NOT FOR PUBLICATION UNTIL RELEASED
BY THE ARMED SERVICES COMMITTEE,
SUBCOMMITTEE ON EMERGING THREATS AND CAPABILITIES
UNITED STATES HOUSE OF REPRESENTATIVES

DEPARTMENT OF THE AIR FORCE

PRESENTATION TO THE HOUSE ARMED SERVICES COMMITTEE
SUBCOMMITTEE ON EMERGING THREATS AND CAPABILITIES
U.S. HOUSE OF REPRESENTATIVES

SUBJECT: Advancing the Science and Acceptance of Autonomy for Future Defense Systems

STATEMENT OF: Dr. Greg L. Zacharias, SES
Chief Scientist of the Air Force

November 19, 2015

NOT FOR PUBLICATION UNTIL RELEASED
BY THE ARMED SERVICES COMMITTEE,
SUBCOMMITTEE ON EMERGING THREATS AND CAPABILITIES
UNITED STATES HOUSE OF REPRESENTATIVES

Introduction

Chairman Wilson, Ranking Member Langevin, Members of the Subcommittee and Staff, I am pleased to have the opportunity to provide testimony on how the Air Force is advancing the science and acceptance of autonomy for future defense systems.

The importance of autonomy to the Air Force and its future systems was emphasized in 2014 when the Secretary and Chief of Staff of the Air Force published *America's Air Force: A Call to the Future*. In laying out a vision for the next 30 years, the United States Air Force strategy places an emphasis on the development of autonomous systems as a means of achieving strategic advantage in future operations. This includes the development of sensors and data gathering technology that can provide the needed information for a system to better understand the operating environment and mission goals, the development and implementation of reasoning systems and software environments to assess situations and make recommendations or decisions, and the refinement of different means for effecting those recommendations/decisions, whether through direct action, such as guiding an unmanned platform, or through recommendations to another human or machine teammate. The overall goal is to enable systems to "react to their environment and perform more situational-dependent tasks as well as synchronized and integrated functions with other autonomous [human or machine] systems". The document also recognizes autonomy's promise to lower system manning costs while making human operators more effective in their missions, increase the range of operations by extending manned capabilities while simultaneously protecting them from harsh or dangerous environments, reducing the time required to conduct time-critical operations, and providing increased levels of operational reliability, persistence, and resilience.

This past summer, my office published the first volume of *Autonomous Horizons: System Autonomy in the Air Force, Volume I: Human Autonomy Teaming*. Authored by my predecessor, Dr. Mica Endsley, this report describes a vision for autonomous systems that work synergistically with our airmen as a part of an effective human-autonomy team in which goals are shared, common situation awareness is maintained, and mission tasking responsibility flows smoothly, simply, and seamlessly between them. In this vision of the future, autonomous systems will be designed to serve as a part of a collaborative team with airmen where flexible autonomy will support the control of missions, tasks, platforms, and subsystems, in a dynamic fashion, sharing information, control, and tasking across humans and other autonomous systems, at different mission levels and over multiple phases or times within a mission. As envisioned, these systems will be able to operate robustly and in a predictable fashion, in uncertain and dynamically changing environments.

During my tenure as Chief Scientist of the Air Force, I plan to continue exploring this technically challenging area with the publication of two follow-on volumes to *Autonomous Horizons*. Volume II will explore nascent or existing science and technology (S&T) areas that can contribute to the creation of a level of perceptual and computational intelligence needed to support the vision described in Volume I, and that can deal effectively with the challenges of uncertainty and variability in “real-world” environments that include countering adversarial activities. Volume III will address key infrastructure needs for autonomous system development, including cyber security/reliability, communications requirements and network vulnerability, and command and control structures and systems. Working in concert with the Air Force Research Lab, I expect these volumes to provide detailed road maps for the basic and applied research and technology demonstrations needed to advance autonomy at across all levels of the Air Force.

As you are aware, six major recommendations were made in the 2012 Defense Science Board (DSB) report on the role of autonomy in DoD systems. The first recommended using a three-facet autonomous systems framework to guide autonomous system development and acquisition, while the second recommended a coordinated autonomous systems S&T effort across the services, in four Tier I areas: Machine Perception, Reasoning and Intelligence (MPRI), Human/Autonomous System Interaction and Collaboration (HASIC), Scalable Teaming of Autonomous Systems (STAS), and Test, Evaluation, Validation, and Verification (TEVV).[1] Because my focus is on S&T, the majority of my comments will focus on this second recommendation, and AFRL's contribution to these areas. However, at the end of this testimony I will provide additional comments related to the other recommendations.

