MP130118

MITRE TECHNICAL REPORT

Anticipating the Onset of Autonomy: A Survey of the DoD, Armed Service, and other Federal Agencies' Outlook on Autonomy

March 2013

Dr. Bob Grabowski
Jessica Rajkowski

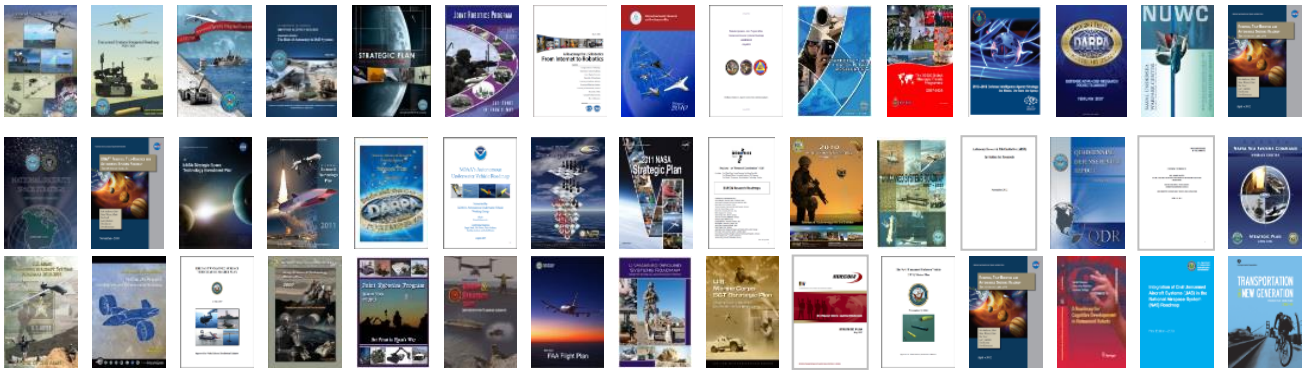

Sponsor: MITRE Sponsored Research
Dept. No.: J82A **Project No.:** 25MSR621-EA
Public Release: 15-1708

The views, opinions and/or findings contained in this report are those of The MITRE Corporation and should not be construed as an official Government position, policy, or decision, unless designated by other documentation.

This document has been approved for public release. Distribution Unlimited

MITRE

Abstract

Autonomy has become a recent focus of attention in each of the military services, the broader Department of Defense community, and many Federal agencies. Even though a clear understanding and definition of autonomy remains elusive, stakeholders are making plans and setting expectations in their roadmaps, strategic plans and research objectives. This summary investigates two emerging perspectives on autonomy; the anticipation of new capabilities and the trepidations of allowing new technologies to make their own decisions. These views collectively express the breadth of stakeholder's views towards the onset of autonomy, what they believe it means, what they think it will offer, and some of the accompanying concerns over adopting this unproven but potentially game-changing capability. The views have been compiled from statements and comments in critical documents and roadmaps published by the DoD, Services, and Federal agencies.

Table of Contents

1 Introduction

In November of 2012, a Blackhawk helicopter glided between two mountain passes in a San Jose valley without any humans onboard. It hugged the ground at 200 feet while scanning the environment using a gimbaled laser range-finder to avoid obstacles, maintain height and constantly assess potential landing areas – all with limited human intervention. For two hours the aircraft navigated the jagged terrain, eventually finding a landing area and setting down on its own in a forest clearing[1].

On the other end of the country, DARPA (Defense Advanced Research Projects Agency) launched a new effort in highly mobile, humanoid machines to navigate rubble, open doors and operate tools and machinery[2]. In another part of DARPA, efforts moved towards construction of a completely unmanned, autonomous, anti-submarine warfare craft intended to stay on station for months and track unknown and potentially hostile submarines across the open seas, through shipping lanes and into harbors and ports[3].

Autonomy is coming and the capabilities it promises are profound. One day these systems will be released into the real world working in close proximity to humans and critical infrastructure. These agile machines will work with dismounted soldiers acting as members of the team rather than extended tools that have to be controlled. They will defend bases and camps and shuttle people and cargo. They will detect and track anomalous or dangerous behavior and engage directly with the enemy. The capability to formulate their own awareness and make complex, goal-oriented decisions will be essential for these missions. This capability is loosely expressed as autonomy.

With the great promises of autonomy come elements of trepidation. If we imbue these systems with the ability to decide their own goals and act on them, what assurances do we have that they will do what we intend or that we can accept what they do? How do we know they won't become a liability during a battle or after the battle? Before these goal-directed machines are released in theater, or powered up near the general public, lawmakers and commanders want a better understanding of the benefits and risks of this new capability and some level of assurance that these machines will not go rogue. They want to trust that these systems will achieve their intended results but also that they will not become a new threat themselves. Trust extends from system themselves to the entire chain of development from designers to acquirers, to testers, and ultimately to the commanders who employ them.

This survey takes a closer look at how the Services and Department of Defense (DoD) are anticipating the onset of autonomy; what they believe it means, what they think it will offer and some of the accompanying concerns they voice over this potentially game-changing capability. It is primarily focused on the statements and comments extracted from critical documents published by the DoD, Federal agencies and the Services.

