



Center for Strategic and Budgetary Assessments

DELIVERING ADVANCED UNMANNED AUTONOMOUS SYSTEMS AND ARTIFICIAL INTELLIGENCE FOR NAVAL SUPERIORITY

THE CASE FOR ESTABLISHING A U.S. NAVY AUTONOMY PROJECT OFFICE



SHARIF H. CALFEE

DELIVERING ADVANCED UNMANNED AUTONOMOUS SYSTEMS AND ARTIFICIAL INTELLIGENCE FOR NAVAL SUPERIORITY

THE CASE FOR ESTABLISHING A U.S. NAVY AUTONOMY PROJECT OFFICE

SHARIF H. CALFEE



Center for Strategic and Budgetary Assessments

2021

ABOUT THE CENTER FOR STRATEGIC AND BUDGETARY ASSESSMENTS (CSBA)

The Center for Strategic and Budgetary Assessments is an independent, nonpartisan policy research institute established to promote innovative thinking and debate about national security strategy and investment options. CSBA's analysis focuses on key questions related to existing and emerging threats to U.S. national security, and its goal is to enable policymakers to make informed decisions on matters of strategy, security policy, and resource allocation.

ABOUT THE AUTHOR

Sharif Calfee is the U.S. Navy Fellow to the Center for Strategic and Budgetary Assessments as part of the Department of Defense Federal Executive Fellowship Program. Afloat, CAPT Calfee is currently serving as Commanding Officer, USS SHILOH (CG 67), homeported and forward deployed in Yokosuka, Japan. He also previously served as Commanding Officer, USS McCAMPBELL (DDG 85), also homeported in Yokosuka, Japan. He has previously served in six shipboard assignments across cruisers, destroyers and frigates. Ashore, he has served as the executive assistant to the Commander, Naval Surface Forces and as a Political-Military Action Officer for counterterrorism security assistance issues on The Joint Staff, J-5 Strategic Plans and Policy Directorate. He is the recipient of the Naval Postgraduate School's (NPS) George Philips Academic Excellence Award, a graduate with distinction from the U.S. Naval War College Joint Professional Military Education program and a graduate of the Joint & Combined Warfighting School at the Joint Forces Staff College (JFSC) where he received the Transformation Writing Award (Fall 2010). He received a 2002-2003 U.S. Navy Space & Naval Warfare Systems Center, Pacific Fellowship for his artificial intelligence research while at NPS in Monterey, CA. He has published articles and research on various military, national security and scientific topics to include in U.S. Naval Institute's Proceedings, Joint Force Quarterly and Naval Engineers Journal.

CAPT Calfee holds an M.S. in Computer Science (Artificial Intelligence concentration) from the U.S. Naval Postgraduate School, a Master's in Public Policy from Princeton University, Princeton School of Public and International Affairs and a B.S. in Computer Science from the U.S. Naval Academy.

ACKNOWLEDGMENTS

The author would like to thank the staff at CSBA for their support and efforts on this study. Additionally, he would like to thank all of the organizations and experts who participated in this study and offered exceptional feedback during the drafting, review, and editing process.

The analysis and findings presented here are solely the responsibility of the author and do not represent those of the Department of Defense (DoD) or the U.S. Navy. CSBA receives funding from a broad and diverse group of contributors, including private foundations, government agencies, and corporations. A complete list of these organizations can be found on our website at www.csbaonline.org/about/contributors.

Cover Graphic: NAVAIRSYSCOM & Northrop Grumman X-47B UCAV; DARPA and the Office of Naval Research's autonomous unmanned surface vehicle, Sea Hunter, underway for at-sea testing and operations; and Metron's Remus 6000 UUV, developed for the U.S. Navy

Contents

EXECUTIVE SUMMARY	i
CHAPTER 1: UNMANNED SYSTEMS RESEARCH STUDY	1
CHAPTER 2: HYPER-FOCUS / STARVATION AMONG KEY UXS R&D LINES OF EFFORT	5
Mission Autonomy	6
Test & Evaluation	10
CHAPTER 3: IMPROVING ORGANIZATIONAL STRATEGY	21
Lack of awareness of key organizational centers of expertise within the R&D ecosystem ..	22
Lack of awareness of complementary R&D efforts and resident expertise at other organizations within the R&D ecosystem	22
Bureaucratic, Administrative, and Risk-Aversion Impediments to R&D	25
The R&D ecosystem is unable to nimbly leverage emerging technologies	27
Better collaboration is needed between Fleet operators and the R&D ecosystem	30
CHAPTER 4: ORGANIZING TO IMPROVE NAVY UXS DEVELOPMENT	33
CHAPTER 5: TOWARD AN AUTONOMY PROJECT OFFICE	37
Autonomy Project Office Organizational Design	39
CHAPTER 6: THE CASE FOR AN AUTONOMY PROJECT OFFICE: THE ADVANTAGES	47
Clear, coherent authority and unity of effort	47
Solidify established UxS gains and make them permanent	48
Bypass institutional resistance and barriers to UxS development	48
An APO can help eliminate impediments to R&D technology transition for prototyping	48
An APO can provide AI R&D innovation and technology for manned platforms & C5I systems ..	49
An APO can more nimbly pivot towards harnessing revolutionary technological breakthroughs ..	49
An APO can significantly increase the pace of experimentation and advance	50
An APO can deliver on effective Operationalization of UxS to the Fleet through Unity of Effort in DOTMLPF development	52
An APO can serve as a service component AI/Autonomy office to the DoD's Joint AI Center ..	53
CHAPTER 7: CONCLUSION	55
APPENDIX	57
LIST OF ACRONYMS	59

FIGURES

FIGURE 1: AUTONOMY PROJECT OFFICE PROPOSED ORGANIZATION	40
FIGURE 2: ENGAGEMENT ACTIVITIES WITH UXS R&D ECOSYSTEM	58

TABLES

TABLE 1: APO DIVISION TITLES & ROLES (PART 1)	42
TABLE 2: APO DIVISION TITLES & ROLES (PART 2)	45
TABLE 3: UNMANNED SYSTEMS R&D LINES OF EFFORT DEFINITIONS	57

Executive Summary

The Center for Strategic and Budgetary Assessments (CSBA) conducted a comprehensive research study from September 2017 to January 2019 in which it engaged a significant portion of the Department of the Navy's research and development (R&D) ecosystem. Specifically, we interviewed and visited over 145 subject matter experts and more than 50 organizations pursuing work in all-domain unmanned autonomous systems (UxS) and artificial intelligence (AI). These organizations included Federally Funded Research & Development Centers (FFRDCs), University Affiliated Research Centers (UARCs), industry, academia, think tanks & independent organizations, Navy and DoD Labs/warfare centers, Navy fleet and operational commands, and Navy policy & research funding offices. These visits enabled a sense of how well this R&D ecosystem (the ecosystem defined as the totality of all the above organizations working to advance autonomy and AI technology for the Navy for UxS, including unmanned vehicles) was performing in this effort. At the individual, organizational level, these organizations house significant talent and engage in innovative research at the cutting edge of a wide range of disciplines and technologies that could maintain the U.S. Navy's technological advantage.

