusion software
 - Autopilot—a guidance and control sensor suite capable of stable, level flight and waypoint navigation
 - Standardized, non-proprietary communications payload bays
 - Non-proprietary remotely piloted command and control system
 - Integrated software switch allowing the remotely piloted command and control system to switch between automatic and autonomous operating modes
 - Minimum 12-hour flight endurance
 - Minimum 100 knot speed
 - Minimum 2 kg payload capacity
 - Modular attachment points to support a gimbaled payload and fixed wing and fuselage payloads
-

Autonomy for logistics

In any military operation, ensuring protection and timeliness of the U.S. supply chain while disrupting the supply chain of the adversary is arguably the most critical element of a successful military strategy. History validates this with numerous examples, such as the defeat of the British in the Revolutionary War and the defeat of Germany in the African theater in World War II. More recently, the victory over Iraq in Operation Desert Storm again proved the value of sound logistics strategy and execution. In a speech to Congress following the war, President Bush stated, “In a very real sense, this victory belongs to them—to the privates and the pilots, to the sergeants and the supply officers, to the men and women in the machines and to the logisticians who made them work.”⁴²

Logistics is the management of the flow of things between the point of origin and the point of consumption in order to meet requirements of consumers. The resources managed in logistics include physical items—personnel, equipment, weapons, ammunition, repair parts, food, water, fuel, medical supplies, and so on. It also includes the integrated flow of critical supply chain information to allow the warfighter to plan accordingly. Said another way, logistics is getting the right stuff to the right place at the right time—the what, where, and when of the resource equation.

Many commercial advances in logistics have been adopted by the Defense Logistics Agency (DLA) and the Military Services, resulting in significant gains over recent years through changes to historical distribution, maintenance, inventory management, and procurement practices. For example, the use of SAP supply chain software makes essentially autonomous many basic logistics functions such as customer order processing, contract solicitation, and contract award. Through these improvements, DLA increased sales to the military services from \$17 to \$46 billion over the past 15 years with no increase in logistics staffing and achieving record levels of readiness.

Employment of logistics autonomy can also be proactively used against an adversary. For example, speeding logistics helps get inside an adversary’s decision cycle. Dynamically distributing logistics operations can increase resilience and help thwart attacks against a single, fixed center of gravity. But as dependence on autonomy increases, adversaries will attempt to attack U.S. logistics autonomy, and so logisticians must learn how to deter, pre-empt, and defend against attacks. At the same time, the U.S. must develop methods and effects to counter adversaries’ employment of logistics autonomy.

Autonomy in commercial logistics

The study considered some of the rapid advances in the use of automation and robotics in the private sector. The competitive drive to reduce labor costs; improve efficiency in storage, distribution, and energy; capital investment utilization; and rapid responsiveness to market demands has driven commercial logistics to widespread automation and increasing autonomy. While DoD has increased logistics efficiency, commercial logistics leadership provides an opportunity for DoD to selectively adopt or adapt key advances. One critical area is the sensing, reasoning, deciding, and

⁴² George H.W. Bush, *Address before a Joint Session of Congress on the End of the Gulf War* [March 6, 1991]. Available at millercenter.org/president/bush/speeches/speech-3430 (Accessed June 2016.)

acting in an anticipatory manner about logistics, both with respect to low-cost consumables and high-value parts and repair.

Consumables: Real time situational awareness and anticipation

Commercial and retail sales in organizations have become highly automated and are increasingly incorporating autonomy to be better informed, anticipatory, and predictive. Low-cost sensors (*e.g.*, RF identification (RFID) location tracking), as well as taggants (*e.g.*, infrared, watermarking, DNA), enable product tracking as well as integrity monitoring. For example, DLA uses vegetable DNA to tag microchips.⁴³ In terms of sales, commercial providers track what consumers browse and purchase, and provide personalized recommendations for relevant items. Analytics leverage big data from purchasing, inventory systems, social media, weather models, and other sources to forecast demand more accurately in order to improve delivery time scales from days down to—eventually—hours.

Amazon recommender systems are well known for building models of a consumer's purchasing behavior using previous orders, product searches, wish lists, returns, or shopping cart data, and predicting what a customer would likely buy next.⁴⁴ Two-day shipping of these items is commonplace now, and the proposed Amazon Air promises even speedier delivery of products up to 5 pounds (86 percent of inventory) within a few years.⁴⁵

Amazon's acquisition of Kiva Technologies further increased the autonomy in fulfillment centers to include not only force multiplication (about four times) of human inventory pickers through robotic movement of materials in warehouses, but also dynamic reconfiguration (and wholesale movement) of inventories within warehouses based on anticipated delivery. Amazon's recent anticipatory package shipping patent extends this dynamic and anticipatory reconfiguration to enroute delivery. The method adds a complete destination address to a package after it has already shipped from the warehouse based on anticipated need reasoning about both a shipping model and a forecasting model. The potential is now emerging for 3D printing in trucks to print customers' product needs enroute, thus reducing storage and transportation costs, delivery time, and carbon footprints.

Walmart has similarly employed enterprise inventory management for predictive supply chain management. Analytics at Walmart Labs gather information from sources, including online purchase transactions, the long-term online shopping records or customer lifecycles of online consumers, and information on industry trends in e-commerce. Walmart recently purchased OneOps, an e-commerce,

⁴³ S. Freedberg, *DLA Demands Chip Makers Tag Products with Plant DNA: A War on Counterfeiters* [October 08, 2012]. Available at breakingdefense.com/2012/10/dla-demands-chip-makers-tag-products-with-plant-dna-a-war-on-co (Accessed June 2016.)

⁴⁴ P. Kopalle, "Why Amazon's Anticipatory Shipping is Pure Genius," *Forbes, On Marketing* [January 28, 2014]. Available at www.forbes.com/sites/onmarketing/2014/01/28/why-amazons-anticipatory-shipping-is-pure-genius (Accessed June 2016.)

⁴⁵ D. Guarini, "Amazon Reveals It Wants To Deploy Delivery Drones. No Joke," *The Huffington Post* [December 1, 2013]. Available at www.huffingtonpost.com/2013/12/01/amazon-prime-air-delivery-drones_n_4369685.html (Accessed June 2016.)

cloud computing development company, as well as Tasty Labs, which specializes in developing ways for retailers to connect with consumers through social media, and Inkiru, which analyzes “big data” to predict likely conversion rates and fraudulent transaction rates for particular promotions and marketing campaigns.⁴⁶

Other retailers anticipate supply chain disruptions. For example, Home Depot employs severe weather prediction and maintains several key locations outside but close to likely impact areas, then strategically places products so they can get goods to impacted areas quickly. The night prior to a storm, trucks loaded with goods proceed to the projected impact area. This is augmented with a post storm upturn in delivery for six weeks or more.⁴⁷ Ace Hardware employs a similar approach, stocking retail support centers with core items such as batteries, flashlights, generators, chain saws, and pumps, as well as clean-up items such as rakes, gloves, and garbage bags. Ace has also invested in a cloud-based transportation management system with a supplier portal, which enables them to interact with suppliers to determine when shipments are ready and send that information to carriers.

Parts and repair: Self-monitoring and anticipatory diagnostics

In addition to forecasting human purchasing behavior, commercial vendors forecast failures of critical parts to prevent failures. Vendors employ built-in, real-time self-monitoring and diagnostics on a variety of high value items. For example, a study by Google found that disc drives with increased heat, noise, and read/write errors detected by self-monitoring, analysis, and reporting technology (SMART) were 39 times more likely to fail, enabling active countermeasures, such as isolating failing sectors to prevent data loss or alerting maintainers to back up the disk.⁴⁸

A broad range of militarily relevant products such as engines, steam turbines, compressors, fans, generators, pumps, heating, ventilation, and cooling can benefit from monitoring and predictive maintenance. Monitoring methods include thermography, tribology (lube oil and wear particle analysis), ultrasonics, visual inspection, and digital testing and analysis.⁴⁹ A range of actions can be taken to generate maintenance alarms, initiate work orders, or recommend operator actions. Potential

⁴⁶ P. Demer, *Wal-Mart buys a ‘predictive analytics’ firm*, *Internet Retailer* [June 10, 2013]. Available at www.internetretailer.com/2013/06/10/wal-mart-buys-predictive-analytics-firm (Accessed June 2016.) See also L. Rao, *Walmart Labs buys data analytics and predictive intelligence startup Inkiru*. [June 10, 2013]. Available at techcrunch.com/2013/06/10/walmart-labs-buys-data-analytics-and-predictive-intelligence-startup-inkiru (Accessed June 2016.)

⁴⁷ J. Brown, *Forecasting the Unexpected: Home Improvement Retailers and Emergency Response*, *Inbound Logistics*. [July 2014]. Available at www.inboundlogistics.com/cms/article/forecasting-the-unexpected-home-improvement-retailers-and-emergency-response (Accessed June 2016.)

⁴⁸ E. Pinheiro, W. Weber, and L. Barroso, *Failure Trends in a Large Disk Drive Population*. Appears in the Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST’07) [February 2007].

⁴⁹ K. Mobley, *An Introduction to Predictive Maintenance*, 2nd Edition [2002].

benefits of anticipatory maintenance include system reliability improvement, operational readiness improvement, life extension, failure prevention, spare reduction, and maintenance facility reduction.⁵⁰

In applications where full autonomy can be employed without adding operator and maintenance burden, this would reduce personnel that could result in a reduction in logistics demands. For example, Amazon's employment of Kiva robots reduces by about a quarter the number of human pickers in a warehouse. An additional military benefit of an autonomous warehouse that required fewer personnel in a deployed area would be a reduction of personnel in harm's way.

Because of the high value of many of these benefits, the aviation industry provides many relevant examples. Customers use advanced diagnostics and engine management to help plan engine maintenance to keep costs down and availability up. Mature predictive techniques in Honeywell's predictive trend monitoring and diagnostics tools estimate how many hours are left before major repairs are necessary to its air-transport auxiliary power units (APUs) and to offer troubleshooting tips, including estimates of the probability of each tip's success. Deeper understanding can result in improved designs that extend lifespan, such as the incorporation of 3D printed fuel nozzles designed to stay on the wings 8 to 10 years before their first major overhaul.