---

[1] tech.slashdot.org/story/12/12/06/0055206/army-tests-autonomous-black-hawk-helicopter

[2] www.darpa.mil/Our_Work/TTO/Programs/DARPA_Robotics_Challenge.aspx

[3] www.darpa.mil/Our_Work/TTO/Programs/Anti-Submarine_Warfare_(ASW)_Continuous_Trail_Unmanned_Vessel_(ACTUV).aspx

These critical documents include science and technology, strategic, research and master plans, as well as technology and investment roadmaps. As a collective, they represent the intentions of the DoD, Services and Federal agencies towards investment in autonomy and unmanned systems. While over 100 documents were reviewed for this survey, less than half contained discussions specifically towards unmanned systems and autonomy. Moreover, only the following critical documents specifically address and expound on autonomy in their writings:

- 2007, Navy Unmanned Surface Vehicle Master Plan
- 2009 DoD Unmanned Systems Integrated Roadmap
- 2009 Air Force Unmanned Aircraft Systems Flight Plan
- 2009 ARCIC Robotics Strategic White Paper
- 2009 DARPA Strategic Plan
- 2010 Air Force Technology Horizons
- 2010 Army Unmanned Air Systems Roadmap
- 2010 NASA, Robotics, Tele-Robotics and Autonomous Systems Roadmap
- 2011 DoD Science and Technology Priorities for Fiscal Years 2013-17 Planning
- 2011 DoD Unmanned Systems Integrated Roadmap
- 2011 RSJPO Unmanned Ground Systems Roadmap
- 2012 DSB The Role of Autonomy in DoD Systems
- 2012 DoD Autonomy Research Pilot Initiative
- 2012 DoD Autonomy in Weapon Systems Directive

2 Defining Autonomy

Autonomy as a distinct capability has only recently become a focal point among the Services and DoD. Prior to 2011, autonomy was seen as a general qualifier on the advancing capabilities of unmanned platforms or as the mechanism by which these systems advanced. However, it was not seen as a particular capability of a systems ability to perceive, reason, or plan.

As recently as 2009, the main focus among the Services was still on robotic platforms and the utility of unmanned systems. The Army's Robotics Strategic White Paper was the first of the Services to define a robot; "a man-made device capable of sensing, comprehending, and interacting with its environment [ARCIC]." This represents early views where the focus of future systems was centered on the capabilities of the physical platform and less on its ability to make local independent decisions and formulate plans. Beyond the recognition of increased autonomy to achieve many of the desired missions, there was limited discussion of autonomy as the focal point of research and little awareness on the need to examine it directly.

Later in 2009, the Unmanned Systems Integrated Roadmap released the first consolidated roadmap. While autonomy was again mentioned as a qualifier to robot platforms, it primarily focused on the physical capabilities of the wide and growing variety of formal acquisition and research platforms. The main focus remained on the capabilities of the platforms themselves not on their ability to reason.

In 2010, the Air Force's Technology Horizons declined an attempt to define autonomy directly but noted its key attributes "[autonomous systems will exhibit] the ability for complex decision making, including autonomous mission planning, and the ability to self-adapt as the environment in which the system is operating changes [TH]."

It wasn't until 2011 that autonomy became a specific focus of the DoD and Services when it was identified as one of the top seven DoD focus areas by the Secretary of Defense for Acquisition, Technology and Logistics (AT&L) [SDP]. Shortly afterward, the second release of the Integrated Systems Roadmap made the first concerted attempt to specify an understanding of autonomy. It made delineations between automated systems (the propensity of currently deployed unmanned systems) and autonomous systems.

As reported in the roadmap, "Automatic systems are fully preprogrammed and act repeatedly and independently of external influence or control [USIR]." They note that these are systems that may be self-regulating (such as speed or elevation control) and self-steering (such as adjusting rudders to maintain a GPS path) which perpetually minimize system error to achieve a plan, but cannot themselves define that plan. They go on to define autonomous systems as those that are "self-directed toward a goal in that they do not require outside control, but rather are governed by laws and strategies that direct their behavior [USIR]. The distinction here is the ability to interpret their environment and, as they put it, "determine what information is important [USIR]." They point out that early designs may be set by humans but that these systems may one day be able to learn their own behaviors to set plans and goals.

This delineation between autonomy and automation is enhanced even further in 2012 in the DoD Autonomy Research Pilot Initiative. They note that automated systems may function with little or no operator intervention but that they have "predetermined responses in reasonably well-known and structured environments" and that "system performance is limited to the specific actions it has been designed to do [ARPI]." The report enhances and expands the understanding of autonomy to include "a set of intelligence-based capabilities that allow it to respond within a bounded domain to situations that were not pre-programmed or anticipated in the design for operations in unstructured, dynamic, uncertain, and adversarial environments." They go on to stipulate that autonomous systems "must be adaptive to and/or learn from an ever-changing environment [ARPI]."

The DoD is not the only Federal agency focusing on autonomy. The National Aeronautics and Space Administration (NASA) is focusing on autonomy for both terrestrial and interstellar applications. In their 2010 Robotics, Tele-Robotics and Autonomous Systems Roadmap, NASA takes a more direct approach at defining autonomy. "Autonomy, in the context of a system (robotic, spacecraft, or aircraft), is the capability for the system to

operate independently from external control [NASA].” They go on to note that an autonomous system is defined as “a system that resolves choices on its own” fundamentally tying autonomy to the decision processes of a system [NASA].