However, the Navy's current UxS R&D construct has an opportunity to continue improving the organization of this effort to further expand and leverage its recent efforts. Although the Navy service ethos embraces decentralization, this report will make the case that a robust and centralized effort is necessary to achieve capability developments that leverage commercial and academic sector advances in autonomous systems. It is within this challenging environment that the U.S. Navy is competing to deliver unmanned autonomous vehicles across all physical domains. The technological advantage the United States currently enjoys over peer/near-peer adversaries in AI and other key UxS supporting technologies may be eroding due to structures that create barriers to innovation, hampering USN efforts to develop, preserve, or extend advantages in unmanned systems. Consequently, the resultant friction is degrading the ability to rapidly invent, innovate, and prototype. These disadvantages make UxS R&D successes challenging to achieve without applying herculean effort and tremendous senior leadership engagement.

In recent years, the U.S. Navy has made notable progress in the organization and pursuit of unmanned vehicles and their critical enabling technology, AI. However, to outpace potential adversaries, the Navy must build upon and accelerate those efforts to not only maintain but increase its technical advantage. The Department of the Navy has published a signed Strategic Roadmap for Unmanned Systems to develop the strategy and vision for deliberate R&D and procurement of unmanned vehicles. However, these plans could benefit from a comprehensive reorganization of the UxS R&D bureaucracy to transition the Navy to an optimal track for success. Without this restructuring, the existing structure may be insufficient for the task, hampering UxS technological progress. The following symptoms of this predicament were identified:

- Resource hyper-focus on some R&D Lines of Effort (LOE) and starvation on other critical LOEs;
- Lack of awareness of key organizational centers of expertise within R&D ecosystem;
- Bureaucratic, administrative, and risk-aversion impediments to R&D;
- Inconsistent processes to evaluate UxS R&D and S&T capability demonstrations for prototyping;
- R&D ecosystem inability to nimbly leverage emerging technologies; and
- Insufficient collaboration between Fleet operators and the R&D ecosystem.

With the surge of R&D in UxS and AI by peer competitors, including for military applications, now is the time for the U.S. Navy to comprehensively reorganize its efforts to ensure its continued technical advantage and outpacing of potential adversaries. Given the potential implications of autonomous systems and the barriers identified in this report, one option for the Navy is to make organizational changes similar to those it has successfully implemented in the past to develop and field advanced systems, such as nuclear reactors, submarine-launched ballistic missiles, and the Aegis Weapon System. In all three cases, the Navy created robust, cross-functional, interdisciplinary organizations consisting of personnel from the military, civilian government service, industry, and academia that were given a broad, strong mandate and authority to research, develop, prototype, and operationalize transformational strategic capabilities. This reorganization is best executed by establishing a dedicated multi-domain Autonomy Project Office (APO) focused on the advancement and delivery of UxS operational prototypes for experimentation, testing, and ultimately operationalization and acquisition into the Fleet. By establishing an APO, the Navy can build unity of direction and effort across a wide spectrum of organizations, staffs, and commands that currently play varying roles in this endeavor. It is recommended the main mission of the APO should be to unify oversight, authority and direction of all the R&D lines of effort across the UxS ecosystem to continuously spiral the latest technological advances into an assembly line of autonomous capability for unmanned air, surface, and undersea vehicle prototypes that can be demonstrated and employed for warfighting value. An APO with the

delivery of UxS prototypes as its primary mandate, can significantly reduce bureaucratic friction, increase component and system commonality, accelerate momentum and truly harness the tremendous innovation and talent of engineers, scientists, and sailors, thereby quickening the Navy's technological progress.

The APO does not need to duplicate the excellent innovation and existing R&D, science, engineering, systems integration, and T&E organizations that already exist within the UxS R&D ecosystem. Instead, the APO will exist to better facilitate orchestration and collaboration among them to produce the unity of effort/direction needed to more rapidly deliver UxS prototypes. It must consist of scientists, engineers, researchers, experts, and managers from all corners of the UxS R&D ecosystem, including Navy and DoD Labs/warfare centers, FFRDCs, UARCs, industry, and academia. Commensurate with its accountability and responsibility, the APO must be invested with the authority to direct the ecosystem in all matters pertaining to UxS, autonomous and AI R&D, prototyping, and T&E, to include funding alignments and allocation. The APO's leadership and staff must demonstrate a deft, light touch whereby they facilitate the appropriate coordination among the ecosystem's constituent parts to achieve increased alignment without concurrently strangling innovation. There are some definite advantages to establishing an APO that will accrue to the Navy's benefit to prototype and operationalize unmanned autonomous vehicles. These include:

- Clear, coherent authority and unity of direction;
- Consolidation and extension of established UxS gains;
- More efficient harnessing of revolutionary technological breakthroughs;
- Significant increases of the pace of experimentation and technology advance; and
- More effective operationalization of UxS to the Fleet.

The Navy should establish the APO to better deliver on technologies for UxS, including autonomy, robotics, and AI, and avoid further eroding the technological advantage upon which U.S. maritime superiority depends. An APO with the delivery of UxS prototypes as its mandate can reduce friction, accelerate R&D momentum and fully leverage the talent of UxS researchers within the ecosystem.

CHAPTER 1

Unmanned Systems Research Study

How can the U.S. Navy best leverage Artificial Intelligence (AI) and autonomous systems? This study, which is based on a comprehensive survey of unmanned autonomous systems (UxS) and AI research and development (R&D) ecosystem — attempts to answer this question. In this report, the ecosystem is defined as the totality of all the above organizations working to advance autonomy and AI technology for the Navy for UxS. Interviews with over 145 subject matter experts and engagements with more than 50 organizations indicate that the current R&D efforts are pursuing cutting-edge research across a wide range of disciplines. Nevertheless, due to several organizational barriers, these efforts are unlikely to generate and sustain a technological lead for the U.S. Navy absent significant reforms.¹ The initial results of this study were reported out in the U.S. Naval Institute's Proceedings, and this study report will explore the subject in depth.² Additionally, this report will explore how the U.S. Navy can more effectively pursue unmanned vehicles, UxS, and other AI systems as different from the Navy's Project Overmatch mission, which is to develop a Naval Battle Network and Naval Operational Architecture.³

1 To date over 145 subject matter experts and 50 plus organizations representing Federally Funded Research & Development Centers (FFRDCs), University Affiliated Research Centers (UARCs), industry, academia, think tanks & independent organizations, Navy and DoD Labs/warfare centers, Navy fleet and operational commands and Navy policy & research funding offices were interviewed and visited across the country. These visits allowed for the development of a sense for how well this R&D ecosystem (ecosystem defined as being the totality of all the above organizations working to advance UxS and AI technology for the Navy) was performing in this effort.