Aircraft maintenance analysis today constantly monitors health and transmits faults or warning messages to ground control for customers. Such tools offer rapid access to maintenance documents and troubleshooting steps prioritized by likelihood of success. Real-time health monitoring systems in technical aircraft-on-ground maintenance control centers can give real-time troubleshooting assistance, guide spare provisioning, and monitor system health to anticipate failure. Regional aircraft manufacturers have also created tools that consolidate aircraft data from onboard systems to monitor and recommend maintenance.

Project #9: Predictive Logistics and Adaptive Planning

Data analytics are critical to state of the art commercial supply chain management.^{51,52} Predictive analytics, data mining, and decision support allow companies to be more agile, more effective, and more efficient; they can help identify opportunities, quickly address crises and support long-term planning of all activities related to logistics. For example, these systems can exploit the availability of historical data to anticipate opportunities and identify necessary actions (*e.g.*, prepositioning blizzard supplies in advance of weather, changing product lines based on buying trends at particular stores, reorganizing shelves to encourage additional purchases).

⁵⁰ H. Canaday, "New Predictive Maintenance Repair Overhaul (MRO) Tools Cut Costs: Gathering aircraft systems data is often easier than analyzing and determining how to act on it," *Aviation Week & Space Technology* [Feb 11, 2013].

⁵¹ J. Baljko, *Betting on Analytics as Supply Chain's Next Big Thing* [May 3, 2013]. Available at www.ebnonline.com/author.asp?section_id=1061&doc_id=262988 (Accessed June 2016.)

⁵² M.A. Waller and S.E. Fawcett, "Data science, predictive analytics, and big data: a revolution that will transform supply chain design and management," *Journal of Business Logistics* 34, no. 2 [2013], 77-84.

The DoD has even more compelling reasons for collecting and exploiting data about logistics. The operational situations are highly diverse. For example, expeditionary forces may need to bring everything with them, with some uncertainty about what “everything” is. Routine operations are much more predictable in their needs and in the modes for acquisition and transportation.

Geopolitical situations introduce constraints in availability of materiel, options for transportation, variability in information available and needed, sensitivity to culture and laws, and local rules of operation. As staff turns over, knowledge can be lost—both knowledge about how the systems function as well as knowledge about what decisions and options have been most successful in the past.

Logistics for the military means deciding on what resources, such as parts, fuel, materials, information, hardware, software, medical facilities, and so on, are needed by whom, at what times, and through what means, such as sources as well as transportation and handling required. The current logistics system (SAP) includes some capabilities for collecting and using historical data that may be helpful in developing predictive analytics and adaptive planning capabilities.

These capabilities require collaboration with operators to collect, exploit, and act on information. Such interaction may require the software to have a model of what the operator is trying to do and how they can be helped in doing so. A project at Carnegie Mellon University has incorporated plan recognition into a proactive agent for assisting planning for applications such as emergency response and peace-keeping missions.⁵³ The “Advisable Planning” project at SRI supported mixed-initiative planning by making the internal representations and reasoning accessible to operators.⁵⁴ In addition, the “Task Assistant” project at SRI enables organizations to capture and exploit knowledge about plans.⁵⁵

Such enhancements will reduce response time to user requests and reduce the cost of operation by impacting inventory, suitability of supplies, and transportation costs. In a larger sense, improved predictive logistics will decrease the time required to generate new plans.

The military has compelling reasons for collecting and exploiting historical data, considering the turnover in staff and knowledge lost as people move on, the high variance in the situations in which the logistics planning is needed, and the variability in the constraints governing different geographic regions. However, DoD currently develops point solutions rather than implementing a systems approach to address supply chain issues.

⁵³ J. Oh, F. Meneguzzi, and K. Sycara, “Probabilistic plan recognition for proactive assistant agents,” In *Plan, Activity, and Intent Recognition: Theory and Practice*, edited by G. Sukthankar, R. P. Goldman, C. Geib, D. V. Pynadath, and H. H. Bui, Elsevier [2014].

⁵⁴ K. Myers and T. Lee, *Generating Qualitatively Different Plans through Metatheoretic Biases*, in Proceedings of the Sixteenth National Conference on Artificial Intelligence [AAAI Press, 1999].

⁵⁵ B. Peintner, *Task Assistant* [2015]. Available at www.sri.com/sites/default/files/publications/1756.pdf (Accessed June 2016.)

Recommendation 20.

Naval Supply Systems Command (NAVSUP) should demonstrate the use of modern intelligent adaptive planning in conjunction with SAP. The cost is estimated to be \$10 million over two years. The key characteristics are to:

- Capture richer historical data, such as constraints that dictated the prior plan, metrics on costs and effectiveness, and logistician’s comments and notes
- Use analytics to make recommendations for new missions based on prior missions
- Autonomously make decisions such as critical supply prepositioning or dynamic plan adjustments

The use of analytics to make starting recommendations for new missions based on prior missions is a goal of the demonstration. The difficulty of resupply for the Navy makes them an excellent first adopter because they are already exploring opportunities such as additive manufacturing for creating parts on demand.

Project #10: Adaptive Logistics for Rapid Deployment

Deploying logistics warehouses is slow, costly, and problematic. Leveraging advances in commercial warehousing and logistics planning could reinvent military logistics deployments. For example, algorithms can use delivery deadlines to select the best combination of long-haul and short-haul shipping to meet all deadlines while minimizing shipping costs for the total flow of material. Machine learning and knowledge of upcoming sales can be used to preposition inventory geographically to anticipate demand while limiting over-ordering. As items become obsolete or as vendors change packaging, automation detects these issues, stops ordering obsolete product, and disambiguates between individual items, packages of items, or cases of packages of items. This keeps inventory accurate, and the right kind of number of items accurate within orders.

Entire warehouses using modular robotic components have been moved from one site to another over 48 hours.⁵⁶ Shelves were shrink-wrapped and shipped as is. Robot stations and chargers were unbolted from the floor and packed for shipping. Reestablishing the warehouse at the new site involved laying stickers on the floor, bolting down robot stations and chargers, then letting the robots autonomously store the shelves on the new floor. This was done within 48 hours in a commercial setting. Similar operations for infantry units took 144 hours. The time for an initial warehouse installation can be three to six months, significantly shortened from 18 months. Robotic sites can be designed, constructed, and brought online much more quickly than conveyance-based systems that require detailed design for every join and split point in material flow, and significant interconnection of components during construction.

⁵⁶ J. Dineen, “Meet the Robot Armies that are Transforming Amazon’s Warehouses,” *Forbes CITVoice*. [2015]. Available at www.forbes.com/sites/cit/2015/03/20/meet-the-robot-armies-that-are-transforming-amazons-warehouses/#781a0db723ee (Accessed March 2016.)

Recommendation 21.

Combined Arms Support Command (CASCOM), Ft Lee 10th Mountain Division, and the Ft Drum Joint Readiness Training Center should develop and deploy adaptive logistics decision support for a relocatable robotic warehouse and trained personnel in preparation for rapid deployment to unstable regions. The estimated cost for this effort is estimated to be \$30 million over three years. The implement steps for this recommendation should include:

- Create an adaptive planning process for rapid deployment
- Develop a relocatable robotic warehouse at Ft Lee
- Train a core group of logisticians to adapt plans
- Assign new logisticians to a unit
- Deploy to Ft Drum Joint Readiness Training Center with new capabilities
- Prepare for rapid deployment

CASCOM will need to baseline current operations and recommend level of adoption of robotic warehouse technologies and adaptive planning.

5 Expanding the Envelope

The previous chapter described a set of representative demonstrations that are “ready now,” based on the commercial and military advances of recent years. Autonomous systems offer the promise of even greater capability in the future, especially as commercial markets continue to drive the advance of underlying technologies—see, for examples, the “Imagine If...” possibilities provided in the Introduction to this report. However, converting commercial advances into military capability such as these requires two additional elements: operational pull and technology maturation in aspects unique to military needs. “Stretch problems,” as described here, are proposed as a mechanism to drive both elements, thereby expanding the envelope of technology available to support military goals. Here, each stretch problem is recommended in support of its own “Imagine If...” possibility. In practice, the recommended stretch problems also provide additional venues for practicing the first 11 recommendations of this study, those designed to “accelerate the adoption of autonomous systems.” Execution of stretch problems, such as those recommended, confronts the issues of trust and cultural barriers by building familiarity and by increasing transparency of autonomous “reasoning.” The problems involve varying types of human-machine teaming, can provide valuable insight into adversaries’ possible use of autonomy, tackle cyber vulnerabilities, emphasize the use of M&S, and offer opportunities for the T&E community to engage with learning systems.

A stretch problem is a goal that is “hard-but-not-too-hard,” and its purpose is to accelerate the process of bringing a new capability into widespread application. The most successful stretch problems are ones that largely leverage existing technology, with additional technology development as the “glue” necessary for integrating an end-to-end solution.

Stretch problems have been used successfully in a variety of implementations. For example, DARPA Grand Challenges (Mojave and Urban) offered cash prizes for successful demonstrations of autonomous navigation of unmanned ground vehicles, initially off-road and subsequently on-road in traffic. These successes galvanized today’s investment and progress in autonomous automobiles. More recently, the DARPA Robotics Challenge stimulated significant progress in controlling humanoid robots to support disaster relief missions. The Ansari X PRIZE awarded \$10 million to the first team to “build a reliable, reusable, privately financed, manned spaceship capable of carrying three people to 100 kilometers above the Earth’s surface twice within two weeks.” Other stretch problems are structured around a game construct: for example, RoboCup and *FIRST*[®] are both designed to allow teams to create their own autonomous players, with each team competing against others in structured events.

These various implementations of stretch problems have some essential commonalities: they accelerate progress by generating excitement, spurring creative approaches, and exploiting dynamic tension among competitors. In formulating these stretch problems, the tenets began with a crisp definition of a mission-relevant goal (or goals). The goal should be bold enough to capture the imagination and attract participants from the full range of technology providers: the fast-moving commercial industry, whether large commercial companies that can operate at scale or small start-up companies that drive some of the best innovations; and the academic community, including both

students and professors. Broad participation accelerates technology maturation as necessary to real the stretch goals.