The Air Force's primary agent for autonomy research, Air Force Research Laboratory (AFRL), commissioned the development of the AFRL S&T Autonomy Vision and Strategy in 2013. This document identifies the major goals, technical challenges, and investment strategies needed to discover, develop, and demonstrate warfighter-relevant autonomy S&T to maintain and enhance air, space, and cyberspace dominance. This strategy has been coordinated with the other services and with OSD through the Assistant Secretary of Defense for Research and Engineering's (ASD(R&E)) Autonomy Community of Interest (COI). A result of this coordination, the DoD Autonomy Roadmap, links directly to the key enabling autonomy technologies identified above and to Recommendation 2 of the DSB report (MPRI, HASIC, STAS, and TEVV).

[1] Dr. Jon Bornstein (Autonomy COI Lead), "DoD Autonomy Roadmap Autonomy Community of Interest", National Defense Industrial Association (NDIA) 16th Annual Science & Engineering Technology Conference/Defense Tech Exposition, 24-26 March 2015

Air Force Autonomy Research Program

In its vision for 2020, AFRL seeks to enable the right balance of human and machine capability to meet Air Force challenges in the future by focusing on growing autonomous system capability, integrated with the human capacity to perform in a high-tempo, complex decision environment, and optimized in a fashion that allows humans to work with machines effectively, efficiently, predictably, and robustly. Simply stated, the AFRL Autonomy S&T vision is:

Intelligent machines seamlessly integrated with humans - maximizing mission performance in complex and contested environments

As machine capabilities advance, AFRL's approach is to keep the airman at the center of the system design to ensure the end result is effective teaming of the airman with an autonomous system to enable greater agility, effectiveness (lethal or non-lethal), and mission utility. Embedded in the vision are the following strategic goals which support the ability to do collaboration and teaming between types of autonomous systems, as well as adjuncts to manned systems: 1) highly effective human-agent teaming will harness a system's ability to digest data and the human's ability to deal with uncertainties to improve human-system performance; 2) multiple autonomous systems will actively coordinate their actions to achieve the mission intent; and 3) systems with enhanced intelligence and self-governance will enable operation in complex, contested environments that create challenges exceeding human (-alone) performance limits. Achievement of these goals will provide the Air Force with autonomous system capabilities that enable greater resiliency, capability, and versatility.

We now describe some of the AFRL autonomous systems research and development efforts, on-going or planned, organized according to the four Tier I areas identified in the DSB 2012 study and which comprise the OSD COI Autonomy.

Machine Perception, Reasoning and Intelligence (MPRI) Research

AFRL has three research areas that directly support the Machine Perception, Reasoning and Intelligence (MPRI) area: mature machine learning, situational understanding of the contested environment, and robust system self-protection. Advancements in computational intelligence underpin a machine's ability to perceive, assess, reason, plan, decide, and act.[2] As a result, AFRL is actively performing research in artificial intelligence, cognitive and computer science, data analytics, and machine and human learning. While much of DoD's focus is in supporting "autonomy in motion" (i.e., in robotics and unmanned platforms), AFRL is also actively supporting "autonomy at rest" needs in the intelligence, surveillance, and reconnaissance (ISR) world. AFRL's Human Effectiveness Directorate has an ISR Analyst Test Bed which provides a research-representative Processing, Exploitation and Dissemination (PED) cell for developing interfaces and technologies. Outputs of this research, the Internet Relay Chat Coordinate Extractor (ICE) and Enhanced Reporting Narrative Event Streaming Tool (ERNEST), not only improve manpower efficiencies and reduce airman workload, but also lay the groundwork for integrated multi-INT autonomous processing and advance analyst cuing via autonomous decision-aiding.

Since autonomous systems must have the ability to anticipate, sense, and respond to dynamically-changing threats, AFRL is also developing technologies that enable situational understanding of the contested environment. This demands an ability of machines to combine and fuse from sensor cues, embedded domain knowledge, and a maintained situational awareness, and to modify their actions in preparation for (or in response to) a wide range of kinetic and non-kinetic threats. This capability is especially critical when autonomous systems

---

[2] Nikita A. Visnevski and Mauricio Castillo-Effen. "A UAS Capability Description Framework: Reactive, Adaptive, and Cognition Capabilities in Robotics", IEEEAC Paper #1259, Version 4 Updated December 3, 2008.

are expected to detect and respond to adversarial manipulation and deception in contested environments.