2 Sharif H. Calfee, "The Navy Needs an Autonomy Project Office," *Proceedings*, December 2018, pp. 32-38.

3 Megan Eckstein, "Navy Focused on Strengthening Networks to Support Unmanned Operations," *USNI News*, October 27, 2020, <https://news.usni.org/2020/10/27/navy-focused-on-strengthening-networks-to-support-unmanned-operations>; and Mallory Shelbourne, "Navy's 'Project Overmatch' Structure Aims to Accelerate Creatign Naval Battle Network," *USNI News*, October 29, 2020, available at <https://news.usni.org/2020/10/29/navys-project-overmatch-structure-aims-to-accelerate-creating-naval-battle-network>

As the Belfer Center Study on Artificial Intelligence and National Security asserted:

AI technology is likely to be a transformative military technology, on a par with the invention of aircraft and nuclear weapons...Though machine learning and AI technology are comparatively young, human and organizational responses to the new technology are often echoes of prior experiences.⁴

Based on this research effort, this report identifies seven Lines of Effort (LOEs) assessed as vital to the development and prototyping of UxS that could be employed in the air, surface, and undersea warfare regimes: Navigation Control, Mission Autonomy, Energy/Endurance, Sensors, Communications, Payloads and Testing & Evaluation (Verification & Validation, Certification/Acceptance, and Operationalization).

In recent years, the U.S. Navy has made progress in UxS technological development, requirements definition, organizational change, and strategy development. These are exemplified by the following accomplishments: Extra Large Unmanned Undersea Vehicle (XLUUV) request for proposal (RFP) release and contract assignment;⁵ Sea Hunter Medium Displacement Unmanned Surface Vehicle (MDUSV) transition from the Defense Advanced Research Projects Agency (DARPA) to the Office of Naval Research (ONR); the MQ-25 Stingray RFP release;⁶ Knifefish Mine Hunting System contractor demonstration;⁷ rebranding of Program Executive Office (PEO) Littoral Combat Ships to PEO Unmanned and Small Surface Combatants;⁸ acceleration of the Snakehead Large Displacement UUV project;⁹ publishing of guidance to accelerate acquisition for rapid development, demonstration and fielding;¹⁰ and

4 Timothy Moy, *War Machines: Transforming Technologies in the U.S. Military, 1920–1940* (College Station, TX: Texas A&M University Press, 2001), p. 4.

5 The Extra Large Unmanned Undersea Vehicle (XLUUV) is the U.S. Navy's acquisition program to develop autonomous, globally operational UUVs. Lee Hudson, "Navy awards two contracts for XLUUV," *Inside Defense*, September 28, 2017; Valerie Insinna, "Navy to Kick Off Extra Large UUV Competition This Month," *Defense News*, January 10 2017; and Lee Hudson, "Navy will use Orca XLUUV acquisition strategy as model for future programs," *Inside Defense*, January 17, 2018.

6 The Medium Displacement Unmanned Surface Vehicle (MDUSV) of which the DARPA/ONR Sea Hunter is one example, is one of the U.S. Navy's USV prototype development programs. The MQ-25 Stingray is the U.S. Navy's carrier-launched autonomous drone designed to conduct aerial refueling of manned aircraft. Pat Host, "US Navy quietly releases MQ-25 Stingray unmanned aerial refueler RFP," *IHS Jane's Defence Weekly*, October 12 2017; and Samantha Masunaga, "Competition to build the Navy's MQ-25 flying tanker shows how drone fighters are taking on new roles," *Los Angeles Times*, April 25, 2018, available at <https://www.latimes.com/business/la-fi-mq-25-competition-20180425-story.html>

7 The Knifefish is the U.S. Navy's autonomous undersea mine-hunting vehicle. Seapower Magazine, "Critical Test Phases Completed for U.S. Navy Knifefish Unmanned Undersea Vehicle," *Seapower Magazine*, June 28, 2014.

8 The Large Displacement Unmanned Undersea Vehicle (LDUUV) project is the U.S. Navy's rapid acquisition program to develop UUVs that can be launched from surface ships or submarines. James F. Geurts, *Establishment of Program Executive Office Unmanned and Small Surface Combatants*, (Arlington, VA: Office of the Assistant Secretary of the Navy for Research, Development, and Acquisition, March 13, 2018).

9 Megan Eckstein, "Navy Accelerating Work on 'Snakehead' Large Displacement Unmanned Underwater Vehicle," *USNI News*, April 4, 2017.

10 Department of the Navy, *Navy Accelerated Acquisition for the Rapid Development, Demonstration, and Fielding of Capabilities* (Arlington, VA: Department of the Navy, March 15, 2017).

the release of the Navy Unmanned Strategy Roadmap (which is transitioning to the Navy Unmanned Campaign Plan).¹¹ Despite these notable achievements, this report urges further transformational restructuring within the U.S. Navy to better deliver on autonomy and AI-based technology for UxS, maritime preeminence and superiority, consolidate and build upon these gains and prevent their erosion.



Boeing Echo Voyager UUV. Photo courtesy of Boeing and the U.S. Navy.

Specifically, this restructuring could take the form of producing only partial, limited success and technological advancement in the long term at a pace much slower than potential adversaries like the People's Republic of China, putting the Navy at increasing risk as its current competitive advantage is eroded, and perhaps reversed to the adversaries' advantage. Unlike the United States, which is generally pursuing AI to address more tactical-level issues, the Chinese are placing a big bet on AI to yield innovative strategic warfighting advantages.¹² As the National Security Commission on Artificial Intelligence (NSCAI) asserted:

The U.S. military has enjoyed military-technical superiority over all potential adversaries since the Cold War. Now, its technical prowess is being challenged, especially by China and Russia. Senior military leaders have warned that if current trend lines are not altered, the

-
- 11 Department of the Navy, *Department of the Navy Strategic Roadmap for Unmanned Systems* (Arlington, VA: Department of the Navy, 2017); and Mallory Shelbourne, "Navy Developing Unmanned Systems Network Key to New Campaign Plan 2020," *USNI News*, September 9, 2020, available at <https://news.usni.org/2020/09/09/navy-developing-unmanned-systems-network-key-to-new-campaign-plan>; and Justin Katz, "Navy's unmanned campaign plan to be 'matured' in fall," *Inside Defense*, September 11, 2020, available at <https://insidedefense.com/insider/navys-unmanned-campaign-plan-be-matured-fall>.
- 12 Elsa B. Kania, *Battlefield Singularity: Artificial Intelligence, Military Revolution, and China's Future Military Power* (Washington, DC: Center for a New American Security, 2017, pp. 13, 28).