These examples provide additional lessons. They typically involve repeated trials, some wildly unsuccessful at first but with later efforts building on the success of prior activities. Success metrics should be simple and clearly defined; they should describe “what” defines success without needlessly restricting “how” the goal is reached. Financial prizes should be awarded when the goals are met. Supporting infrastructure is required—competition or training ranges, starter-kits or basic platforms, simulation capabilities, data sets. The stretch problems that follow are structured around repeated competitions, and include a brief description of some of the supporting infrastructure anticipated to be necessary.

Another essential aspect of the proposed stretch problems is the participation of the full range of DoD stakeholders throughout their execution. Getting operators involved in these stretch problems will give them hands-on exposure to better understand autonomy’s value in military missions, creating “operational pull” and shaping the requirements for future procurements. The acquisition establishment must participate, to ensure the right capabilities are developed and fielded and to better assess the make/buy (or adopt/adapt) trades. The stretch problems also provide a venue for developing the testing methodology appropriate to autonomous systems. In addition, because autonomy is expected to provide the greatest value by enabling new missions (rather than in simply substituting machines for humans), it is essential that this hands-on experimentation explicitly consider the CONOPs, doctrine, and policy implications for new ways to use new systems. This requires up-front involvement from those communities, as well. Thus, a diverse set of DoD stakeholders must participate simultaneously as the operational pull is being created, so that all relevant equities are considered and iterated together.

This approach is in contrast to traditional DoD processes used for acquisition in which both operational needs and technology enablers are well understood. Because the types of autonomous capabilities described here are so unfamiliar, and the CONOPs for their use so undefined, the traditional, more sequential approach to stakeholder involvement would be fraught with opportunities for failure and delay. The reasons could be whether operational and policy conflicts were only identified after procurement, the development community did not make use of fast-moving commercial developments, testing methods were not suited to inherent system characteristics, or because of myriad other misalignments that could occur.

For these reasons, the study recommends that each stretch problem include active engagement by the full range of DoD stakeholders, as shown in **Table 2**. This aspect sets our recommendations for stretch problems somewhat apart from prior DoD experience with challenges and competitions. It may be most akin to the Army Warfighting Assessment at the U.S. Army Training and Doctrine Command (TRADOC), an annual wargame planned to start in 2017 and intended to bring together operators and industry to explore ideas for using new

Defense Science Board Summer Study on Autonomy

Table 2 Value in participating in stretch problems

DoD Stakeholder	Insight gained by participating in stretch problems:
Operators Doctrine writers	Potential uses, limitations, and vulnerabilities of autonomous systems
Policy makers	Implications for policies related to autonomous systems (<i>e.g.</i> , rules of engagement, etc.)
Testers	Testing methodologies suited to complex, software-intensive, and learning systems
Requirements community	Requirements for autonomy and counter autonomy systems
Acquisition community	Identification of new, high-payoff programs Limitations of commercially available technology, to clarify the adopt/adapt/develop acquisition strategy
DoD S&T community	Identification of priority focus areas for aligning technology investments
Non-DoD Providers	Attractiveness of participating in stretch problems
Academia	Maturing technology to enable solutions to important, hard problems
Start-up companies	Potential new markets where they have essential differentiation
Commercial industry	Potential new markets they can evaluate without normal burden of government contracting
Defense industrial base	Understanding system integration opportunities for future programs of record

technologies, as well as understanding their impact on tactics and CONOPs.⁵⁷ However, the study recommendation goes a step further in also advocating the involvement of policy makers and the testing community—the latter being involved not to impose traditional test methods, but to learn how to create and use test methods suited to software-centric, adaptive, and learning systems. The full spectrum of participants is a way to build trust, as described in Chapter 2.

Each stretch problem that follows is motivated by a vision for an important new military capability. Current technology enablers and shortfalls are identified. The essential role of autonomy in developing the capability is articulated and each stretch problem is outlined to show how it can accelerate the development and use of the envisioned capability.

Finally, it is important to emphasize what the stretch problems are not. They are not traditional Programs of Record that result in the procurement of materiel vetted for all the “-ilities”

⁵⁷ S.J. Freedberg, Jr., “AWA is not NIE: Army tries to buy weapons that work,” *Breaking Defense* [2015]. Available at breakingdefense.com/2015/04/awa-is-not-nie-army-wrestles-with-requirements-reform (Accessed June 2016.)

necessary for fielded systems. Rather, they are carefully constructed opportunities for the Department to engage with the full range of technology providers in a way that purposely accelerates technology maturation as they stimulate and clarify operational pull.

Generating future loop options

Imagine if national leaders had sufficient time to act in emerging regional hotspots to safeguard U.S. interests using interpretation of massive data including social media and rapidly generated strategic options.

Accurate and timely understanding of global social movements is critical for protecting U.S. interests. How many lives might have been saved with a timely anticipation of Arab Spring, or a clear understanding of situations unfolding around our embassies? Providing our national leaders with a well-considered slate of strategic options—diplomatic, information, military, economic—requires improved early recognition of emerging geopolitical events as well as an understanding of event drivers and repercussions, causal linkages, and possible non-kinetic solutions.

Such a capability may soon be achievable. Massive datasets are increasingly abundant and could contain predictive clues—especially social media and open-source intelligence. The U.S. uniquely enjoys access to open-source data as well as the full cadre of DoD’s multiple intelligence sources. Recent advances in data science, ever-improving computational resources, and new insights into social dynamics offer the possibility that we could leverage these massive data sources to make actionable predictions about future world events.

An autonomous early awareness system could:

- Ingest and event-code the wide array of data sources available today, including multiple intelligence sources, open source data, and social media, all in real-time with minimal human intervention to...
- Identify causal linkages between actions and outcomes globally to...
- Use these linkages to identify possible future outcomes, assess “what-if” scenarios, and analyze candidate U.S. courses of action.

The purpose of such a system would be to better understand possible future trajectories of unfolding events, and to help decision-makers assess various shaping options by estimating their likely impacts and repercussions.

Some essential first steps have recently been taken along the path towards creating such a system. For example, military and commercial systems have demonstrated the ability to forecast sentiment, threats, and disease outbreaks. DoD’s program on Integrated Conflict Early Warning System (ICEWS) uses static models and various raw data sources, news text, and econometric models to generate monthly forecasts for individual countries. IARPA’s Open Source Indicators (OSI) program, including Virginia Tech’s Early Model Based Event Recognition using Surrogates (EMBERS) efforts, have used high-velocity ingest of open source and social media data to demonstrate seven-day lead time for civil protest and three-week lead time for disease outbreaks, as

compared to the World Health Organization assessments. These and other efforts demonstrate it is possible to forecast future levels of instability—sometimes even insurgencies and rebellions—with remarkable accuracy, months into the future.

While impressive, these systems have limitations. They are based on correlations between situations and outcomes, which limits their ability to elucidate underlying causes. By analogy to weather forecasting, ICEWS can make predictions such as “40% chance of rain on Monday” but does not provide a clear explanation for why it will rain, or how the weather will unfold from Monday to Tuesday. OSI provides leading indicators of unrest (*e.g.*, indications of the arrival of a storm front), but not a predicted event trajectory or insight into what could influence the trajectory.

Creating the type of system envisioned here requires moving from today’s correlation-based “forecasting” models to models that identify the causal linkages that underlie emerging social movements. Such models would illuminate the interconnected drivers of observed behavior, and thus provide a basis for enumerating possible future event trajectories and assessing the impact of various courses of action. Continuing the weather analogy, causal linkages would allow us to go beyond simply forecasting a storm to understanding the possibilities for how storm might evolve over time, and to assessing the impact of options under our control, such as moving civilians to safer locations, releasing water from a dam, and so forth. But the envisioned system must also account for the fact that, unlike weather-related events, human behavior changes in response to our actions, and correctly capturing this feedback is essential. For DoD applications, such a system would require algorithms that sense the state of the world and build an internal representation of the underlying causal linkages. These algorithms would use statistically based extrapolation to identify possible future event trajectories and their likelihoods; algorithms for planning and analysis that generate large numbers of possible courses of actions and assess likely outcomes; and algorithms to assess the impact of an intervention to learn how to have the desired effect in the target country. These capabilities build upon but go beyond what is available today.

Causal models have a further advantage over correlation-based models in their ability to handle rare events. Correlation-based models require large training sets, so they have difficulty forecasting events that occur infrequently, such as *coups d’état*. Causal models replace the requirement for many *coup* examples in the training set with a requirement to clearly understand the precursors for a *coup* to take place.

Weather is a helpful analogy, but it is much simpler to predict the weather and its consequences than it is to unwind the myriad, interrelated elements that impact the trajectory of social movements. Nonetheless, recent advances suggest the creation of such causal models is becoming possible. An explosion of techniques in machine learning and deep learning; cognitive and social science modeling of populations, groups and individuals; and our growing understanding of the role and impact of social media on society and culture are key enablers.

A major difficulty in creating an autonomous system for identifying causal relationships is the absence of ground truth regarding predictions for future events. The recommended approach deals with that problem by relying on ground truth where available (*i.e.*, for historical events) and by

emphasizing transparency about model “reasoning” and evidentiary support for predictions regarding possible future trajectories.

Autonomy is essential in the envisioned system, first and foremost, because of the scope, variety, and complexity of data that must be continuously and quickly analyzed. In addition, the number of branching paths for future event trajectories and their various probabilities would be far beyond the ability of humans to track manually. Finally, autonomy will be necessary for both creating the underlying models and in generating their training data by event-coding various data sets. Today, each aspect requires significant human involvement, which represents a stifling bottleneck that must be overcome before it will be possible to operate at the scale and complexity envisioned in this application. Instead of manually creating each model and coding individual events, humans will use their expertise to create the causal schema (universe of candidate causal architectures) from which the machines can learn and select the appropriate causal models.

Recommendation 22.

DARPA should initiate a stretch problem designed to create a system that autonomously, globally, and in real-time identifies the causal linkages behind emerging social movements, and helps leaders understand the impact of possible courses of action along various possible future event trajectories. It is estimated this will take four years and cost approximately \$75 million in total.