For several years, Agency and Service working groups and steering committees debated the definition of autonomy and methods to define levels within it. In 2012, the DSB attempted to break the cycle of definitions; “The attempt to define autonomy has resulted in a waste of both time and money spent debating and reconciling different terms and may be contributing to fears of unbounded autonomy. The definitions have been unsatisfactory because they typically try to express autonomy as a widget or discrete component, rather than a capability of the larger system enabled by the integration of human and machine abilities [DSB].” Instead, the DSB took the discussion in two new directions. First, they looked at autonomy as a capability not a property; “Autonomy is better understood as a capability (or a set of capabilities) that enables the larger human-machine system to accomplish a given mission, rather than as a ‘black box’ that can be discussed separately from the vehicle and the mission [DSB].” Consequently, they defaulted to a more mundane definition, “For the purposes of this report, a capability that is delegated to the machine is considered autonomous.” Second, the DSB made the argument that autonomy must be separated from the platform. “It [autonomy] is also primarily a software endeavor, which is a shift from traditional hardware oriented, vehicle-centric development [DSB].”

Ultimately whether the community comes to a common understanding or not, autonomy will move forward. Whether it is a characteristic, a property or a capability, the true question is; how do we prepare for it? Autonomy represents great promise but also potentially great peril. As echoed by the Technology Horizons back in 2010, autonomy represents the “single greatest theme" for advancing warfighting capability into the future [TH].

3 Anticipation

While a formal definition is desirable, a more informative understanding of the impact of autonomy can be found by looking at what stakeholders expect from autonomy. These expectations reflect the nature of what the DoD and Services believe to be the properties, characteristics and capabilities of an autonomous system and give insights into the technical challenges that researchers will have to overcome to achieve them. These expectations can be broken into 3 general categories; reduced operator burden, tactical advantage and trusted companion.

- Reduction of Manpower: how autonomy is expected to reduce the manpower required in increasingly complex missions, especially those related to unmanned systems.
- Tactical Advantage: how autonomy is expected to increase the warfighters’ tactical advantage with its ability to provide greater competent standoff, access denied areas and operation for extended periods of time.

- Trusted Companion: how autonomy will enable increased interaction and cooperation that will complement the warfighter and eventually become an ally on the battlefield in direct contact with adversaries.

3.1 Reduction of Manpower

One of the most pragmatic and immediate advantages of autonomy is the reduction of manpower. That is, stakeholders are looking for ways to reduce all aspects of operations, manned and unmanned, from the workflow of the warfighter to the resources needed to support them.

3.1.1 Mitigation of Unmanned

In large part, the call for autonomy is a desire to rebalance the impact of unmanned systems themselves. As noted by the Unmanned Systems Integrated Roadmap, "The rapid proliferation of unmanned systems and the simultaneous operation of manned and unmanned systems as unmanned systems expand into additional roles have created a manpower burden on the Services [USIR]." This burden persists even today as reflected by the DSB, "We have compensated for the challenges of our UAVs with an extraordinary level of manning and sustainment investments and we need to move forward to meet the next challenges in our national security landscape [DSB]."

Unmanned systems have proven themselves to be invaluable additions to the warfighter but have come at a cost. Instead of initially reducing the logistical cost and cognitive load of the warfighter as expected, it has done the opposite. A single unmanned platform requires new logistics support and in many cases additional dedicated operators. For example, the Predator requires four operators at constant vigil to control a single vehicle. Even small UAV deployment requires special training and special operators. In addition to manpower, these highly managed UAVs put a large burden on resources such as communications spectrum and bandwidth. Competent local decision capability has the potential to greatly mitigate both the expertise to control these systems and the resources they require. As noted by the Army UAS roadmap, "Increased autonomy will significantly reduce operator workload, increase reliability and speed of mission performance [and] reduce demands upon bandwidth or allow more capability with the same bandwidth [AUAS]."

3.1.2 Reduction of Operators

Beyond regaining balance of current manning, autonomy is seen a strong driver to further reducing current manning. Reduction in manning and dedicated resources has a direct impact on reducing costs. As noted by the Technology Horizons, "dramatically increased use of autonomy offers potentially enormous increases in capabilities, and if implemented correctly can do so in ways that enable manpower efficiencies and cost reductions [TH]." This is echoed by the USIR; "introducing a greater degree of system autonomy will better enable one operator to control more than one unmanned system, and has the potential to significantly reduce the manpower burden [USIR]."

DARPA also recognizes the potential cost savings from advanced autonomy, "Cognitive systems [advanced autonomous systems] will give military commanders and their staffs better access to a wide array of rapidly changing information, reduce the need for skilled computer system administrators, and dramatically reduce the cost of system maintenance [DARPA]." They go on to make the critical link between cost savings and autonomy; "Without learning through experience or instruction, our systems will remain manpower-intensive and prone to repeat mistakes, and their performance will not improve [DARPA]."

With NASA's move towards unmanned systems for future exploration, it is not surprising that they make a direct correlation between autonomy and cost, identifying one of the three fundamental benefits of autonomy to be "cost savings via increased human labor efficiencies and reduced needs [NASA]."

This potential endures today as noted by the DSB, "there is significant potential for increased use of autonomy to have a dramatic impact on the manning requirements for unmanned systems. Manpower costs are a large part of the DoD budget and the fiscal constraints of the pending budget environment will provide a strong motivation to increase efficiencies and add capability to unmanned systems to free people for more critical purposes [DSB]."

This reduction of human burden is expected to continue for more complex autonomous systems even as the number of unmanned systems increases. The Air Force Flight Plan predicts the utility of autonomous swarms and the ability to be controlled by a single operator; "The near-term concept of swarming consists of a group of partially autonomous UAS operating in support of both manned and unmanned units in a battlefield while being monitored by a single operator. [AFFP]" While it would seem swarm technology is a far off concept; there are already efforts underway in Europe to control a convoy of high speed vehicles (platooning) in engineered highways [SARTE]. Teaming and eventually swarming are actively being investigated.