U.S. military will lose its military-technical superiority in the coming years....Both of our great power competitors believe they will be able to offset our military advantage using AI-enabled systems and AI-enabled autonomy...DoD lags far behind the commercial sector in integrating new and disruptive technologies such as AI into its operations...Visionary technologists and warfighters largely remain stymied by antiquated technology, cumbersome processes, and incentive structures that are designed for outdated or competing aims.¹³

Until recently, the U.S. Navy has not faced a peer/near-peer threat. The long period of relative stability in the post-Cold War era reshaped R&D structures, which were optimized accordingly to that reduced major power threat. Although U.S. scientific and engineering talent remains strong, the ecosystem would benefit from increased centralized leadership oversight to synchronize and integrate disparate efforts in pursuit of clear and relevant objectives. Component ecosystem R&D organizations are engaged in activities that are sometimes disjointed or duplicative. Moreover, the collaboration that does occur is often driven by personal relationships rather than formal processes. Although improved collaboration is beneficial, this report also highlights that rapid technological advances have frequently been achieved through competitive R&D activity. Increased oversight and centralization of effort could smartly optimize competitive R&D activity within the ecosystem, providing a complementary avenue for UxS R&D advancement. Taken together, centralized leadership oversight of collaborative and competitive R&D ecosystem activity could reduce organizational friction and enhance the USN's ability to invent, innovate, and prototype in a way that meets the Chief of Naval Operations' challenge in his recently released NAVPLAN 2021 to leverage AI, Machine Learning, and autonomous system advances to enhance U.S. Navy decision superiority in combat.¹⁴

13 Eric Schmidt and Robert Work, *National Security Commission on Artificial Intelligence Final Report* (Washington, DC: National Security Commission on Artificial Intelligence), 2021, p. 61.

14 Admiral Michael M. Gilday, *CNO NAVPLAN*, January 11, 2021, available at <https://www.navy.mil/Press-Office/Press-Releases/display-pressreleases/Article/2467465/cno-releases-navigation-plan-2021/>

CHAPTER 2

Hyper-focus / Starvation among Key UxS R&D Lines of Effort

The overall diminished collaboration and unity of effort of the UxS R&D Ecosystem have resulted in a hyper-focus (defined as an overinvestment of resources)/starvation environment for research priorities. This situation occurs because LOE research within the ecosystem is not orchestrated with a centralized, top-down prioritization that could ensure a deliberate, coherent assignment of UxS LOE R&D activities. What exists is akin to a *laissez-faire* environment where the ecosystem organizations are sometimes free to self-assign and pursue their own projects (especially when guiding funds from organizations like ONR, DARPA, NISE, etc. are not in play). When such research guidance is absent, this unstructured decentralization and delegation prevents the efficient assignment of research efforts and resources (time, people, and money) on the LOEs in a prioritized manner. This has resulted in some organizations pursuing focused research in areas when one of the following conditions exists: it is more appropriate for other facilities to pursue; the return on investment is low; or the work is duplicative to a significant degree. For example, the Navigation Control LOE (the development of programs/modules that can autonomously execute precision navigation utilizing sensors and onboard inertial systems) for unmanned vehicles was researched at nearly two-thirds of the ecosystem organizations, frequently independently (i.e., non-collaboratively). While some organizations were conducting valuable and cutting-edge research in this LOE, there were a surprising number of other organizations whose LOE work in this area could be assessed to likely would yield little R&D value ROI when considering the resources expended.

Conversely, Mission Autonomy and Testing & Evaluation (T&E) are two critical LOEs where the Navy has the greatest deficit of expertise and has made the least amount of progress, yet is investing the fewest resources. This is evident in the ecosystem, as R&D efforts to

advance progress in these LOEs were assessed as sparse, especially when compared to the hyper-focused resources expended on the Navigation Control LOE. The Mission Autonomy LOE is the most important aspect of UxS that will make them transformative, revolutionary, and highly valuable in future warfare. The T&E LOE is vital to ensuring UxS attain the high level of trust by Fleet operators through critical verification, validation, and certification testing. The remaining LOEs, relative to the Mission Autonomy and T&E (for UxS) LOEs, although undoubtedly necessary to the development and prototyping of UxS, are less complex comparatively and have achieved far more technological (scientific & engineering) progress than them. In other words, the other LOEs represent the “low-to-medium hanging fruit” in which resources have been over-dedicated to pursue while the “high hanging fruit” of mission autonomy and T&E receive substantively fewer resources, yet require the most significant R&D effort to attain.

Mission Autonomy

Mission Autonomy is the development of programs or modules that can autonomously execute key mission actions/behaviors such as patrolling, anti-submarine warfare, and tracking, using all data and sensors available to the system.

Software is the ‘brain’ or key component of complex autonomous systems. Autonomous systems of any significant complexity are built of software...Autonomy in motion is closely associated with robotics, where the software is responsible for the higher-order behaviors we humans observe as a robot accomplishes its tasking. Software is, in some sense, the ‘brain’ of autonomous systems... As the demands upon autonomous operation increase, software necessarily becomes more advanced (and generally more complex and large) to meet the demands.¹⁵

Mission Autonomy is essentially the cognitive work that humans have performed in warfighting for millennia. Because of its fundamental complexity, intelligent systems are only now reaching the threshold of replicating these processes. Therefore, R&D resource starvation in this area is particularly concerning as the advancement and development of this LOE are derived directly from the exceptionally complex Artificial Intelligence discipline. Within AI, two broad categories define the ongoing research efforts: General AI and Narrow AI. “General AI (sometimes called Artificial General Intelligence, or AGI) refers to a notional future AI system that exhibits apparently intelligent behavior at least as advanced as a person across the full range of cognitive tasks.”¹⁶ General AI systems would think and reason like humans and, when provided with vague or incomplete information, deduce answers or make follow-on decisions. Examples include the HAL-9000 Supercomputer from the classic

15 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy* (Arlington, VA: Department of the Navy, 2017), p. 8.

16 National Science & Technology Council, Committee on Technology, *Preparing for the Future of Artificial Intelligence* (Washington, DC: Executive Office of the President, 2016), p. 7

book and movie *2001: A Space Odyssey* or The Terminator robot from the movie of the same name.

Conversely, the Narrow AI field seeks to develop systems that are experts or can reason within limited subject area domains where their computing power and ability to process exponentially vast quantities of data produce performances that may be superior to that of humans. Narrow AI operate systems such as strategic games and language translation programs, which require specific AI methods or toolkits to solve, thus preventing any single solution from being extrapolated or generalized to solve a broad range of problems that require intelligent behavior to resolve.”¹⁷ Prominent examples of commercial AI include the following: IBM’s Deep Blue chess-playing system, Google Deepmind’s AlphaGo and AlphaStar systems, Apple’s Siri, Amazon’s Alexa, DARPA’s AlphaDogfight, and driverless/autonomous vehicles by companies such as Uber, Google Waymo, Tesla, Nissan, and Toyota.