- A critical requirement of this project is the construction of a digital test range as a scale model of society, analogous to the National Cyber Range, the scale model of the Internet used to carry out cyber wargames. DoD should build out the test range and equip it with huge volumes of unclassified data about recent historical and ongoing “events” from a wide variety of open sources, including social media. The test range must also include models that provide the best available simulation of societies and their interactions at multiple scales (*i.e.*, groups, regions, countries, and ideologies). The range should be staffed by government subject matter experts with backgrounds in social sciences, who operate the range; find, acquire, and curate the data; and procure and validate simulations. They facilitate and manage community involvement in the competitions.
- The competition should be structured in two parts: an ongoing “qualification” phase and a “prize” phase consisting of regularly scheduled competitions. The “qualification” phase is focused on historical data, where ground truth is available. In this stage the full complement of technology participants (*i.e.*, government, FFRDC, defense industrial base, commercial companies, academia, startups) are encouraged to use the test range to develop and calibrate their causal models, with the goal of each developing a system that accurately captures the dynamics that drive large-scale social, societal, or government movements. The systems ingest historical data and attempt to predict likely future trajectories for the movements. Participants qualify for the next round whenever they can accurately “predict” the outcome of a particular social or societal movement that has been withheld from the

training set and for which, during the qualifying test, they are provided data that stops short of the actual event to be predicted.⁵⁸

- The “prize” phase is available to all participants who successfully qualify as above, and this phase is repeated on a regular basis. Its focus is on predictions about the unfolding of events underway at the time of each competition, so absolute ground truth is elusive. In this phase, the predictive systems compete head-to-head, and prizes are awarded for best performance in two distinct categories. The first category focuses on the persuasiveness of a system’s predictions. Here, competitors’ models are judged against each other for their clarity and succinctness in identifying the causal relationships that drive the unfolding events that are the focus of the current challenge, and in providing the evidentiary support for those relationships. This focus on transparency of model “reasoning” helps create trust in the model predictions, and allows users to better recognize when the model is appropriate or not. The second category of award focuses on completeness in enumerating future event paths. Here, competitors’ models are judged against each other for their ability to identify a full range of possible future outcomes. More weight can be given to future outcomes that are strongly supported by evidence, as identified in the first type of competition. Prizes are awarded for the best performers in each category during every round of competition. The government user community should play a key role in judging and selecting the winners.
- Although beyond the scope of the stretch problem described here, it is expected that the most compelling models will attract further interest from the user judges. This could lead to an expansion of the test range to include classified data, so that the users and (appropriately cleared) competitors could further test and tune the most successful models based on the full data sets available to operational users.

Some critical elements of this program may include:

- A strong emphasis on the development of autonomous capabilities for both model building and event-coding, which are missing enablers for the overall system and missing from current programs in this area.
- Model development that emphasizes the generation of human-understandable explanation of the causal linkages discovered by models, rather than allowing the “explanations” to be implicit and hidden within proprietary code.
- Access to large volumes of high-quality training data. Just as access to training data ushered in an explosion in computer vision technologies, so will the creation, curation, and maintenance of a global event database serve to foster innovative approaches to building and testing the types of models required for this capability—and for other government needs.⁵⁹

⁵⁸ For example, all data associated with Crimea/Ukraine may be withheld from the training set, and the start of hostilities may be the event to be predicted. During qualification, “predict” means the system identifies the “actual” outcome as one of the possible likely outcomes generated by the system’s underlying causal models.

⁵⁹ The study notes that the government often pays multiple times for the same data. For example, one program manager suspected that the government must own 100+ separate contracts for Twitter data. In creating the training database for this program, the Department should determine how to host the data (data.gov? dataverse?) to make it available for other appropriate uses.

- Person-to-person interaction as an important contributor to innovation suggests the testbed be housed in Silicon Valley; Cambridge, Massachusetts; or other innovation hub in order to maximize in-person involvement by both commercial and academic participants, increase the exchange of ideas, and speed innovation.
 - While meeting these goals will build on recent commercial and military accomplishments, success will nonetheless require substantial, purpose-built technology development.
-

Enabling autonomous swarms

Imagine if commanders could deny the enemy's ability to disrupt through quantity by launching overwhelming numbers of low-cost assets that cooperate to defeat the threat.

In military applications, swarming is a convergent attack from many directions. As described in *Swarming & The Future of Conflict*, “Swarming is seemingly amorphous, but it is a deliberately structured, coordinated, strategic way to strike from all directions, by means of a sustainable pulsing of force and/or fire...”⁶⁰ Many scholars have considered the role that swarming might play in future warfare, including the role for autonomous weapons systems as part of a swarm (or constituting a swarm in its entirety). For instance, *Robots on the Battlefield II: The Coming Swarm* provides a recent, comprehensive discussion on the topic.⁶¹ As robotic platforms become increasingly capable, some role for autonomous swarms seems highly feasible—although not yet proven.

Robotic swarms could be constituted over a wide range of characteristics, as shown in **Error! Reference source not found.** The Department has a number of “swarm” efforts underway, and these are making important progress in understanding how robotic platforms can cooperate with each other and with humans in the performance of military missions. However, the study observes that the current efforts tend to align towards the *left* end of the possible range for the attributes itemized in Table 4. This alignment is remarkable for two reasons. First, the left end of the range is the more difficult to implement in autonomous systems. Second, the emphasis of current work is at odds with the admiration that military proponents of swarms often express when observing nature’s examples of swarms, such as hive insects, which are aligned with the *right* end of the table’s range of attributes. Hives consists of massive numbers of genetically identical organisms that obey very simple rule sets and have limited (or no) direct, peer-to-peer communication. Nonetheless, they can demonstrate sufficient coordination to accomplish tasks far beyond the capability of any individual, and even allow adaptation to changing environments.⁶² For example, fire ants use a few simple rules to build bridges out of their own bodies so that the colony can float or cross bodies of water, as shown in Figure 17.

⁶⁰ J. Arquilla and D. Ronfeldt, *Swarming and the Future of Conflict*, RAND/D8-311-OSD [RAND CORP Santa Monica CA, 2000].

⁶¹ P. Scharre, *Robotics on the Battlefield Part II: The Coming Swarm*, The Center for New American Research [2014].

⁶² While the organisms are identical, the functions they perform for the hive vary.

The simple self-construction methods result in proliferated air pockets that promote buoyancy and allow ants on the bottom layer to breathe.⁶³

Such significant accomplishments by organisms as simple as insects offer an intriguing possibility: that very large numbers of platforms with limited individual capability might be able to accomplish meaningful military missions. The possibility is attractive for several reasons. As limited-capability platforms, they could be cheap enough individually to be affordable for proliferation in overwhelming numbers; this would allow truly attritable assets, enabling exploration of missions and tactics in which “quantity has a quality all its own”. Limited communication or direct, peer-to-peer coordination between platforms offers a built-in robustness to the jamming and contested electromagnetic spectrum that is expected to be a constant on future battlefields. Finally, the simple rules followed by individual insects give rise to emergent behaviors, *i.e.*, the collective behavior is different than that exhibited by the individuals. The collective behavior of an emergent system can depend strongly on the environmental conditions, even when the basic rule set followed by individual members is essentially constant. In principle, emergent behavior could lead to highly adaptive military systems. However, predicting collective behaviors from the rules followed by individual entities is difficult, and today it would be difficult to know *a priori* if the collective’s adaptive responses would be beneficial or detrimental to a military mission.

For the remainder of this discussion, the term “swarm” is used to mean a massive collection of hundreds or thousands of simple autonomous systems with characteristics described by the *right* side of **Error! Reference source not found.** This is not to suggest that other uses of the term “swarm” are inappropriate or to question the value of missions they might carry out. The purpose of this definition is simply to clarify the type of swarm that is the focus here.



Figure 17 Raft built entirely of fire ants, where the building follows a few simple rules and results in a buoyant structure that allows ants to survive until they reach dry land.

SOURCE: National Park Service, available at www.nps.gov/akr/photosmultimedia/photogallery.htm?id=385E5498-1DD8-B71C-073997EB3E9682E1

⁶³ N. Mlot, C. Tovey, and D. Hu, *Dynamics and Shape of Large Fire Ant Rafts*. *Commun. Integr. Biol.* 5 [2012], p. 590-597.

Table 3 Forms of Robotic Swarms

Attribute	Harder to implement <i>Most DoD "swarm" efforts</i>	Easier to implement <i>Natural examples</i>
Diversity	Heterogeneous <i>e.g.</i> , mixed ground and air platforms	Homogeneous <i>e.g.</i> , standard platform, perhaps with modular payloads
"Intelligence"	High <i>e.g.</i> , complex reasoning	Minimal <i>e.g.</i> , simple, pre-defined rule sets
C2/decision making	Complex <i>e.g.</i> , highly interactive decisions	Minimal <i>e.g.</i> , implicit C2
Communications bandwidth	High <i>e.g.</i> , to provide detailed intra- (or extra-) swarm updates	Low <i>e.g.</i> , stigmergy (environmental marking)
Complexity of human interaction	High <i>e.g.</i> , could require advanced human-machine interface	Minimal <i>e.g.</i> , limited to human giving "Go" command

Massive swarms of this type might be used for various offensive or defensive missions (examples are provided below). Indeed, the most effective counter to massive swarms may be other massive swarms. The possibility of having to face massive swarms is a good reason for the Department to accelerate its understanding of these systems.

One type of mission in which an overwhelming number of simple platforms might be effective is in disrupting operations at a forward arming and refueling point (FARP). Exposed fuel, munitions, and runways (or landing areas) all represent points of vulnerability that could be disrupted with relatively small explosive payloads. These payloads could be delivered by a large number of small fixed- or rotor-wing aircraft using crude targeting methods, relying on random delivery within a defined area to statistically ensure coverage of the FARP, rather than pinpoint targeting or between-platform coordination and deconfliction. While these threats could be simple to target individually, doing so would distract from executing the FARP's mission and therefore degrade operations. And in sufficiently large numbers, they could overwhelm the FARP's defensive capacity by depleting magazine depth, or presenting more targets than could be prosecuted at once. In this case, one element of countering such a swarm might be another swarm, consisting of platforms that are at least as agile as the offensive swarm and that intercept and detonate the intruders before they can reach the FARP's sensitive points.