In addition, the autonomous systems must operate safely and competently when communication links with humans or other autonomous systems are broken or jammed. To ensure effective operations in an adversarial environment, AFRL is pursuing assured, secure, robust communication, surveillance, and navigation technologies. Specific areas of research in this area include GPS degraded/denied navigation for weapons, laser communication, and robust satellite operation in the presence of anomalies and threats.

Human/Autonomous System Interaction and Collaboration (HASIC) Research

A goal of AFRL's autonomy program is to deliver flexible autonomy systems with highly effective human-machine teaming. The keys to maximizing the human-machine interaction are: instilling appropriate levels of confidence and trust among the team members[3]; understanding of each member's tasks, intentions, capabilities, and progress; and ensuring effective and timely communication needed to support effective teaming. All of which must be provided within a flexible architecture for autonomy that facilitates different levels of authority, control, and collaboration, along the lines identified by the "cognitive echelon" described in the 2012 DSB Autonomous Systems Report (Recommendation 1). Targets for reaching this goal range from near-term analysis of cyber operator stress and vigilance indicators, to more long-term aspirations of trusted autonomy where thousands of adaptive collaborating autonomous agents team with humans.

The current integration of an Auto Ground Collision Avoidance System (Auto GCAS) into the Air Force's operational F-16 fleet is an example of how the focus on human-machine

---

[3] *Appropriate* is the key qualifier here, since overtrust of an underperforming system can lead to misuse and mishaps, while undertrust of a competent system can lead to disuse and inefficiencies.

teaming and the need to develop trust across the team can build acceptance of autonomous systems within the Air Force. The system was developed jointly by five organizations working closely together: AFRL; Lockheed Martin's Advanced Development Programs (ADP), also known as the Skunk Works®; the Office of the Undersecretary for Personnel and Readiness; NASA's Armstrong Flight Research Center; and the Air Force Test Center. The system has already been credited with saving the lives of two F-16 pilots (and their aircraft), by preventing a tragic accident called controlled flight into terrain (CFIT), which happens when a pilot, in full control of his or her properly functioning aircraft, inadvertently flies the aircraft into the ground.[4] One of the keys to acceptance of this Auto GCAS technology was a two-year study sponsored by AFRL's Air Force Office of Scientific Research (AFOSR) and its Human Effectiveness Directorate.[5] As the Auto GCAS system neared implementation in 2014, the study worked with test pilots, engineers, managers, and operational pilots to identify lessons learned from technology development, to determine real-world perspectives on trust evolution, and to examine potential influences on trust. These results are now being used by a follow-on field study is currently underway for the Air Force Flight Test Center, the F-16 System Program Office Safety Office, the Air Combat Command Safety Office, the AFRL Aerospace Systems Directorate, and the F-35 Collision Avoidance Technology Program. This four-year study plans to examine pilot trust in Auto GCAS, and the evolution of trust with continued experience with the system. The study will also identify and document user experience, concerns, impact, and benefits of the technology as they emerge, to provide a basis for understanding critical trust issues with more complex autonomous systems to be fielded in the future.

---

[4] https://www.faasafety.gov/files/gslac/library/documents/2006/Oct/6583/AC%2061-134.pdf

[5] "Influence of Cultural, Organizational and Automation Factors on Human-Automation Trust: A Case Study of Auto-GCAS Engineers and Developmental History" in *Human-Computer Interaction. Applications and Services* (Volume 8512 of the series Lecture Notes in Computer Science), 2014.

Scalable Teaming of Autonomous Systems (STAS) Research

AFRL's Autonomy Research Program is focused on creating actively coordinated teams of multiple autonomous systems to achieve mission goals that are better carried out by larger teams of smaller disaggregated systems or platforms (commonly referred to as "swarms"). Advanced autonomous systems must be able, as individuals and as teams, to analyze their missions, goals, and responsibilities; and decompose them into actionable, individual tasks and functions. They must be able to dynamically organize into a team to effectively perform the mission tasks, efficiently allocate and use their collective resources in real-time, and communicate as necessary to replan and reallocate tasking as the mission or adversarial environment changes.

Air Force research in this area includes near-term goals of collaborative control of six or fewer autonomous platforms and space-to-space asset collaboration and cross queuing with more far-term goals of agile swarming weapons and large numbers of adaptive C2ISR collaborating agents. As an enabler to this research, AFRL is currently collecting proposals for a Low Cost Attritable Strike Unmanned Aerial System (UAS) Demonstration that will design, develop, assemble, and test a technical baseline for a high speed, long range, low cost, limited-life strike Unmanned Aerial System (UAS). The program will also identify key enabling technologies for future low cost attritable aircraft demonstrations, and provide a vehicle for future capability and technology demonstrations. The goal of this program is to establish a benchmark, concluding in a flight demonstration that will test the bounds of what can be accomplished in a short time to establish a baseline system cost against a notional set of strike vehicle requirements.