3.1.3 Information Filtering

In addition to the manpower operating the unmanned systems, further human attention is being dominated by the sheer volume of the data these systems collect. As noted by the Army UAS Roadmap, "The Army's ability to collect information far outpaces its ability to use the information collected [AUAS]." The number and duration of unmanned missions is resulting in unprecedented collections of data that have to be analyzed by humans. As noted by the Unmanned Systems Integrated Roadmap; "This challenge [overload of data] is not unique to the unmanned environment, but it has been exacerbated by the large numbers of ISR-capable, long-endurance unmanned systems being fielded [USIR]." Technology Horizons echoes the concerns of data overload; "Although humans today remain more capable than machines for many tasks, natural human capacities are becoming increasingly mismatched to the enormous data volumes, processing capabilities, and decision speeds that technologies offer or demand [TH]." The DSB notes "about a third of the staff required to support Air Force UAVs are devoted to processing sensor data and exploiting them to create useful information. Even with this staffing level, the rapid growth in data volume is making it very difficult to keep up. [DSB]."

Solutions to this growing problem will come from the ability to smartly process the data into consumable information. However, this capability will have to grow to where decisions are made about what is important and needs to be passed on. The USIR succinctly puts it; Autonomous systems will be able to "determine what information is important in making a decision [USIR]." The DSB echoes this observation; "Identifying more efficient ways of processing the increasing volume of data collected by various platforms will be essential to realizing the platforms' benefits [DSB]." They sum up by saying, "there are many opportunities to use autonomy capability to increase the capacity of the intelligence analysts assigned to the exploitation function [DSB]." They make the critical link to autonomy, "Autonomy has a role in advancing both collection and processing capabilities toward more efficient, integrated ends [DSB]."

3.2 Tactical Advantage

The tactical advantages that autonomy can offer are undeniable. Unmanned systems are already providing unprecedented standoff and access for the warfighter. Adding the ability to make local decisions increases the competence of these systems giving them the ability to operate in conditions where direct human supervision is not possible either because of increased operational tempo, extreme distance or environmental conditions, mission duration or mission complexity.

3.2.1 Faster Reaction Time

Machines can operate at cycle speeds that far surpass human response while accurately tracking large volumes of information. As noted by the Technology Horizons, "such advanced autonomous systems will be powerful force multipliers and enable operations at timescales far faster than possible with human in-the-loop control" and "the increased operational tempo that can be gained through greater use of autonomous systems itself represents a significant capability advantage. [TH]"

Adding the ability for machines to make informed local decisions (autonomy) will only increase this capability. NASA, in their Autonomous Systems Roadmap, has identified autonomy as being a critical component to operation in space where support from earth is greatly hampered due to communications latency; "Greater use of highly adaptable and variable autonomous systems and processes can provide significant time-domain operational advantages to robotic systems or crewed systems that are limited to human planning, decision, and data management speeds [NASA]."

3.2.2 Deeper Penetration

In addition to faster decisions, autonomy will enable deeper penetration into spaces and environments currently beyond the scope of existing capability. This includes operation in new and inaccessible domains such as deep water, high altitudes and inhospitable environments such as accident sites, disaster areas and adversary-held land.

The Navy in particular is looking to autonomy to enable operations in inaccessible environments; "The need for long-term independent operation is essential for the MS

[maritime security], ASW [anti-submarine warfare], and MCM [mine countermeasure] missions where the requirement exists to transit long distances, detect, assess, and avoid potential threats and collect information independent of direct human operators [NUSV]." They tie this capability directly to autonomy; "Autonomy offers the benefit of minimizing manning and bandwidth requirements while extending the tactical range of operations beyond the line of sight [NUSV]."

As forecasted by the USIR; "Autonomy can also enable operations beyond the reach of external control or where such control is extremely limited [USIR]." In the extreme, NASA is looking to autonomy to enable distant operations; "Autonomy should make human crews independent from Earth and robotic missions more capable [NASA]."

3.2.3 Extended Operation

This capability was recognized by the Unmanned System Integrated Roadmap, "Autonomy can help extend vehicle endurance by intelligently responding to the surrounding environmental conditions (e.g., exploit/avoid currents) and appropriately managing onboard sensors and processing (e.g., turn off sensors when not needed)."

DARPA has made a real emphasis on increased duration of unmanned vehicles. As they note in their Strategic Plan, advances in unmanned systems will "enable entirely new design concepts unlimited by the endurance and performance of human crews [DARPA]." For example, DARPA is currently funding a program, called Vulture, designed to keep a high altitude UAV aloft for over five years[4]. DARPA is also funding a project to enable an anti-submarine USV to remain on station for months to find and track submarines[5].

3.2.4 Agility and Adaptability

Technology Horizons recognizes a necessary driver for autonomous systems is agility and that they will "demand a shift from systems designed for fixed purposes or limited missions to ones that are inherently agile in their ability to be readily and usefully repurposed across a range of scenarios [TH]." The Unmanned Integrated System Roadmap makes a similar argument; "The special feature of an autonomous system is its ability to be goal-directed in unpredictable situations. This ability is a significant improvement in capability compared to the capabilities of automatic systems [USIR]." They go on to note that this capability will be enabled by the ability to learn; "While robustness in adaptability to environmental change is necessary, the future need is to be able to adapt and learn from the operational environment because every possible contingency cannot be programmed a priori [USIR]."