Another mission for a massive number of simple platforms could be in agile mines. An adversary could use mobile mines to continuously self-replenish a mined area, or to mine an area previously determined to be mine-free by U.S. forces. Such mines could use very simple rules and coordination. For example, they might detect whether they are close to any other mines; if so they could move in some predetermined manner, and if not they could remain in place. While we have the ability to deal with mines, an apparently endless self-replenishment capability increases their

disruptive capacity. Again, there might be a role for a counter-mine-swarm to augment our existing counter-mine capability.

As a final example, swarms may be well suited for roles in the radio frequency domain, where having a large number of platforms allows geographic diversity to be exploited in new ways. Swarms could enable a new and fatal form of communications “jamming”, in which a large number of quadcopter-borne RF sensors are pre-positioned to blanket a contested area. When a sensor detects adversary communications coming online in its vicinity (*e.g.*, simply by detecting an increase in signal intensity in the appropriate communication bands), the quadcopter could fly into the communications emitter and self-detonate. If the emitter stops transmitting before the quadcopter self-detonates, the quadcopter could assume another got there first and it could re-settle until needed again. Thus, peer-to-peer coordination could be unnecessary. A similar technique might be effective to counter a swarm of low-power, proliferated barrage jammers that are interfering with U.S. or allied communications and that are difficult to defeat using traditional methods. More advanced versions of swarms that use DRFM techniques could exploit geographical diversity in additional ways. For example, swarms could use DRFMs across multiple, mobile, blinking, and cooperating emitters to spoof U.S. radars or screen high-value targets. Counter-swarms having their own mobility and operating within the area of interest could help unwind ground truth from false data, by resolving and geolocating the full variety of emitters.

Recent advances driving the feasibility of such swarms include the proliferation of unmanned platforms—ground, air, and sea. For example, sales of quadcopters are exploding worldwide. Precise figures are difficult to determine, but estimates are that consumer drone sales have grown to over a million a year recently.⁶⁴ Autonomous navigation and control systems are readily available, as are small, high-quality cameras and other electronic components. Buzz, an open-source programming language specifically designed for understanding and predicting swarm behavior, may also speed innovation in this area.⁶⁵

The swarms contemplated here are autonomous, by assumption. They are intended to have the simplest possible interaction with humans, along the lines of receiving a “go” command from their human operators. Even without this restriction, the hundreds-to-thousands of individual platforms would be beyond the ability of humans to control directly.

⁶⁴ A. Amato, *Drone Sales Numbers: Nobody knows, so we venture a guess*. [April 16, 2015.] Available at dronelife.com/2015/04/16/drone-sales-numbers-nobody-knows-so-we-venture-a-guess (Accessed June 2016.) See also R. Lever, *Drones swoop into electronic show as interest surges*, YahooTech. [January 7, 2015.] Available at www.yahoo.com/tech/s/drones-swoop-electronics-show-interest-surges-061549575.html (Accessed June 2016.)

⁶⁵ C. Pinciroli, A. Lee-Brown, and G. Beltrame, *Buzz: A novel programming language for heterogeneous robot swarms*. Available at robohub.org/buzz-a-novel-programming-language-for-heterogeneous-robot-swarms (Accessed May 2016.)

Recommendation 23.

The Assistant Secretary of the Army for Acquisition, Logistics, and Technology (ASA(ALT)), with close participation by the Army Capabilities Integration Center (ARCIC), should establish an annual “swarm games” challenge. This is expected to cost about \$25 million per year, including funding for facilities, participants, and equipment.

- This stretch problem is envisioned as an annual game to encourage open exploration of a variety of concepts. Each event should define specific mission goals that are appropriate for massive swarms, and participants may field teams to either accomplish or defeat those mission goals (“offense” and “defense”). This asymmetry makes these swarm games different from other robotic games: the two sides have different purposes, they do not have to play by the same rules, and they need not field identical teams. To further enhance the realism of these games, the organizer may choose to impose different rules of engagement on the two sides.
- The games can be structured as a series of head-to-head competitions, with prizes given for the best offense, best defense, and the overall winner. The game organizer is responsible for clearly defining the metrics for each category, as appropriate for each mission. Candidate missions include disrupting operations at a fuel depot, attacking a fixed facility, disrupting ground-force maneuvers, jamming communications, and spoofing and decoying sensors.
- There are several critical considerations for these games. First, it requires an outdoor test arena suitable for (simulated) kinetic operations; the facility must include instrumentation to measure impact on operations, score the games, and play back events. Second, testing should include persistent RF jamming, to enforce the strong limitations on intra- and inter-swarm communications necessary for operations in RF-denied environments. Third, the government should offer to furnish participants with basic platforms and payloads to encourage focus on the development of algorithms and CONOPs. Finally, the government should make available a common simulation environment for use by participants and use simulation results to determine which participants may progress to the live games.

Intrusion detection on the Internet of Things

Imagine if commanders could defeat adversary intrusions in the vast network of commercial sensors and devices by autonomously discovering subtle indicators of compromise hidden within a flood of ordinary traffic.

The Internet of Things (IoT) is the set of IP-addressable devices that interact with the physical environment. IoT devices typically contain elements for sensing, communications, computational processing, and actuation. They span a range of complexity and physical size—from thermostats to traffic lights to televisions, from mini-drones to full-size vehicles. Applications include media targeting, data capture, environmental monitoring, infrastructure management, manufacturing, energy management, medical and healthcare systems, building and home automation, and




Figure 17 The Internet of Things is increasing rapidly in both numbers and types of smart objects.

transportation. In principle, the IoT can incorporate almost any device or function imaginable—and countless not yet imagined.

The IoT is vast and growing rapidly, as shown in Figure 18. According to one estimate, more than 50 billion IoT devices will generate more than \$3 trillion in spending by 2020.⁶⁶ The Department of Defense is already a major presence on the IoT as a large purchaser of many kinds of innovative commercial devices, and it is likely to develop or exploit many more IoT devices in the future. Additionally, the myriad devices of more than three million service members and employees also reside in this globe-spanning network.

The seemingly limitless opportunities afforded by the Internet of Things also bring deeply embedded risks. Here we focus on just a few aspects of those risks: vast scale, limited configurability, and already demonstrated security flaws.

With the shift to the IPv6 addressing standard, the number of devices that can be networked is many orders of magnitude larger than the current Internet. IPv6 uses 128-bit addresses, meaning that it can, in principle, handle up to 340 undecillion unique devices (340 with 36 zeroes after it); this compares to fewer than five billion devices under the outgoing IPv4 standard. This immense, sparsely populated space of interconnected devices could serve as a globe-spanning, multi-sensing surveillance system or as a platform for massively proliferated, distributed cyber-attacks—or as an immense test range for real-world, non-permissive testing of large-scale autonomous systems and swarms.

By their nature, the IoT's small, networked devices are designed for simple operation. The devices' inner workings are hidden to the degree possible, and configuration options are limited. This makes it very difficult for end users to mitigate risks or to detect when devices or data have

⁶⁶ *Gartner Says 6.4 Billion Connected "Things" Will Be in Use in 2016, Up 30 Percent From 2015.* Available at www.gartner.com/newsroom/id/3165317 (Accessed June 2016.)

if the effects of a recent demonstration where attackers took over a Jeep's control systems were replicated across hundreds of vehicles at a time.⁷⁰

The widespread and growing adoption of the IoT, along with its inherent and largely unaddressed security issues, mean that it represents a threat that could soon dwarf that of the internet. Vint Cerf, one of the fathers of the Internet, was recently quoted about his views on the IoT. While acknowledging its potential benefits he also admitted, "Sometimes I'm terrified by it. It's a combination of appliances and software, and I'm always nervous about software — software has bugs."⁷¹

For the nascent opportunities in the IoT to be fully realized, the public must be convinced that what has been called "the greatest mass surveillance infrastructure ever conceived" will not share information gleaned from our devices unwittingly with "the original manufacturer, the information services we subscribe to, national security agencies, contractors, cloud computing services, and anyone else who has broken into, or been allowed into, the data stream".⁷² This is especially important for DoD, given its large and growing exposure to the IoT and its special attractiveness as a target.

As a first step, the Department should develop a capability to detect large-scale intrusions on the IoT without having direct access to the individual devices. This is different from attempting to detect compromise of individual devices. Instead, it should focus on the characteristics and signatures associated with the remote activation of massive and/or coordinated intrusions of IoT devices. Developing this capability would require new algorithms and techniques to detect the changes in operating behaviors of large numbers of IoT devices as seen from various observation probes.

The feasibility of such a detection scheme is based on large-scale techniques already in use today by ISPs, multinational companies, and private threat intelligence companies to detect large botnets, worm outbreaks, and other major Internet events. A number of techniques, such as border gateway protocol (BGP) event monitoring, unused IP address space monitoring, passive domain name service (DNS) analysis, netflow changes, and information sharing clearing houses, all play a role in large-scale event detection on today's Internet. New approaches can build upon these proven techniques to help achieve the IoT system envisioned.

The scale and speed of IoT attacks will overwhelm human-in-the-loop defenses. Autonomous systems will be necessary to deal with the massive amount of data to be processed, as well as the speed necessary to defend and act within the difficult and diverse ecosystems of the IoT.

⁷⁰ A. Greenberg, "Hackers Remotely Kill a Jeep on the Highway—With Me in It," *Wired*, [July 21, 2015]. Available at www.wired.com/2015/07/hackers-remotely-kill-jeep-highway (Accessed June 2016.)

⁷¹ K. Noyes, "Sometimes I'm terrified of the Internet of Things, says father of the Internet," *IoT Council: A Thinktank for the Internet of Things* [August 26, 2015]. Available at www.theinternetofthings.eu/katherine-noyes-sometimes-im-terrified-internet-things-says-father-internet (Accessed June 2016.)

⁷² See, for example, J. Powles, "Internet of Things: the greatest mass surveillance infrastructure ever?" *The Guardian*. Available at www.theguardian.com/technology/2015/jul/15/internet-of-things-mass-surveillance (Accessed June 2016.)

Recommendation 24.

DARPA should develop autonomous systems that detect large-scale intrusions on the IoT, by passively and remotely monitoring bulk network traffic, and identifying aggregate indicators of compromise hidden within the flood of ordinary traffic. The program should include a series of competitions over a period of five years at an estimated cost of \$80 million, which includes funds for a testbed, real and synthetic data sets, prizes, and block grants.