Test, Evaluation, Validation, and Verification (TEVV) Research

Test and evaluation and validation and verification (TEVV) is vital to maximizing the operational gains of advanced autonomy, but provides significant challenges over more traditional systems because of four key discriminators: 1) "high-functioning" autonomous systems will necessarily be complex software systems, capable of demonstrating unanticipated "emergent" behavior, and will suffer the key validation and verification (V&V) issue associated with such systems, notably an inability to exhaustively test operation under all foreseeable conditions; 2) autonomous systems interact with their environment, via sensors, effectors, and communications links, and the explicit specification of all combinations of the system inputs/outputs and environmental variables is combinatorially impossible; 3) autonomous systems will interact with human teammates at some level, so that real-time operator in- or on-the-loop experimentation will be necessary, thus limiting the dimensionality of the test space because of time/resource constraints; and 4) sufficiently high-functioning systems will evolve into systems that learn over time and scenario exposure, so that behaviors will change over time, necessitating retest and possible recertification as a function of system "experience." As a result, the TEVV working group of the DoD Autonomy COI, following the recommendations of the 2015 DSB Autonomy Study has adopted a "continuous/life cycle approach" to T&E.

Additionally, research is underway on alternative means of autonomous agent licensure similar to pilot or operator licensure vs. traditional certification approaches for non-autonomous components. AFRL is also developing methods for verification early in the design cycle (e.g., to formally validating requirements for consistency and completeness). Finally, AFRL is developing run-time validation techniques to ensure the system is operating safely and effectively during use, even if design-time verification is not complete due to system learning.

Air Force Response to 2012 DSB Report on the Role of Autonomy in DoD Systems

I now return to the recommendations made in the 2012 Defense Science Board (DSB) report on the role of autonomy in DoD systems, and the follow-on efforts made by the Air Force in that regard.

The first recommendation proposed the use of a three-facet autonomous systems framework to guide autonomous system development and acquisition. This is illustrated in Figure 1, and is comprised of three views: a cognitive echelon; mission dynamics; and trade space. In general, Air Force S&T efforts in this area espouse this framework, if not explicitly acknowledging its specific structure and nomenclature. In particular, the *cognitive echelon* view is used to parse many of the autonomous system development efforts, ranging from mission-level ("at rest") planning aids to help develop a mission, to section-level ("in motion") capabilities needed for collaboration of multiple autonomous platforms, to vehicle-level ("in motion") systems like Auto GCAS. The *mission dynamics* view is also embraced by AFRL efforts, since many of the efforts focus on individual portions of the mission (e.g., pre-mission planning, in-mission flight management, post-mission ISR processing). In addition, researchers have embraced the notion of flexible levels of autonomy (as described in *Autonomous Horizons Volume I*), in which those levels will vary, and task responsibility will change over mission phases, depending on a number of factors, such as operator workload, system competence, task tempo, etc. Finally, the *trade space* view can be viewed as simply good systems engineering, where systems attributes, such as optimality vs. robustness, are traded off in an iterative design-prototype-experiment-evaluate fashion to arrive at a compromise solution that accounts for wide variety of factors and characteristics that constitute a well-designed system that can be operated effectively and maintained efficiently. While this is normally the purview of the Air Force's acquisition arm, the S&T community is well aware of the constraints that go into developing a successful weapons system through the use of an overarching system engineering process which includes the trade space view noted in the DSB report. In short, AFRL embraces the three-facet autonomous systems framework recommended by the 2012 DSB study.