3.3 Trusted Companion

Perhaps the most anticipated (and most uncertain) employment of autonomy is the creation of competent unmanned allies on the battlefield. From systems that work behind the lines assisting in transport, logistics or protection, to companions that work side by side in direct

[4] http://www.darpa.mil/Our_Work/TTO/Programs/Vulture.aspx

[5] http://www.darpa.mil/Our_Work/TTO/Programs/Anti-Submarine_Warfare_(ASW)_Continuous_Trail_Unmanned_Vessel_(ACTUV).aspx

contact with the adversary, they could provide advice on assessing threat, operational cover from ambush or attack and coordinated battle maneuvers.

3.3.1 Faithful Servant

As articulated by the Unmanned Ground Systems Roadmap, autonomous systems may provide capabilities including "vehicles to serve as robotic 'mules' to take on multiple soldiers 'loads.'" In addition to labor intensive roles, these autonomous systems are expected to one day act as local sentries guarding bases or vessels. As noted by the DoD memorandum on Autonomous Weapons Systems; "Human-supervised autonomous weapon systems may be used to select and engage targets, with the exception of selecting humans as targets, for local defense to intercept attempted time-critical or saturation attacks for static defense of manned installations and onboard defense of manned platforms [AW]."

3.3.2 Loyal Wingman

As these autonomous systems become more competent, they will work their way into the battlefield working side by side in maneuvers with humans in direct contact with the adversary. The Air Force in particular is looking to the capabilities unlocked by smart collaboration; "Loyal wingman technology will accompany and work with a manned aircraft in the AOR (Area of Responsibility) to conduct ISR (Intelligence, Surveillance and Reconnaissance), air interdiction, attacks against adversary integrated air defense systems, offensive counter air missions, command and control of micro-UAS, and act as a weapons 'mule,' increasing the airborne weapons available to the shooter. This system is capable of self-defense, and is thus, a survivable platform even in medium to high threat environments [AFFP]. "

The Autonomy Research Planning Initiative notes; "As unmanned systems become more autonomous, machines (agents) will be delegated increased decision making authority by their human operators. In some operational settings, an autonomous system may act in a peer-to-peer relationship with the warfighter, needing to interact naturally with the warfighter at full operational tempo [ARPI]."

4 Trepidation

Complimentary to the anticipation of autonomy is concern about the complications or consequences of unmanaged autonomy. This section examines some of the concerns in the community regarding the adoption of new capabilities and systems that are not adequately understood or lack proper safeguards. Rather than looking at what is gained from employing these new autonomous, self-regulating, goal-seeking systems, this section looks closer at what tenuous confidences might be lost. That is, what well-established, well-practiced processes which typically instill confidence in high-tech solutions might be weakened or invalidated by this new technology:

- Unmanaged Complexity – concern of the exceedingly large, complex and coupled decision state space that makes behavior prediction, verification and validation all but impossible in autonomous systems
- Unintended Consequences – concern that autonomy will expose the warfighter to unintended consequences and unintended engagements due to unpredicted emergent behavior
- Asymmetry of Development – concern that adversaries will corrupt, infiltrate, or develop autonomy, unburdened by consequences of unstructured autonomy

4.1 Unmanaged Complexity

There is a significant and founded concern about the growing complexity of highly capable systems. Understanding and ultimately trust of these systems is derived from the ability to predict and track how a system behaves through the mired of states that it must traverse. Unfortunately, the combinatorial effects of inputs, outputs and state transition can result in systems with far too many states to explore methodically. As noted by the Technology Horizons this results in "the near-infinite state systems that result from high levels of adaptability [TH]." They go on to say this complexity makes them "inherently unverifiable by today's methods [TH]." NASA also recognizes this possibility; "large software projects have such complex software that exhaustive and manual exploration of all possible cases is not feasible [NASA]."

4.1.1 Coupling of Sub-systems

For well-structured systems, mitigation to the complexity problem has been decomposition. Each sub-system is tested independently from other systems with known and established inputs and outputs. However, autonomous systems are increasingly designed to interact more directly with the environment and actors, resulting in a deep concern about the interaction and dependency between sub-systems. The Unmanned Systems Integrated Roadmap reflects this concern; "failures often occur at the interfaces between system elements, in many cases, between interfaces thought to be separate. The exponential trends in software and network communications increasingly mean that many elements of a system can now affect one another [USIR]." They go on to say, "as systems get much of their functionality from software and multisystem interactions, complexity is no longer separate and distinct [USIR]."

This coupling has profound impact on methodologies currently being employed for testing and certification. One mechanism to reduce testing cost and complexity is to test sub-systems separately and disconnected from the remainder of the system. This complexity mitigation is reduced with coupled systems. Changes in one part of the system may change behaviors in another. As noted by the Unmanned Systems Integrated Roadmap, "Unmanned systems raise new issues of artificial intelligence, communications, autonomy, interoperability, propulsion and power, and manned-unmanned (MUM) teaming that will challenge current T&E capabilities. These problems will get more serious as systems become more interactive and more automated [USIR]."

This concern of interdependencies has led the DoD to release requirements specifically targeted towards the impact of changes in the operation of autonomous systems software; "regression testing shall identify any new operating states and changes in the state transition matrix of the autonomous or semi-autonomous weapon system [AW]."