- This stretch problem is envisioned as a series of competitions to meet increasingly difficult intrusion challenges. Participants win prizes for outstanding performance as either intruders or detectors, with prizes for successful detection being ten times greater than for successful intrusion. This difference is because intrusion is inherently easier, and the prize differential will encourage innovation in detection approaches while still rewarding those who help elucidate new threat (intrusion) pathways.
- The periodic challenges would become progressively more difficult. For example, the first challenge may only require the participants to be able to detect a “noisy” and overt attack on the IoT test range using a large number of passive sensors. The next challenge may require participants to build upon the previous success and detect a “quieter” and stealthier intrusion using fewer sensors. Subsequent challenges would increase the complexity of the detection problem while reducing the quantity or quality of the sensing data relative to the size of the test range.
- To encourage a diverse range of participants, refreshed over time, the program should establish a regularly scheduled series of head-to-head competitions and, at the start of each new challenge, provide a round of block grants to support the development of the most promising approaches.⁷³ Block grants, rather than contracts, are recommended to maximize the likelihood of engaging commercial, start-up, and academic groups. In order to maximize innovation, each competition event should be open to any qualified participant regardless of whether they had received a grant that round or whether they had been involved in earlier competition cycles.
- An essential aspect of this challenge is a test range that simulates the scale and heterogeneity of full IoT. This range must be purpose built (perhaps built out from an existing cyber range) and include a large, diverse set of IoT devices within an instrumented network that can ingest, store, and process extremely large data sets generated by the devices and their network traffic. In addition, synthetic IoT data is likely to be required to enhance the realism of the test environment by increasing its scale virtually. This facility will be used to conduct the periodic competitions, and in between these events it should be made available to qualified competitors to allow them to test out their approaches under the most realistic conditions possible.

⁷³ This approach has been proposed in DARPA’s Cyber Grand Challenge, a DARPA program planned to launch at DEFCON in August 2016.

Building autonomous cyber-resilient military vehicle systems

Imagine if commanders could trust that their platforms are resilient to cyber-attack through autonomous system integrity validation and recovery.

The vulnerability of networks to cyber-attacks is increasingly understood, and new methods are being developed to handle cyber-threats. Today the most robust solutions rely on large quantities of data collected from globally distributed, cooperative sensors, and advanced analytical methods carried out using high-performance computing. The best techniques not only carry out real-time cyber-defense, they also extract useful information about the attacks and generate signatures that help predict and defeat future attacks across the entire network. They are powerful, resource-intensive and reliant on high-bandwidth network access. Said differently, they exploit network resources to protect networks.

Autonomous and semi-autonomous vehicles, like networks today, also face cyber-threats and, like networks a decade ago, they often have limited defenses. For example, in 2011 the Predator and Reaper UA cockpits at Creech Air Force Base were infected with malware that proved very difficult to remove.⁷⁴ As the degree of autonomy increases in U.S. platforms, the cyber-vulnerability of subsystems will have increasing impact. Increased autonomy is inevitable—even today, humans cannot operate some high-performance vehicles without autonomous subsystems to maintain platform stability; the loss of these subsystems would force the vehicle to operate in a degraded state, if it could operate at all. Autonomous (unmanned) platforms lack a human operator to take over in the event of subsystem compromise. And further, fully or partially autonomous platforms may have to operate in communications-denied environments, limiting the value of defensive measures that are off-board or reliant on networking.

Thus, the network-centric cyber defenses that provide the best defense for networks are not well suited for providing cyber protection to autonomous platforms (or autonomous subsystems). Protecting autonomous platforms requires a different paradigm.

Rather than focusing on *robustness*, as is traditional in cyber defense of networks, the cyber-protection of autonomous (or semiautonomous) platforms should focus on *resilience*. Robustness seeks to ensure resistance to an attack, whereas resilience emphasizes rebound from attack and/or operating through the attack with as much mission performance as possible.⁷⁵

An emphasis on resilience opens new options for protection approaches. Because a system under attack would not seek to fully understand or decisively defeat the attack, it could take a more limited approach to defending itself. It would need to detect the fact—and unfolding—of the attack, for example, by run-time integrity validation. It would need to recognize which subsystems were corrupted, and be able to autonomously determine the criticality of each affected system for the

⁷⁴ N. Shachman, *Computer Virus Hits U.S. Drone Fleet* [October 7, 2011]. Available at www.wired.com/2011/10/virus-hits-drone-fleet (Accessed June 2016.)

⁷⁵ D.D. Woods, “Four concepts for resilience and the implications for the future of resilience engineering,” *Reliability Engineering & System Safety*, 141 [September 2015], pp. 5-9.

current mission set. It would need the capacity to restore essential subsystems from known good images, and to isolate or shut down non-essential systems as appropriate.

Some aspects of this type of system have already been demonstrated. For instance, Volexity advertises a run-time system that continuously validates the integrity of a computer's operating system via analytics carried out in random access memory (RAM). DARPA's Clean-Slate Design of Resilient, Adaptive, Secure Hosts (CRASH) is creating new computing architectures that focuses on system security, processing/memory segmentation, and resilience.

Further work remains in optimizing methods for hardware- and software-based integrity validation, autonomous assessment of subsystem compromise, and autonomous adaptation, including the restoration or shutdown of subsystems. It may be useful to develop so-called trusted "sidecar" modules that can easily integrate with various vehicle platforms under meaningful size, weight, and power constraints. These modules could execute out-of-band system-integrity assessments as well as host and restore the known good subsystem images. Such sidecars could also hold slight variations in subsystem images, to increase the likelihood of resistance to any specific attack. As well, a sidecar architecture could facilitate between-mission updates.

Autonomous systems, especially those unable to communicate with humans, require the ability to defend themselves autonomously. Even for autonomous subsystems that are components of larger systems with humans in the loop, the timescale required to respond to cyber-attack can be far too short to allow human involvement.

Recommendation 25.

DARPA should implement a stretch problem to demonstrate autonomous cyber-resilient systems (ACRS) for autonomous military vehicles. A competition should be run annually for six years at an estimated cost of \$60 million.

- This stretch problem should be structured as a series of competitions to meet increasingly difficult cyber-attacks. Participants are awarded prizes for outstanding performance as either cyber-attackers or defenders, where defenders create resilient systems that enable autonomous vehicles to operate through the attack. Prizes for resilience are ten times greater than prizes for successful attacks. This difference is because cyber-attack is inherently easier than autonomous cyber defense, and the prize differential will encourage innovation in defensive approaches while still rewarding those who help elucidate new vulnerabilities.
- The annual competitions would be progressively more difficult. For example, the first competition may only require a candidate ACRS to detect and recover from an obvious disruption-style attack that is attempting to disable critical platform subsystems. Subsequent competitions might require candidate ACRS' to detect and recover from increasingly stealthy compromises that attempt to subvert system functions versus simply trying to disable subsystems. Later competitions could also increase the attack rates and sources to further validate the ACRS' ability to deal with a highly contested cyber-

environment, possibly with multiple adversaries using different attack techniques and approaches.

- To encourage a diverse range of participants that is refreshed over time, the program should create a regularly scheduled series of head-to-head competitions and, at the start of each new challenge, provide a round of block grants to support the development of the most promising approaches. (Block grants, rather than contracts, are recommended to maximize the likelihood of engaging commercial, start-up, and academic performers). In order to maximize innovation, each competition event should be open to any qualified participant regardless of whether they had received a grant that round or whether they had been involved in earlier competition cycles.
- Participants must be provided limited access to military vehicle operating systems and hardware architectures. The block-grant approach also provides an opportunity for screening the suitability of potential participants, independent of whether they receive government funding. In addition, the program requires a range for the autonomous operation of military ground and air vehicles, which is instrumented to capture ground truth of mission effectiveness. The test vehicles must be allowed to be subjected to cyber-attack. The range must also be suited to cyber-attack, including appropriate instrumentation to assess cyber-“health” of the platforms. Between competitions, the range should be made available periodically to qualified participants for ongoing development and testing.

Planning autonomous air operations

Imagine if commanders could operate inside adversary timelines by continuously planning and replanning tactical operations using autonomous ISR analysis, interpretation, option generation, and resource allocation.

The Joint Air Tasking Cycle is illustrated in Figure 20, which highlights the centrality of the Master Air Attack Plan (MAAP). Decisions made during the MAAP become the core of the daily Air Tasking Order (ATO). The MAAP process currently takes 12 hours and must be completed 24 hours prior to execution to allow time for subsequent ATO generation and dissemination to units for detailed mission execution planning.

The range of functions carried out in the MAAP generation process is shown in the lower right of Figure 20. Today, the process is heavily manual, with as many as 40-50 people required for generating the master plan for a large operation. Human planners do mission and resource planning, aided by stand-alone (non-integrated) models for individual platforms and effects.

The current timelines to complete the MAAP/ATO process are too long to effectively counter an adaptive adversary. In fact, there is ample evidence that the timeline has been too slow even for the relatively modest threats faces over the past several decades.⁷⁶ For example, by the end of Operation Enduring Freedom the time between target identification and target destruction had

⁷⁶ P. Winkler, *The Evolution of the Joint ATO Cycle*, Joint Advanced Warfighting School, Joint Forces Staff College [2006]. Available at www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA451239 (Accessed June 2016.)

Defense Science Board Summer Study on Autonomy




Figure 19 The range of functions carried out by the MAAP team within the Joint Air Tasking Cycle is a manual, slow process.

shrunk to under 20 minutes. This acceleration was the basis for the formalization of the concept of kill box interdiction and close air support (KI/CAS) during Operation Iraqi Freedom, which enabled aircraft to be launched without preplanned targets in order to handle dynamic targeting. A subsequent analysis showed that almost 80 percent of the targets struck were of this class, *i.e.*, they were outside of the formal ATO process (although they did rely heavily on the MAAP to assign a geographic kill box to platforms).

Experience proves we have adapted our planning process well to defeat the regional powers and non-state threats of the past several decades. As we prepare, however, to meet a near-peer threat we must expect the planning problem to become more complex, with more platforms of different types (including a mix of manned and unmanned assets), potentially executing multiple missions or operating in tight collaboration with other platforms, all within a congested and contested electromagnetic spectrum. Planning complexity will increase, and at the same time, a near-peer can be expected to use tactics specifically designed to defeat a multi-day planning cycle such as we use today. A clever adversary will generate multiple plans, initiating one as a foil for the U.S. planning function and then switching to another one after it is too late for the U.S. to adapt. We must prepare to counter such tactics.