Within Narrow AI is the discipline of Machine Learning, which includes the specialization, Deep Learning. Future Mission Autonomy systems are likely to be developed utilizing these Narrow AI specialties as their foundation, especially given some of the success they have garnered with driverless/autonomous vehicles and other complex systems. As the 2016 White House report, “Preparing for the Future of Artificial Intelligence,” explains:

Machine learning is one of the most important technical approaches to AI and the basis of many recent commercial applications of AI. Modern machine learning is a statistical process that starts with a body of data and tries to derive a rule or procedure that explains the data or can predict the future. This approach — learning from data — contrasts with the older “expert system” approach to AI, in which programmers sit down with human domain experts to learn the rules and criteria used to make decisions, and translate those rules into software code. An expert system aims to emulate the principles used by human experts, whereas machine learning relies on statistical methods to find a decision procedure that works well in practice.¹⁸

Compared with other problem-solving software development methods that require the developers to explicitly design the software coding framework, Machine Learning has the advantage in solving problems and deducing solutions from highly complex scenarios or environments that defy the finite design limits of humans.¹⁹ Andrew Ilachinski’s CNA report further discusses the intersection of Machine Learning and autonomous systems:

Machine learning (ML) lies at the heart of most state-of-the-art research into developing intelligent, autonomous systems...Indeed, some of the best-known AI “successes” in recent

¹⁷ Ibid., p. 7.

¹⁸ Ibid., p. 8.

¹⁹ Ibid., p. 8.

years (e.g., IBM's *Watson*²⁰ and Google's *AlphaGo*²¹) were made possible by advances in ML. However, a limitation that applies to *all* extant ML methods as they apply to “narrow AI” problems is that they are effectively “black boxes” that do not easily reveal the “logic” behind the “reasoning.”²² This may be innocuous when playing an AI-system in chess, say, but assumes an entirely new (and serious) dimension if the “narrow AI” in question is embedded within a military autonomous system. For example, how does one ensure (during, say, the testing and evaluation phase of DoD's acquisition process) that whatever autonomous system being developed will not perform “surprising” (i.e., unanticipated) actions during a mission?²³

Deep Learning is a subspecialty of Machine Learning that has yielded some notable successes in the development and prototyping of autonomous systems/vehicles.

Deep learning uses structures loosely inspired by the human brain, consisting of a set of units (or “neurons”). Each unit combines a set of input values to produce an output value, which in turn is passed on to other neurons...Deep learning networks use many layers... and often use a large number of units at each layer, to enable the recognition of extremely complex, precise patterns in data.²⁴

Systems that employ Machine and Deep Learning require massive amounts of data, first to train or teach the systems the skills, expertise, or operation that they are designed to perform; and second to test and evaluate those trained systems to assess that performance. “Data is the ultimate resource for learning algorithms” that are central to those systems.²⁵

To apply machine learning, a practitioner starts with a historical data set, which the practitioner divides into a *training set* and a *test set*. The practitioner chooses a *model*, or mathematical structure that characterizes a range of possible decision-making rules with adjustable *parameters*...In practice, a model might have many millions of parameters... *Training* the model is the process of adjusting the parameters to maximize the objective function. Training is the difficult technical step in machine learning...Once a model has been trained, the practitioner can use the test set to evaluate the accuracy and effectiveness of the model. The goal of machine learning is to create a trained model that will *generalize* — it will

-
- 20 Anya Sostek, “Human Champs of ‘Jeopardy!’ vs. Watson the IBM Computer: A Close Match,” *Pittsburgh Post Gazette*, February 13, 2011, <https://www.post-gazette.com/ae/tv-radio/2011/02/13/Human-champs-of-Jeopardy-vs-Watson-the-IBM-computer-a-close-match/stories/201102130241>.
- 21 Sam Byford, “Google's AlphaGo AI beats Lee Se-Dol again to win Go series 4-1,” *The Verge*, March 15, 2016, <https://www.theverge.com/2016/3/15/11213518/alphago-deepmind-go-match-5-result>
- 22 Tao Lei, Regina Barzilay, and Tommi Jaakkola, *Rationalizing Neural Predictions* (Boston, MA: Massachusetts Institute of Technology Computer Science and Artificial Intelligence Laboratory, 2016).
- 23 Andrew Ilachinski, *AI, Robots, and Swarms: Issues, Questions, and Recommended Studies* (Arlington, VA: Center for Naval Analyses, 2017), p. 9.
- 24 National Science & Technology Council, Committee on Technology, *Preparing for the Future of Artificial Intelligence*, p. 10.
- 25 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, p. 16.

be accurate not only on examples in the training set, but also on future cases that it has never seen before.²⁶

The data needed to train and test Machine and Deep Learning-based autonomous systems must be properly labeled and annotated to be effective. Without access to a massive database of this data, which also encompasses all of the anticipated scenarios and conditions, autonomous systems cannot be effectively designed. As the NRAC report highlights, “In general, data curation is the process of turning independently created data sources (structured and semi-structured data) into unified data sets ready for analytics, using domain experts to guide the process...To train the machine learning algorithms one needs enormous annotated data sets. If one has exclusive access to the data, then one has the potential to build AI systems with capabilities that no one else has. In the future, data will win wars.”²⁷

The Mission Autonomy LOE is the most crucial aspect of UxS that will make them transformative, revolutionary, and highly valuable in future warfare. As the previous paragraphs have underscored, the cutting edge of this advanced field involves extraordinarily detailed research, most of which is being performed by academia and corporations/industry. The DoD and DoN retain very little organic expertise in this area and therefore are dependent on outsourcing to leverage the development of mission autonomy systems for UxS.

Mission Autonomy systems are the most difficult, complicated aspect to deliver, yet the Navy (and its R&D ecosystem) are dedicating only a small amount of resources to its development. Vastly greater resources are required to achieve necessary technological advancements, but this LOE is under-resourced due to insufficient prioritization. The National Artificial Intelligence R&D Strategic Plan noted the critical deficiency in AI technology at the national level needed to ensure the safety and performance of AI systems given their complexity and capability for software evolution.²⁸

To emphasize the seriousness of this underinvestment, during research study engagements within the ecosystem, an instance was encountered in which funded research produced a mission autonomy system working prototype. However, the Navy did not transition the prototype for follow-on simulation testing or incorporate it into a small-scale UxS prototype vehicle to evaluate its capability and utility for further future development. Instead, as soon as the research funding concluded, this functional prototype was abandoned without any further consideration of its potential application. Despite a paucity of similar mission

26 National Science & Technology Council, Committee on Technology, *Preparing for the Future of Artificial Intelligence*, p. 9.

27 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, pp. 16-17.