4.1.2 Unstructured Adaptation

This concern is amplified where learning and adaptation may occur. In these cases, large portions of code cannot be verified and tested because it has not yet been formulated. Not only will adaptable systems dynamically adjust weighting between importance of data and tasking, but they will start to write the code that governs them as well. As pointed out by the Unmanned Systems Integrated Roadmap, "while robustness in adaptability to environmental change is necessary, the future need is to be able to adapt and learn from the operational environment because every possible contingency cannot be programmed a priori. [USIR]"

4.1.3 Brittleness

The flip side of adaptability is brittleness. Brittleness is a property where complex systems behave correctly and appropriately under nominal conditions but fail catastrophically under less than optimal conditions. As reflected by the Unmanned Systems Integrated Roadmap, "because artificial systems lack the human ability to step outside a problem and independently reevaluate a novel situation based on commander's intent, algorithms that are extremely proficient at finding optimal solutions for specific problems may fail, and fail badly, when faced with situations other than the ones for which they were programmed [USIR]." Brittleness is not a conscious choice (no one chooses to be brittle) but rather a failure to account for the inherent complexities when operating in a dynamic, unstructured environment. As noted by the Defense Science Board, "Current designs of autonomous systems, and current design methods for increasing autonomy, can create brittle platforms, and have led to missed opportunities and new system failure modes when new capabilities are deployed [DSB]." They go on to report, "Brittle autonomous technologies result in unintended consequences and unnecessary performance trade-offs, and this brittleness, which is resident in many current designs, has severely retarded the potential benefits that could be obtained by using advances in autonomy [DSB]." Brittleness results from assumptions and tradeoffs that are made during the design process in order to reach a desired functionality. The concern is whether these tradeoffs are sufficiently understood and consequently whether they are worth the risk.

4.2 Unintended Consequences

The most feared (and least understood) trepidation that accompanies autonomy is the unintended consequence. Warfighers have always been willing to incur risk to gain tactical advantage as long as that risk is known and can be managed. However, autonomy represents a new capability that is not fully understood. The benefits of employing machines with autonomy must outweigh the risks both at the systems level but also at the mission and theater level. The concern is not the magnitude of the consequence but the ability to predict and assess them. This concern was noted by the Defense Science Board, "For commanders, a

key challenge presented by the complexity of software is that the design space and tradeoffs for incorporating autonomy into a mission are not well understood and can result in unintended operational consequences."

4.2.1 Ambiguity in Command

An autonomous system, by definition, makes decisions on its own. The consequences of introducing this independent entity into the military, where it may not conform to the long-standing command structure of the Services, are unknown. Questions will arise about where an autonomous system fits in the command structure, who can command an autonomous system, and can the system give commands of its own. As noted by the Army UAS Roadmap "The technology implications on military structure, particularly as unmanned systems become increasingly autonomous will be significant [AUAS]." This is echoed by the Unmanned Systems Integrated Roadmap, "Advances in autonomy at the system level must proceed with awareness of potential disadvantages and vigilance for unintended consequences, which may include diminished command over parts of the forces structure. [USIR]."

The Defense Science Board has acknowledged the need to rethink CONOPS and procedures, stating that "the Task Force urges caution against falling into the 'Substitution Myth' by trying to replace humans with autonomous systems without considering how machines change work patterns, responsibilities and training." Further, they state that "when operational training of human-autonomous system teams begins, it is likely that new top-level requirements or changes to the CONOPS will be identified that will improve future teams." As noted by the Unmanned Systems Integrated Roadmap, "surrendering decision trust to a software-based and self-learning design outside the context of specific operations is a matter of high rigor, and must be examined in the context of organizational-unit and theater CONOPS. [USIR]"

4.2.2 Safety and Liability

The Ground Roadmap states that "safety is a critical concern and one of the most significant issues autonomous vehicles must overcome before they can be widely accepted and fielded. Currently, autonomous vehicles operate within restricted areas in which all operators are fully aware of the vehicles‟ limitations. Before operation in crowded urban environments can happen, advancements in navigation and sensor fusion algorithms are needed to allow robots to distinguish humans from other objects, negotiate complex terrain, and operate."

There are real concerns about the liability, both economic and political, of the release of machines capable of formulating their own actions outside of the normal command structure. Unintended engagements can have long-term ramifications. The Navy Surface Vehicle Roadmap recognized the accompanying legal concerns of unattended platforms, "legal issues will continue to drive the general rules of engagement for weaponized USVs [unmanned surface vehicles]" and "drive much of the technical and operational problem associated with weapon release, autonomous or man-in-loop [NUSV]."

This concern has caused the DoD to directly call for and establish "guidelines designed to minimize the probability and consequences of failures in autonomous and semi-autonomous weapon systems that could lead to unintended engagements [AW]." This requirement was further elaborated; "Systems will go through rigorous hardware and software V&V and realistic system developmental and operational T&E, including analysis of unanticipated emergent behavior resulting from the effects of complex operational environments on autonomous or semi-autonomous systems [AW]."

Early failure, especially in with a public audience can have long lasting effects including placing overly restrictive regulations on the use of autonomy. This is particularly a concern for the automotive industry. As noted by the Center for Automotive Research, "the ramifications of an early autonomous or connected-vehicle traffic crash could be calamitous. Bad publicity is a significant risk for the deployment of innovative automotive technology, even if the technology itself is not the cause [CAR]." With the pervasiveness of media and reporting, this caution is of equal concern to military leaders.