Autonomy could enable a much faster planning cycle through the use of integrated tools and models that can handle the allocation of large numbers of resources in complex sequences with high interdependencies. It will be essential in managing and optimizing the complexity of branching scenarios associated with what-if analyses and pre-planning for contingencies. And, during execution, it enables a running comparison of “plan” to “actual,” to assess whether large-scale re-planning is merited as events unfold. In this sense, planning is no longer a phase that happens before execution. Instead it becomes a continuous, background process that assists the commander in redirecting assets when needed during execution.

The tools to shorten the MAAP/ATO process by an order of magnitude do not yet exist. However, there is evidence of enabling progress in many underlying capabilities. For example, DARPA’s Distributed Battle Management program is developing automated decision aids for managing air-to-air and air-to-ground combat.

Autonomy is required to handle the scale and complexity of the planning at speed. The study notes that a better planning tool on its own is not sufficient to shorten the planning timeline. That will also require changes in workflow, staffing, and other factors.

Recommendation 26.

The AFRL/RIS office should undertake a stretch problem to generate a new MAAP/ATO within one hour of target development. The cost is estimated at \$25 million per year, with tests every 18 to 24 months until the objective is achieved.

- This stretch problem differs from some others in that it is tied to a specific operational process. Thus, its form appears most like a standard program—but with some essential features to drive innovation from the commercial world. One critical element is the necessity for an open architecture, to enable participation by a variety of performers. An open architecture also supports the development of modular mission capabilities, so that the mission-planning scope can be expanded over the course of the program. The program manager should spend the first year of execution defining the open architecture to be used within the program.
- Another essential feature is a program structured around a series of milestones that assign increasingly challenging time goals for ATO generation, *e.g.*, 12 hours, six hours, one hour. The program manager should apportion an expanding scope of mission capabilities to each milestone. Once the architecture is defined, the government should open a competition to fund at least two competing systems integrators, each leading an innovative team of subcontractors that bring diverse and leading technology expertise. Periodically (*e.g.*, every 18 months) the sponsor should run a head-to-head competition of performance against the current milestone goals. The winning team should be assigned a prize, possibly in the form of an award fee. Then each prime should be allowed to reselect a new team of subcontractors, based on performance to date and emerging technologies.
- The program should carry out the assessments at interim milestones using a combination of live, virtual, and constructive facilities. The graduation exercise should be carried out at Red

Defense Science Board Summer Study on Autonomy

Flag, to enforce performance assessment in a realistic combat environment that mimics the fully complexity of modern operations—fighter interdiction, attack/strike, air superiority, enemy air defense suppression, airlift, air refueling, and reconnaissance missions.

Summary

The study concluded that autonomy will deliver substantial operational value—in multiple dimensions—across an increasingly broad spectrum of DoD missions, but the DoD must move more rapidly to realize this value. Allies and adversaries alike also have access to increasingly rapid technological advances occurring globally, which are driven by commercial market forces stemming from a diverse array of commercial markets. While difficult to quantify, the study concluded that autonomy, fueled by advances in artificial intelligence, has attained a “tipping point” in value, and that autonomous capabilities are increasingly ubiquitous.

Over-arching themes that emerged from the study included:

- The need to build trust in autonomous systems while also improving the trustworthiness of autonomous capabilities
- The need to accelerate adoption of autonomous capabilities through DoD enterprise-wide enablers
- The need to strengthen the operational pull for autonomy by demonstrating operational value across a broad range of missions
- The need to expand the technology envelope to help the U.S. sustain military advantage through the increasing use of autonomy

While DoD is already embracing the value of autonomous capabilities, in both fielded systems and developmental programs, it has not yet adapted its enterprise processes to effectively support the rapid and widespread adoption warranted by the potential benefits—and made imperative by the potential perils of autonomy in the hands of adversaries. The study therefore concluded that action on the enterprise-level recommendations is of far greater importance—and urgency—than the implementation of any single program. These interdependent recommendations focus on the enablers needed to accelerate adoption of autonomous capabilities.

An important objective of this study was to identify opportunities for DoD to more rapidly exploit ongoing technological advances. By selecting several demonstrations of autonomous systems with near-term benefits, the study intends to illustrate the operational value across a diverse array of missions, thereby strengthening the operational pull for autonomous capabilities. It should be noted, however, that the full value of such demonstration programs will be realized only if they are conducted in concert with—and used to refine and mature—the recommendations focusing on the enterprise enablers.

The study also observed that DoD has research efforts underway that will, over time, expand its envelope of technological options. The study’s recommendations, which focus on a set of stretch problems, are intended as a supplement—not a replacement—for such ongoing research. While focusing primarily on expanding the technology envelope, execution of the stretch problems as designed yields broader benefits. By engaging a broad array of non-DoD providers, together with

Defense Science Board Summer Study on Autonomy

the diverse spectrum of DoD stakeholders, the stretch problems are intended to foster relationships that not only accelerate innovation but also accelerate DoD’s ability to exploit that innovation.

At its core, autonomy is about decision-making. The working definition used during this study was “autonomy results from delegation of a decision to an entity that is authorized to take action within specific boundaries.” As used in this report, the “entity” to which decision authority is delegated is a software algorithm. A key benefit is that the use of autonomy can increase decision speed—enabling the U.S. to act inside an adversary’s operations cycle.

But speed is equally important in a second dimension—rapid transition of autonomy into warfighting capabilities in order for the U.S. to sustain military advantage. And this dimension requires a DoD enterprise that is both ready and eager to realize the benefits of autonomy across its entire mission set.

Table 4 summarizes the recommendations in the report. Details for implementation of each recommendation are found throughout the report on the pages listed in the table.

Table 4: Summary of Recommendations

No.		Page
Accelerating Adoption of Autonomous Capabilities		
1	USD(AT&L) should require that best practices be developed and applied to all software dominated systems and, in particular, autonomous systems.	28
2	USD(AT&L) should address the special issues associated with cyber resiliency in autonomous systems.	30
3	DOT&E in conjunction with DT&E should establish a new T&E paradigm for testing software that learns and adapts.	34
4	The DoD test and evaluation community should establish a new paradigm for T&E of autonomous systems that encompasses the entire system lifecycle.	34
5	USD(AT&L) should require the acquisition community to establish and implement a consistent and comprehensive M&S strategy throughout the lifecycle of the system.	37
6	Military Service Chiefs should integrate technology insertion, doctrine, and CONOPs by ensuring early experimentation that uses alternative sources and informs employment doctrine.	38
7	USD(P&R), working with USD(AT&L) and Military Service Chiefs, should develop an autonomy-literate workforce.	40
8	ASD(R&E) should improve global autonomy technology discovery by encouraging personnel exchanges and coordinating partner organization efforts in FFRDCs, UARCS, and the IC.	41
9	The Deputy Secretary of Defense should establish departmental governance of autonomy by creating an EXCOM for oversight and establishing advocates in the Military Services.	43
10	USD(AT&L), USD(P), and ASD(PA) should take a proactive, two-pronged approach to anticipate cultural objections to the use of autonomy.	43

Defense Science Board Summer Study on Autonomy

11	The Deputy Secretary of Defense should take immediate action to counter adversary autonomy.	46
Strengthen Operational Pull for Autonomy		
12	NSA, in partnership with DARPA and IARPA, should fully develop the means to tip and cue DISA and the defense industrial base to defend the DoD information infrastructure, extending to U.S. government and private sector support as appropriate.	52
13	DARPA, working with AFRL and the 711th Human Performance Wing, should initiate a new program to adapt existing ISR data screening and fusion tools, such as the Air Force’s Dynamic Time Critical Warfighting Capability (DTCWC) or PCPAD-X, or DARPA’s Insight for autonomous, real-time use.	54
14	DIA and USSOCOM should integrate commercial components and build a new machine-learning analysis tool, and prototype the resulting system using existing historical data, seized media, and commercial (collateral) sources.	55
15	CERDEC, AFRL, and SPAWAR should develop Military Service prototypes for local, agile spectrum deconfliction and control among a few systems; concurrently DARPA should develop an architectural framework and algorithms for near-real time, theater-level spectrum deconfliction and control for a full ensemble of joint, coalition systems.	58
16	The Navy PEO for Littoral Combat Ships should conduct a user operational evaluation system program run by PEO-LCS in partnership with ONR.	60
17	USCYBERCOM should take the lead to develop an automated cyber-response, in partnership with CIA, FBI, NSA, DARPA (Plan X), key cyber-security industry players, and DISA.	63
18	U.S. Navy and DARPA should collaborate to conduct an experiment in which assets are deployed to create a minefield of autonomous lethal UUVs.	66
19	The U.S. Marine Corps, DARPA, ONR Code 30, and an FFRDC or UARC develop and experiment with a prototype heterogeneous, autonomous UAS support team that includes ten or more UA.	70
20	NAVSUP should demonstrate the use of modern intelligent adaptive planning in conjunction with SAP.	77
21	CASCOM, Ft Lee 10th Mountain Division, and the Ft Drum Joint Readiness Training Center should develop and deploy adaptive logistics decision support for a relocatable robotic warehouse and trained personnel in preparation for rapid deployment to unstable regions.	78
Expand Technology Envelope for Autonomous Systems		
22	DARPA should initiate a stretch problem designed to create a system that autonomously, globally, and in real-time identifies the causal linkages behind emerging social movements, and helps leaders understand the impact of possible courses of action along various possible future event trajectories.	84
23	ASA(ALT), with close participation by ARCIC, should establish an annual “swarm games” challenge.	90
24	DARPA should develop autonomous systems that detect large-scale intrusions on the IoT, by passively and remotely monitoring bulk network traffic, and identifying aggregate indicators of compromise hidden within the flood of ordinary traffic.	94

Defense Science Board Summer Study on Autonomy

25	DARPA should implement a stretch problem to demonstrate autonomous cyber-resilient systems (ACRS) for autonomous military vehicles.	96
26	The AFRL/RIS office should undertake a stretch problem to generate a new MAAP/ATO within one hour of target development.	99

Terms of Reference



ACQUISITION,
TECHNOLOGY
AND LOGISTICS

THE UNDER SECRETARY OF DEFENSE
3010 DEFENSE PENTAGON
WASHINGTON, DC 20301-3010

NOV 17 2014

MEMORANDUM FOR CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: Terms of Reference – Defense Science Board 2015 Summer Study on Autonomy

The technology of autonomy is rapidly advancing and finding widespread private sector and public sector application. Relevant capabilities span the spectrum from autonomy, i.e. the brains, to autonomous systems (e.g. robots, drones, etc.) which integrate autonomy into physical systems. Applications include IBM's Watson, the use of robotics and automation in ports and mines, autonomous vehicles (from UAVs to Google's self-driving car), automated logistics and supply chain management, and many more.