28 National Science & Technology Council, Networking & Information Technology Research & Development Subcommittee, *Artificial Intelligence Research and Development Strategy Plan* (Washington, DC: Executive Office of the President of the United States, 2016), p. 15.

autonomy prototypes available, the Navy essentially left this unevaluated system on the shelf, orphaning it.²⁹

To be tactically useful in the future, UxS mission autonomy systems must operate in a Human-supervised/Human-on-the-Loop (HOTL) mode and possess the capability to communicate to allow collaboration and swarming with other friendly UxS.³⁰ Designing and producing UxS imbued with these characteristics/capabilities will substantially amplify their operational value to the Fleet while simultaneously greatly reducing the large manpower command and control (C2) footprint that current generations of human-in-the-loop UxS require. One way to leverage AI for increased operational effectiveness is to allow single operators to control multiple air, surface, or undersea unmanned platforms. Doing that requires advances in the Mission Autonomy LOE. “AI applications in this field are similar to commercial self-driving vehicles, which use AI technologies to perceive the environment, recognize obstacles, fuse sensor data, plan navigation and even communicate with other vehicles.”³¹ However, it is also within an adversary’s capabilities as well, and they are unequivocally committed to achieving the same AI/autonomy breakthrough. The fundamental difference between the United States and its adversaries is the difference in resourcing levels and support dedicated to this goal. Of all of the LOEs that comprise the R&D and manufacturing of actual unmanned systems/vehicles (this excludes the T&E LOE), the Mission Autonomy LOE easily surpasses the other LOEs as the most important and challenging, complex aspect to this endeavor. The other LOEs all can demonstrate associated components/systems that are currently in existence and have been demonstrated in a practical, real-world environment. The same cannot be asserted for the Mission Autonomy LOE, which underscores the need to allocate increased resources towards its R&D and prototyping.

Test & Evaluation

Development of the UxS Test & Evaluation (T&E) LOE is another case where resource starvation is occurring in an area that requires substantial funding, talent, and time. Verification and Validation (V&V) of traditional software systems (to include operational

29 Interviews with Autonomy, UxS, and AI subject matter experts at the Pennsylvania State University Advanced Research Laboratory, 2017. The author conducted these SME interviews in confidentiality, and the names of interviewees are withheld by mutual agreement.

30 Human-supervised/Human-on-the Loop (HOTL) is an autonomous system that is designed to provide human operators with the ability to intervene and terminate engagements/actions. Absent human intervention, the system will performed its assigned tasks or missions under human supervision. HOTL allows for the UxS to perform a majority of their mission actions independently (simply reporting back their activity) while reserving a vital subset of significant mission actions to the purview of humans (e.g. target selection/engagement, shifting missing priorities). Ilachinski, *AI, Robots, and Swarms*, p. 146-148.

31 Bill Canis, *Issues in Autonomous Vehicle Development* (Washington, DC: Congressional Research Service), pp. 2-3.

testing & evaluation) is considered a complex endeavor in its own right.³² Those types of systems typically encompass the testing of hundreds of thousands to millions of lines of code that demand tremendous resources to comprehensively accomplish with a high degree of confidence.

T&E of critical software for conventional (i.e., non-autonomous) systems has been estimated to cost *seven times* that of software development costs.³³ The associated T&E costs for software driving *autonomous* behavior will be at least this great. Also, current practice calls for *full* system tests, which will become increasingly infeasible as autonomous systems attain ever greater levels of self-governance. In designing and certifying conventional software, one typically needs only to address the issues involved in answering *what* (the software does) and *when* (the software does it). Since autonomous systems must make their own decisions, autonomy inevitably introduces the more complex *why* (the software chooses to do something).³⁴

The herculean effort involving the ongoing T&E of driverless vehicles by corporations and academic institutions is instructive. For at least a decade, these organizations have been devoting massive resources (billions of dollars, millions of man-hours, thousands of experts) to T&E for driverless vehicles. After such a substantial outlay, they have only recently attained a confidence level that supports deployment of these vehicles in limited settings.³⁵

To be operationally useful, UxS will need to perform an array of complex behaviors (missions) in a dynamic environment in which these vehicles will often need to make decisions with ambiguous, incomplete, erroneous, or unknown information. C2 opportunities and direction may be unavailable in a benign environment (e.g. undersea, far over the horizon, etc.) or denied in a challenging environment. Within this situation, UxS T&E will involve the highly complex integration of communications, sensors, navigation, payloads, energy systems, and propulsion into vehicles that are controlled by a supervising mission autonomy system required to operate unsupervised anywhere from days to months within a dynamic environment. As noted above, the Mission Autonomy systems of the near future for UxS will almost certainly rely on the complex AI discipline of Machine/Deep Learning as a sizable component, which brings all of the associated technical challenges for T&E.

32 *Verification* is the process of determining that a model or simulation implementation and its associated data accurately represent the developer's conceptual description and specifications; i.e., verification effectively answers the question, "*Did we build the system correctly?*" *Validation* is the process of determining the degree to which a model or simulation and its associated data are an accurate representation of the real world from the perspective of the intended uses of the model; i.e., validation answers the question, "*Is the system the right solution to the problem?*" Department of Defense, *DoD Instruction Number 5000.61: DoD Modeling and Simulation (M&S) Verification, Validation, and Accreditation (VV&A)* (Arlington, VA: Department of Defense, December 9, 2009).

33 Christine Hang, Panagiotis Manolios, and Vasilis Papavasileiou, "Synthesizing Cyber-Physical Architectural Models with Real-Time Constraints," *Computer Aided Verification* (Springer Berlin Heidelberg, 2011).

34 Ilachinski, *AI, Robots, and Swarms*, p. 200.

35 Global automotive research and development expenditures are estimated at \$94.2 billion for 2016 — more than three times the 2016 global expenditures on aerospace and defense. Missy Cummings, "The Brave New World of Driverless Cars: The Need for Interdisciplinary Research and Workforce Development," *TR News*, March-April 2017, pp. 34-37.

With these variables in play, conducting the T&E for UxS run by mission autonomy systems will demand an order of magnitude more resources than driverless vehicle T&E. The NRAC report noted that historically,

evolving Verification, Validation, & Accreditation (VV&A) methods and tools lagged the developments of the systems to be evaluated, but through deliberate investment in research, they were developed and eventually employed. The VV&A methods and tools for learning systems need similar deliberate investment. These methods and tools do not exist today. The DoD Autonomy Community of Interest (COI) has taken a significant step focused on Verification and Validation (V&V) Strategy, including the testing and evaluation (T&E) of autonomous learning systems.³⁶

Despite this exponentially more difficult requirement, the U.S. Navy is not dedicating sufficient resources toward making progress in designing, establishing, and implementing a comprehensive T&E Trust-Building Organizational Regime. This regime includes not only modeling and simulation, but also extensive component and full-system vehicle integration tests on operational ranges in dynamic environments/scenarios. To accomplish its implementation, the regime must include not only scientists, engineers, and operational and tactical experts from the Fleet at every step of the process. Incorporating this essential aspect of the process is part of the overarching operationalization of UxS in the Fleet.