4.2.3 Ethics of Weaponization

Concerns for liability go beyond damage and casualties. This is echoed today in the DSB report, "In addition to technical limitations and vulnerabilities, UxVs are operationally hampered by doctrinal and cultural issues [DSB]." Employment of autonomous machines that may make their own decisions on the application of force is a concern at the highest levels of warfare. The employment and application of autonomous, thinking machines becomes an ethical concern especially as these systems are weaponized. These concerns need to be addressed before they become a point of contention. As articulated by the Air Force Flight Plan, "ethical discussions and policy decisions must take place in the near term in order to guide the development of future UAS capabilities, rather than allowing the development to take its own path apart from this critical guidance [AFFP]." Recently the DoD has released new directives on the use of autonomized weapons systems, "autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force [AW]." While ethics is certainly not a direct technical issue, design and innovation can significantly support or mitigate the evolving landscape and ease public concerns.

4.2.4 Complacency

Autonomous systems hold great promise in reducing the cognitive load of operators from command and control to information exploitation. However, there is an awareness of the need for the human to stay "on the loop." As noted by the Air Force Flight Plan, "the trust required for increased autonomy of systems will be developed incrementally [with] humans monitoring the execution of operations and retaining the ability to override the system or change the level of autonomy instantaneously during the mission [AFFP]." As humans become comfortable with autonomy, they will increasingly delegate authority and action to these systems. Unfortunately the requirement of human oversight and override capability is at odds with systems that which take over the dull and mundane. Ironically, the very nature of autonomy may inadvertently lead to a reduction in vigil. Autonomy allows systems to run

for greater periods of time regardless of how dull the task may be. This has the potential to result in oversight or complacency. The Unmanned System Integrated Roadmap recognizes this danger and cautions, "implementing autonomy can lead to a loss of human attention to vital oversight in matters having potentially dangerous or lethal consequences [USIR]." Ironically, part of the solution may be in augmenting human attention with autonomy on the oversight task as well giving the monitoring systems capability to assess potential danger and alert operators – essentially re-sensitizing them to important activities.

4.3 Smart Adversary

Most of the early concerns for autonomous systems have centered about the competence of the system to achieve its goals and the ability to adequately predict and control its consequences. However, there is a growing awareness of the concern of these systems as they come in contact with or are developed by a smart adversary. As noted by the DSB, "The use of autonomous UxVs [by adversaries] may be the next "knowable" capability surprise [DSB]." This concern spans the entire spectrum of contact from corruption and loss of control to asymmetry in advancement.

4.3.1 Corruption

Autonomy may represent a technology that once released, causes more problems than it is worth. While few fear that machines will take over, as depicted in fiction entertainment like the Terminator series, we are very aware of the effects of malicious software in the form of viruses and malware such as Trojan horses. These small software programs can wreak havoc on computer networks and systems, and could have similar consequences for autonomous systems. The potential results of a compromised system are that they can behave in unexpected or malicious ways. As reflected by the Defense Science Board; "key external vulnerability drivers for unmanned systems include communication links, cyber threats and lack of self-defense [DSB]." As these systems become more complex, their vulnerabilities become more complex and venues of attack harder to predict. This concern may be compounded when dealing with systems designed to adapt and learn if that learning can be used to hide or defend corruption.

Corruption extends to the supply chain as well. Without active measures, adversaries can exploit vulnerabilities during the design process by tampering with the manufacturing process either to cripple functionality or install backdoors around safeguards. As noted by the DSB, "The dependence on commercial information technology hardware (processors, etc.) also exposes the UxV to the cyber vulnerabilities of the global supply chain."

4.3.2 Loss of Control

Beyond corruption, there is the beginning of concern about the enemy gaining control of an autonomous asset. Now instead of simply contending with a system that is behaving abnormally, warfighters would have to contend with a competent, adaptable system with malicious intent. As foreshadowed by the DSB, "another, serious emerging vulnerability is from all forms of cyber attacks—from denial of service to taking over C2 of the actual platforms [DSB]." This concern has led the DoD to require that autonomous systems are

"sufficiently robust to minimize failures that could lead to unintended engagements or to loss of control of the system to unauthorized parties [AW]." In addition to contending with the enemy, commanders would have to respond to their own competent and capable technologies.

4.3.3 The Aware Adversary

Designing systems capable of executing complex maneuvers in a dynamic environment is itself incredibly challenging and will require significant advances in autonomy. The problem becomes even more uncertain when these systems come in contact with an aware adversary. How will rules of engagement change when an adversary becomes aware that they are contending with an autonomous system? Vigilance will then have to shift towards assuring an adversary cannot use gaps in autonomy against itself. For example, if a submarine captain becomes aware that he is being followed by an autonomous surface ship, can he use this knowledge to lure the system into a crowded shipping lane in order to cause an incident?

The recent directive on autonomous weapons systems hints at the need to consider adversary reaction requiring that "operators and commanders understand the functioning, capabilities, and limitations of a system's autonomy in realistic operational conditions, including as a result of possible adversary actions [AW]" and as they add later "against adaptive adversaries [AW]."