The purpose of this study is to identify the science, engineering, and policy problems that must be solved to permit greater operational use of autonomy across all warfighting domains. The study will assess opportunities for DoD to enhance mission efficiency, shrink life-cycle costs, and reduce loss of life through the use of autonomy. Emphasis will be given to exploration of the bounds—both technological and social—that limit the use of autonomy across a wide range of military operations. The study will ask questions such as: What activities cannot today be performed autonomously? When is human intervention required? What limits the use of autonomy? How might we overcome those limits and expand the use of autonomy in the near term as well as over the next 2 decades?

Applications to be considered include decision aids, planning systems, logistics, surveillance, and war-fighting capabilities. The study will also identify cost-imposing opportunities such as the use of autonomy to spoof adversaries, creating confusion and consuming their resources; and will also consider potential threats stemming from the use of autonomy by adversaries.

The study will examine the international landscape, identifying key players (both commercial and government), relevant applications, and investment trends. Considerations will include "baked-in" security, scalability, and variable cooperation between autonomous algorithms/systems and humans.

The study will consider opportunities such as: the use of large numbers of simple, low-cost (i.e. "disposable") objects vs. small numbers of complex (multi-functional) objects; use of "downloadable" functionality (e.g. apps) to repurpose basic platforms; and an ability to vary the degree of autonomy vs. human supervision/control for specific missions rather than developing mission-specific platforms.

The study will deliver a plan that identifies the barriers to increased operational use of autonomy and ways to reduce or eliminate those barriers. The study report should include: recommendations to reduce or eliminate the barrier, an assessment of risks to successful

implementation of the recommendation, and an estimate of resources required to implement the recommendation.

I will sponsor the study. Dr. Ruth A. David and Dr. Paul D. Nielsen will serve as Co-chairmen of the study. Dr. Jonathan Bornstein, US Army Research Laboratory, will serve as Executive Secretary. Lt Col Michael Harvey, USAF, will serve as the DSB Secretariat Representative.

The study will operate in accordance with the provisions of P.L. 92-463, the "Federal Advisory Committee Act" and DoD Directive 5105.04, the DoD Federal Advisory Committee Management Program." It is not anticipated that this study will need to go into any "particular matters" within the meaning of title 18, United States Code, section 208, nor will it cause any member to be placed in the position of action as a procurement official.



Frank Kendall

Members of the Study

Study Chairs

Dr. Ruth David	Private Consultant
Maj Gen Paul Nielsen, USAF (Ret.)	Software Engineering Institute

Executive Secretaries

Dr. Jonathan Bornstein	U.S. Army Research Lab
Lt Col Scott Heritsch	Joint Staff J8

Members

Mr. James R. Allard	Amazon
Dr. Amy E. Alving	Private Consultant
Dr. Michael R. Anastasio	Los Alamos National Laboratory
Dr. Brent Appleby	Draper Laboratory
Dr. Wanda M. Austin	Aerospace Corporation
Mr. Michael Bayer	Private Consultant
Dr. Jeffrey M. Bradshaw	Institute for Human and Machine Cognition
Mr. Frank J. Cappuccio	Private Consultant
Mr. James F. Carlini	Private Consultant
Gen Michael P. C. Carns, USAF (Ret.)	Private Consultant
Dr. Arup K. Chakraborty	Massachusetts Institute of Technology
Dr. Mark Chevillet	Applied Physics Laboratory
Hon. David S.C. Chu	Institute for Defense Analyses
Ms. Victoria Coleman	Technicolor
Mr. Christopher W. Day	Packet Forensics
Dr. Eric D. Evans	MIT Lincoln Laboratory
ADM William J. Fallon, US Navy (Ret.)	Countertack, Inc.
Dr. Craig Fields	Private Consultant
CAPT James R. Gosler, US Navy (Ret.)	Applied Physics Laboratory
Mr. Alfred Grasso	MITRE Corporation
Dr. S.K. Gupta	University of Maryland
CAPT Mark R. Hagerott, US Navy (Ret.)	U.S. Naval Academy
Mr. Paul J. (Page) Hoeper	Onpoint Technologies, Inc.
Dr. Adele Howe	Colorado State University
Brig Gen Chris Inglis, USAF (Ret.)	U.S. Naval Academy
Dr. Miriam E. John	Lawrence Livermore National Laboratory
Hon. Anita K. Jones	University of Virginia
Hon. Paul G. Kaminski	Technovation, Inc.
Dr. Ronald L. Kerber	Private Consultant
Gen Paul Kern, US Army (Ret.)	Cohen Group

Defense Science Board Summer Study on Autonomy

Dr. Amy A. Kruse	Cubic Global Defense, Inc.
Hon. Zachary J. Lemnios	IBM
Mr. Ashley J. Llorens	Applied Physics Laboratory
Dr. Robert Mandelbaum	Lockheed
Dr. John L. Manferdelli	Google
Dr. Joseph Markowitz	Private Consultant
Dr. Mark T. Maybury	MITRE
Gen James P. McCarthy, USAF (Ret.)	U.S. Air Force Academy
Hon. Judith Miller	Private Consultant
Hon. James N. Miller, Jr.	Applied Physics Laboratory
Dr. Robin R. Murphy	Texas A&M University
Mr. Robert F. Nesbit	Private Consultant
Mr. David H. Scheidt	Applied Physics Laboratory
Hon. William Schneider, Jr.	International Planning Services, Inc.
Dr. Ralph D. Semmel	Applied Physics Laboratory
Dr. Ross D. Shachter	Stanford University
Mr. James D. Shields	Private Consultant
Mr. Robert M. Stein	Private Consultant
VADM Edward M. Straw, US Navy (Ret.)	Private Consultant
Dr. James A. Tegnalia	University of New Mexico
Dr. Anthony J. Tether	Sequoia Group
Mr. David M. Van Buren	L-3
Mr. Vincent Vitto	Private Consultant
Mr. Lewis Von Thae	DynCorp International
Dr. David A. Whelan	Boeing
Dr. Robert L. Wisnieff	IBM
Dr. Greg L. Zacharias	U.S. Air Force

Defense Science Board

Mr. David Jakubek	Executive Director
LtCol Michael Harvey	Deputy for Operations, US Air Force
Mr. Robert Ramsay	
Mr. Tom Simms	

Government Advisers

Mr. Maynard Holliday	ODUSD(AT&L)
Dr. Greg Hudas	U.S. Army
Mr. Scott Littlefield	Tactical Technologies Office, DARPA
Dr. James Overholt	Air Force Research Laboratory
Dr. Brian Sadler	Army Research Laboratory
Dr. Frederick E. (Fritz) Schultz	OASD(R&E)
Mr. Donald Sofge	Naval Research Laboratory

Defense Science Board Summer Study on Autonomy

Dr. Marc Steinberg

Office of Naval Research

Staff

Ms. Rosemary Battles

Strategic Analysis, Inc.

Ms. Amy Cauffman

Strategic Analysis, Inc.

Ms. Erin Erickson

Strategic Analysis, Inc.

Ms. Meghan Fitch

Strategic Analysis, Inc.

Ms. Hannah Freeman

Strategic Analysis, Inc.

Ms. Ashlee Gilligan

Strategic Analysis, Inc.

Ms. Sarena Harvey

Strategic Analysis, Inc.

Mr. Marcus Hawkins

Strategic Analysis, Inc.

Mr. Leland Lambert

Strategic Analysis, Inc.

Dr. Toni Maréchaux

Strategic Analysis, Inc.

Ms. Diane O'Neill

Strategic Analysis, Inc.

Ms. Jen Schimmenti

Strategic Analysis, Inc.

Ms. Jeray Simms

Strategic Analysis, Inc.

Ms. Stephanie Simonich

Strategic Analysis, Inc.

Ms. Melissa Smittle

Strategic Analysis, Inc.

Mr. Ted Stump

Strategic Analysis, Inc.

Briefers to the Study

February 25–26, 2015

Dr. Lynn Parker	National Science Foundation
CPT Matt Pregmon	National Defense University
Mr. Paul Scharre	Center for a New American Security
Dr. Brad Tousley	Tactical Technology Office, DARPA
Dr. Richard Voyles	White House Office of Science and Technology

March 17–18, 2015

Dr. Paul Cohen	Information Innovation Office, DARPA
Mr. Jim Geurts	U.S. Special Forces
Mr. Bob Work	Deputy Secretary of Defense
Mr. Eric Sundberg	Aerospace Corporation
LTC Matt Dooley	U.S. Army Capabilities Integration Center
Gen Mike Hostage	USAF Air Combat Command (retired)
Mr. Frank Kendall	Under Secretary of Defense (AT&L)

April 29–30, 2015

Dr. Guru Banavar	IBM
Dr. Steve Chien	NASA Jet Propulsion Laboratory
Dr. George Ka'iiliwai	U.S. Pacific Command

May 18–19, 2015

Mr. Jean-Charles Ledé	Tactical Technologies Office, DARPA
-----------------------	-------------------------------------

June 16–17, 2015

Mr. Carl Johnson	Northrop Grumman Electronic Systems
Dr. Gill Pratt	Defense Sciences Office, DARPA
Mr. Stephen Randich	FINRA
Dr. Peter W. Singer	New America Foundation

July 28–29, 2015

Mr. Michael Gilmore	DoD Operational Test and Evaluation
Mr. Jehezkel Grizim	Israel Aerospace Industries
Mr. Oskar Levander	Rolls Royce Marine
LTG H.R. McMaster	U.S. Army Capabilities Integration Center
Mr. Brian Pierce	Tactical Technologies Office, DARPA