Current engineering methods for T&E/VV&A of hybrid hardware-software systems lack sufficient fidelity and robustness to deal with highly complex, software-intensive systems. The most difficult and challenging component of autonomous weapon systems to certify is the AI/machine-learning/adaptive *software* embedded within them.³⁷

Consequently, this situation is solidifying the likelihood of a scenario in which the Navy receives prototype vehicles but will not yet have developed the robust T&E process to evaluate them to a high degree of confidence, thus setting the stage for additional delays measured in years while it pursues the critical work needed. During research visit engagements with the UxS R&D ecosystem, when interviewees were queried about their thoughts regarding T&E, they would often report that the T&E Trust Building Operationalization Regime would be developed once the Navy received autonomous UxS prototype vehicles. Such statements indicate a “working in series” mindset in which critical R&D for one phase of work does not begin until a previous phase is completed — an approach that can add several years to the fielding of unmanned autonomous vehicles relative to a “working in parallel” strategy.

To fully appreciate the impact of the lack of resources assigned to T&E, examining the driverless vehicle industry is instructive as its T&E outlay serves as a baseline and reasonable

36 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, p. 22.

37 U.S. Department of Defense, Defense Science Board, *Summer Study on Autonomy* (Arlington, VA: Department of Defense Office of the Under Secretary of Defense of Acquisition, Technology, and Logistics), June 2016; Ilachinski, *AI, Robots, and Swarms*, p. 199.

proxy for the scale of effort/resources the Navy will need to commit. Since the original DARPA driverless vehicle Grand Challenge over a decade ago, the sheer magnitude of resources committed by companies to the T&E process has been enormous in cost, time and manpower. Different companies have approached T&E using a variety of methodologies that range from majority computer modeling & simulation up front (with on the road live tests occurring later in the process) to the reverse in which on the road testing is the initial evaluation means.³⁸ Most companies fall somewhere in between. No matter which path companies choose, all of them have an urgent need to acquire massive amounts of data related to vehicle performance, whether it be simulated or live; driverless mode, semi-autonomous mode, or manual mode with the human fully in control.

To begin evaluating what measure of performance is needed to quantify the T&E effort and characterize the safety of driverless vehicles as compared to humans, industry settled on two statistics, “miles driven” (the total miles a vehicle model has driven — either simulated or actual — and a comparison, “miles driven without fatalities”).³⁹ The latter statistic was used to compare how many miles a driverless vehicle operated before a crash fatality occurred versus the same metric for human-driven vehicles. These metrics, defined in a RAND study, sought to bound the T&E problem by quantifying how much testing would be required to demonstrate a theoretical degree of safety that was at least on par with human driving performance. That report determined autonomous vehicles would require between hundreds of millions to possibly billions of millions live, on-road testing to properly demonstrate their safety reliabilities, an impossible prospect if the assurance was required to precede actual deployment.⁴⁰ In other words, autonomous vehicle developers will need to rely on other methods to complement on-road testing.

Besides calculating a theoretical miles-driven value, the RAND study asserted another key finding — it was impossible to realistically conduct sufficient live, on-the-road miles-driven testing to achieve the theoretical value. To meet the RAND proposed standard,

38 “Across the range of different types of stakeholders we interviewed, stakeholders consistently indicated that the complexity of automated technologies and the difficulties in demonstrating their safety are a challenge for existing public-sector safety programs...Vehicle developers are likely to need a combination of on-road testing, simulations, and other techniques to validate their technology functions as intended. For example, one study estimates that millions of on-road testing miles could be needed to prove that automated vehicles perform safely; so many, in fact, that the study concludes on-road testing alone may not be practical. Of particular concern might be situations that are atypical and may not be routinely experienced in on-road testing, but may actually be sometimes or commonly experienced by drivers (e.g., a child running into the street or a work zone).” Government Accountability Office, *Comprehensive Plan Could Help DOT Address Challenges* (Washington, DC: Government Accountability Office), 2017, p. 11, Nidhi Kalra and Susan M. Paddock, *Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability?* (Santa Monica, CA: RAND Corporation), 2016, Philip Koopman and Michael Wagner, *SAE World Congress and Exhibition*, April 14, 2016.

39 Driverless car companies have generated a “miles driven” metric to assure safety. RAND Corporation has stated that driverless cars must drive 275 million miles without a fatality to prove that these cars are as safe as human drivers, at a 95 percent confidence level. (Kalra and Paddock 2016). Tesla logged 130 million miles before the U.S. fatality — the most of any company — falling significantly short of the RAND Corporation metric. Cummings, “The Brave New World of Driverless Cars”, 2017, p. 36.

40 Kalra and Paddock. *Driving to Safety*, p. 10.

companies would need to pursue their driverless vehicle T&E through a combination of simulation (which logically would constitute the majority of those tests) and actual operations in the field. For those simulated miles driven to be effective, the driverless vehicle simulators required a very high degree of fidelity to accurately represent a real-world operating environment for vehicles. Furthermore, that environment must be populated with realistically-behaving vehicles and precisely-replicated phenomena such as weather, road conditions, and vehicle performance. This requires the collection of tremendous amounts of data by crewed vehicles followed by the curation of that product by data scientists before simulation engineers can develop simulations attaining the appropriate level of fidelity. One article discusses how Waymo, one of the leaders in driverless simulation, integrates this method into its T&E program:

Waymo, comparatively, sounds more confident about its simulations. The company re-creates full computer models of the cities it's testing in, and sends 25,000 'virtual self-driving cars' through them each day [and] created a tight feedback loop by recreating real-world driving data on the computer, where 'thousands of variations' of a scenario can be run. The data is then downloaded back into Waymo's test cars. Waymo has also built a dedicated test facility in California, where it can build out particular street features or stage scenarios that seem to give its vehicles the most trouble.⁴¹

The importance of simulation has been highlighted in the aftermath of the Uber collision in March 2018 that resulted in the death of a pedestrian.

When a self-driving car prototype operated by Uber fatally struck a pedestrian in Tempe, Ariz., in March, Uber quickly identified the likely cause in software that caused the vehicle to ignore certain objects that its sensors detected. But a further realization dawned on some executives and members of the team: The rush to develop a commercial self-driving vehicle had led Uber to de-emphasize computer simulation tests that attempt to anticipate how autonomous vehicles would react in millions of driving scenarios. That stood in contrast to the process at Alphabet's Waymo and some other major companies developing self-driving cars, where simulation testing was a top priority.⁴²

The driverless vehicle's demanding, challenging journey towards researching T&E principles, determining the scale and scope of the T&E process, and then implementing it through a comprehensive program provides an invaluable lesson and point of departure for how the Navy will need to conduct its autonomous UxS T&E program. The Navy's program will likely

41 Sean O'Kane, "How Tesla and Waymo are Tackling a Major Problem for Self-Driving Cars: Data," *The Verge*, April 19, 2018, <https://www.theverge.com/transportation/2018/4/19/17204044/tesla-waymo-self-driving-car-data-simulation>

42 The failure to prioritize simulation testing stemmed from earlier decisions by the business unit and made worse by the growing pressure within Uber to move quickly to launch robo-taxis into the ride-hailing market by this year. Simulation testing is time-consuming, as engineers need to figure out how to create a virtual world that mimics the real one. There aren't many engineers with expertise in the field, and constructing a useful simulation system is considered an art form. Amir Efrati, "Uber Neglected Simulation on Self-Driving Cars, Insiders Say," *The Information*, June 19, 2018, <https://www.theinformation.com/articles/uber-neglected-simulation-testing-on-self-driving-cars-insiders-say>

be more complex in its scope and scale than the driverless vehicle program for all of the reasons discussed in the previous mission autonomy section. Although progress has been marginal for Navy UxS T&E efforts, valuable activity has occurred in various areas, which this report will discuss in the following section.