4.3.4 Asymmetry of Development

It would be naive to think that the United States and its allies are the only organizations investigating autonomy. Moreover, ubiquitous access to online resources that include source code and advanced manufacturing services has considerably lowered the barrier to entry. Complete systems can be developed, procured and assembled with little or no infrastructure requirements. Couple these resources with a malevolent intent and a higher tolerance of uncertainty, autonomy could represent a possible asymmetrical threat. This threat was first voiced in Technology Horizons; "the relative ease with which autonomous systems can be developed, in contrast to the burden of developing certifiable V&V methods, creates an asymmetric advantage to adversaries who may field such systems without any requirement for certifiability; countering this asymmetry will require access to as-yet undeveloped methods for establishing certifiably reliable V&V [TH]." They continue to point out that adversaries may "gain potential capability advantages that we deny ourselves." The lesson is that we cannot take an overly cautious approach to the advancement of autonomy but must move out in a methodical and disciplined fashion.

5 Conclusion

Excitement in the DoD community is growing as autonomous capabilities edge us closer to systems that can independently perceive their environment, develop complex internal understandings, reason on that understanding and formulate appropriate plans to prosecute its intended mission. Autonomy promises to reduce the burden of the warfighter both on cost and cognitive load. It promises to provide tactical advantages from faster decisions to

extended operation and access to currently denied areas. If realized, autonomy promises to enable faithful servants and eventually battlefield wingmen. However, this new capability comes with its own perils. Can we manage their complexity? Can we anticipate or even bound their behavior to mitigate unintended consequences? Can we guarantee control and oversight when engaging with adversaries? Part of the concern is embedded in our understanding of what we mean by autonomy. Not its definition but the many different aspects that are bound to it. Ultimately the question comes down to whether autonomy simply represents an evolution of increasingly complex and competent automation or will it require a change in the way we develop and adopt it?

6 References **(Critical Documents, roadmaps, strategic plans, directives)**

1. [NUSV] U.S. Navy, 2007, "Unmanned Surface Vehicle Master Plan"
2. [ARCIC] U.S. Army, 2009, ARCIC "Robotics Strategy White Paper"
3. [AFFP] U.S. Air Force, 2009," Unmanned Aircraft Systems Flight Plan 2009-2047"
4. [DARPA] Defense Advanced Research Agency, 2009, "Strategic Plan"
5. [USIR09] Department of Defense, 2009, "Unmanned Systems Integrated Roadmap FY 2009-2034."
6. [AUAS] U.S. Army, 2010, "Unmanned Aircraft Systems Roadmap 2010-2035"
7. [TH] U.S. Air Force, 2010, "Technology Horizons," Volume 1; AF/ST-TR-10-01
8. [NASA] National Aeronautics and Space Administration, 2010, "Robotics, Tele-Robotics and Autonomous Systems Roadmap"
9. [SDP] Department of Defense, 2011, "Science and Technology Priorities for Fiscal Years 2013-17 Planning", DoD Memorandum
10. [USIR] Department of Defense, 2011, "Unmanned Systems Integrated Roadmap FY 2011-2036."
11. [RSJPO] Robotics Systems Joint Project Office, 2011, "Unmanned Ground Systems Roadmap"
12. [DSB] Defense Science Board, 2012, "Task Force Report: The Role of Autonomy in DoD Systems."
13. [AW] Department of Defense, 2012, "Autonomy in Weapons Systems," DoD Directive 3000.09
14. [ARPI] Department of Defense, 2012, "Autonomy Research Pilot Initiative"

7 References (Other Roadmaps – references to but no clarity on autonomy)

1. U.S. Navy, 2000, "Unmanned Undersea Vehicle (UUV) Master Plan"
2. U.S. Navy, 2004, "Unmanned Undersea Vehicle (UUV) Master Plan"
3. U.S. Navy, 2008, "Naval Sea Systems Command Strategic Plan"
4. National Oceanic and Atmospheric Administration, 2009, "Autonomous Underwater Vehicle Roadmap"
5. Academia, 2009, "A Roadmap for US Robotics: Roadmap for U.S. Robotics From Internet to Robotics"
6. Department of Defense, 2010, "Strategic Plan for DoD Test and Evaluation Resources"
7. National Science and Technology Council, 2010, "National Aeronautics Research and Development Plan"
8. U.S. Marines 2011, "Vision and Strategy 2025"
9. U.S. Air Force, 2011, "Science and Technology Plan"
10. U.S. Navy, 2011, "Naval Air Warfare Center Aircraft Division Strategic Plan"
11. U.S. Navy, 2011, "Naval S&T Strategic Plan"
12. Joint Planning and Development Office, 2011, "NextGen Avionics Roadmap"
13. Joint Planning and Development Office, 2012, "NextGen Unmanned Aircraft Systems Research, Development and Demonstration Roadmap"
14. U.S. Army, 2012, "RDECOM Strategic Plan"
15. National Aeronautics and Space Administration, 2012, "Space Technology Roadmaps and Priorities"
16. U.S. Navy, 2012, ONR Code 31, "Science and Technology Strategic Plan"
17. Joint Planning and Development Office, 2012, "NextGen UAS Research, Development and Demonstration Roadmap"
18. National Aeronautics and Space Administration, 2013, "Strategic Space Technology Investment Plan"
19. Academia, 2013, "A Roadmap for U.S. Robotics: From Internet to Robotics"

8 References (Additional)

20. [NRS] Nevada Revised Statute 482A http://leg.state.nv.us/NRS/NRS-482A.html#NRS482ASec050
21. [CAR] KPMG and Center for Automotive Research, “Self-driving cars: The next revolution”
22. [SARTE] Bergenhem, C., Huang, Q., Benmimoun, A., Robinson, T. (2010) “Challenges of Platooning on Public Motorways”, proceedings of the 17th ITS World Congress, Busan 2010