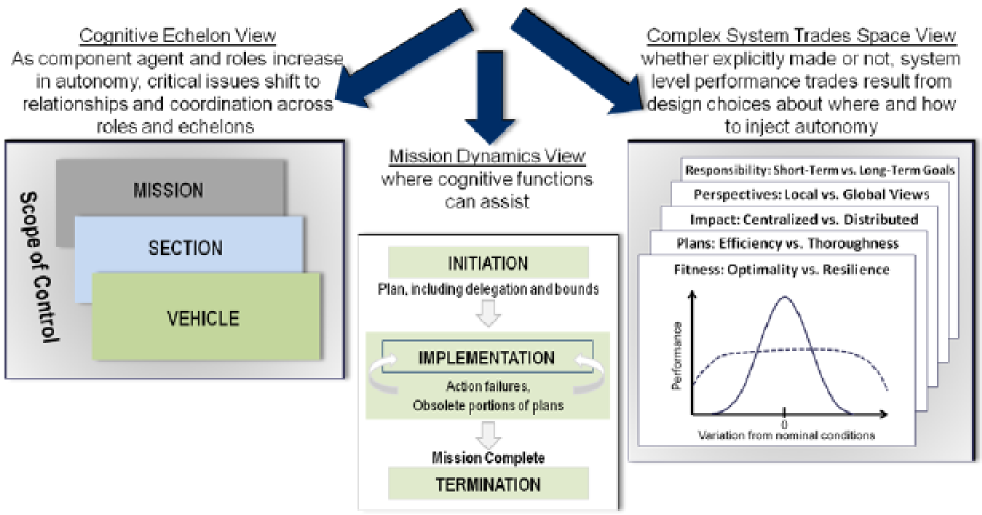

Figure 1: Framework for the Design and Evaluation of Autonomous Systems

The second recommendation proposed a coordinated S&T program guided by feedback from operational experience and evolving mission requirements. The Air Force is contributing to this in three ways. First, the Air Force has been and continues to be a strong and active member of the OSD Autonomous Systems COI and its four Tier I areas described earlier. The AFRL's current Chief Technology Officer, Dr. Morley Stone, was a previous Autonomy Community of Interest (COI) Lead and was instrumental in building a coordinated DoD roadmap for Autonomy research. AFRL members currently co-lead three of four of the COI's working groups. Second, as described earlier in the bulk of my testimony, AFRL, via its 2013 S&T Autonomy Vision and Strategy document, has explicitly laid out how it will support the major technology development recommendations made in the DSB Report and explicitly formulated in the Tier I areas of the Autonomy COI. Finally, in late 2013, under ASD(R&E)'s Reliance 21 initiative, DoD allocated

$15 million per year for three years to competitively fund autonomy projects run by the DoD laboratories. The Autonomy Research Pilot Initiative (ARPI) program selected seven projects, one of which is led by AFRL and is supported by the other services. The initiative's focus is on research and development for instantiating an "Intelligent Multi-UxV Planner with Adaptive Collaborative/ Control Technologies (IMPACT)" by combining flexible play-calling, bi-directional intuitive human-autonomy interaction, cooperative control algorithms, intelligent agent reasoning and machine learning technologies to enable cooperative multi-air/ground/sea unmanned vehicle missions.

The third recommendation proposed that acquisition programs emphasize the separation of the software from the platform hardware and the use of open software architectures. As these are ongoing efforts within the acquisition branch of the Air Force and not direct S&T efforts, I will refrain from commenting, except to note that the S&T community is aware of these acquisition goals and will support them as needed.

The fourth recommendation focused on addressing the unique challenges posed by autonomous systems for the developmental and operational test communities. As noted above, the Air Force S&T community acknowledges these challenges, and is directly supporting the Test, Evaluation, Validation, and Verification (TEVV) Autonomy COI Working Group, across multiple fronts: exploring concepts for "continuous V&V" throughout a system's lifetime as missions change and system experience grows, developing methods for model-based simulation and verification at all stages of development, acquisition, testing, and training ("the "digital thread"), and co-development of concepts of operations via extensive man-machine teaming.

The fifth recommendation focuses on using lessons learned based on the use of autonomous systems in recent conflicts, in professional military education, war games, exercises and operational training. While not directly the purview of Air Force S&T, AFRL has a long history of working with remotely piloted aircraft (RPA) operators to address shortcomings in existing systems, particularly the Ground Control Stations of Predator and Reaper. This work has led to the development of new concepts for command and control of these and future systems like them. In addition, the S&T community's engagement in Future Games 2014 wargaming efforts (with operators) led directly to the publication of *Autonomous Horizons Volume I*.

Finally, the sixth recommendation encouraged the Defense Intelligence Agency (DIA) and the Intelligence Community to track adversarial capabilities with autonomous systems and the Services to include these threats in war games, training, simulations and exercises. As these are ongoing efforts within the Intelligence Community, I will refrain from commenting.

Conclusion

In conclusion, I hope that I have been able to convey to you the Air Force's progress in advancing the science and acceptance of autonomy in our future systems and the close coordination of our research portfolio with the other services through OSD's Autonomy COI. In addition, I am confident that I have demonstrated that the Air Force S&T enterprise has been responsive to recommendations from the 2012 DSB Autonomy Study.

On behalf of the dedicated scientists and engineers of the Air Force, thank you for the opportunity to testify today and for your continued support of United States Air Force science and technology. I look forward to answering any questions you may have.