During site visits, CSBA discovered T&E centers of expertise focused on R&D to advance autonomous UxS T&E. In many cases, the autonomous UxS R&D center was either resident in the organization based on historical expertise (e.g. the Jet Propulsion Laboratory, NAVAIRSYSCOM, NAVSEASYCOM, FFRDCs, and UARCs) or had been developed in parallel with ongoing UxS testing at the organizations. For example, the U.S. Navy commissioned the Unmanned Undersea Squadron (UUVRON) ONE to serve as the Navy's focal point in T&E for UUVs.⁴³ As a consequence, their host facility and partner, the Naval Undersea Warfare Center Keyport has lent its submarine testing expertise to the UUVRON while simultaneously developing its UxS experience and knowledge through various means. In the UUVRON's case, they have embarked on a plan to rapidly and aggressively ramp up their own experience and expertise by both inviting experts from industry, academia, Fleet, and FFRDCs/UARCs to visit the UUVRON to train and work alongside their personnel and also send many of those personnel to back these organizations for extended immersion training.⁴⁴

The UUVRON has demonstrated the worthwhile progress and value produced by a modest, limited investment in establishing this T&E organization. Its forerunner, the Submarine Development Squadron (SUBDEVRON) performed aspects of this UxS T&E previously, but the creation of the UUVRON has allowed for the development of organic T&E UxS expertise and experience within the Navy.⁴⁵ With its Air Test and Evaluation Squadrons (AIRTEVRON), the Navy has long recognized the value of having organic aviation T&E experts, and the AIRTEVRONs have incorporated unmanned air systems/vehicles (UAV) into the responsibilities for the squadrons.⁴⁶ To build similar critical expertise and experience for unmanned surface vehicles (USV), the Navy should continue investing the proper resources in the Surface Development Squadron (SURFDEVRON), which has assumed the primary role in the T&E of MDUSVs like Sea Hunter.⁴⁷

43 Joseph Trevithick, "The US Navy Has Created Its First Ever Underwater Drone Squadron," *The Drive*, September 28, 2017, available at <https://www.thedrive.com/the-war-zone/14733/the-us-navy-has-created-its-first-ever-underwater-drone-squadron>

44 Organizations that provided this type of training support included Woods Hole Oceanographic Institute (WHOI), Naval Oceanographic Office, Penn State University ARL, Scripps Oceanographic Institute, Hydroid, Boeing and other industries.

45 Submarine Development Squadron FIVE (COMSUBDEVRON FIVE), available at <https://www.csp.navy.mil/csds5/>

46 Global Security, "VX-1 Air Test and Evaluation Squadron ONE [AIRTEVRON ONE] Pioneers," available at <https://www.globalsecurity.org/military/agency/navy/vx-1.htm>

47 Similar to the SUBDEVRON and AIRTEVRON, a SURFDEVRON would focus on the exploration, development and evaluation of advanced tactics and operational capabilities for the Surface Navy.

The UUVRON, AIRTEVRONS, and SURFDEVIRON have the potential to develop significant UxS expertise and experience. However, they will be limited in achieving this knowledge without substantial partnership and collaboration with the UxS R&D ecosystem organizations that possess the highly technical T&E knowledge, especially related to mission autonomy systems. As discussed previously, the ecosystem organizations along with the Navy still require significant resource investment in T&E research and development to gain sufficient expertise on how to comprehensively conduct those evaluations on future UxS prototypes. Some of the ecosystem organizations have collaborated on individual autonomous UxS projects, but when assessed collectively, the pursuit of UxS T&E and Trust Building Regime expertise is disjointed, fragmented, and under-resourced. To deliver progress and capability, the U.S. Navy must provide substantively greater resources (people, funding, time) and increase its priority (along with Mission Autonomy) among the various UxS R&D LOEs. The NRAC report highlighted this shortcoming as well and provided corrective recommendations:

We recommend that DoN create a world-class VV&A research program [in 6.1 to 6.3] for autonomous systems by dramatically expanding the work being done by the DoD Autonomy Community of Interest [IV]... In the future, where autonomous learning systems supplant human operations, VV&A becomes extremely critical and more difficult to achieve.” [Pg. 21]...Autonomous systems are complex, and some of the most advanced elements, such as Deep Learning elements, are effectively “black boxes”. Consequently, Validation, Verification, and Accreditation (VV&A) of these systems is particularly challenging. Many of the mature VV&A methodologies for complex systems (e.g., those followed for deploying aircraft software), do not apply to autonomous systems. Furthermore, we believe that the ability to upgrade autonomous systems quickly may be essential for battlefield success, which presents new VV&A challenges.⁴⁸

DoD Directive (DoDD) 3000.09H (“Autonomy in Weapon Systems”) should be applied as a framework to understand how the previously discussed T&E issues must fit together to support the Navy’s development of a T&E Trust Building Operationalization Regime. This document serves as the top-level policy guidance on how autonomous systems must be tested. DoDD 3000.09 requires that weapons systems:⁴⁹

- Go through “rigorous hardware and software verification and validation (V&V) and realistic system developmental and operational test and evaluation (T&E), including analysis

48 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, pp. iv-v.

49 Enclosures 2 and 3 of DoD Directive 3000.09 (*Autonomy in Weapon Systems*, Nov 2012) address T&E and V&V issues, and general review guidelines, respectively. Department of Defense, *Department of Defense Directive Number 3000.09: Autonomy in Weapons Systems*, (Arlington, VA: Department of Defense), November 21, 2012.

of *unanticipated emergent behavior* (emphasis added) resulting from the effects of complex operational environments on autonomous or semiautonomous systems.”⁵⁰

- “Function as anticipated in realistic operational environments against *adaptive adversaries* (emphasis added).”
- “Are sufficiently robust to minimize failures that could lead to *unintended engagements* (emphasis added).”

Of note, UxS T&E will be distinctly different from the process employed for conventional military platforms and systems. The difference is driven by the Mission Autonomy systems that will be embedded within the UxS, which are derived from artificial intelligence-infused algorithms. Conventional software engineering techniques follow a cycle of development in which the software realism/accuracy matures with each iteration, but as highlighted in the following, such a process is problematic for autonomous systems. However, artificial intelligence and machine learning-based Mission Autonomy systems will be significantly more challenging to evaluate given they will possess the ability to learn (without explicitly being programmed), which will allow them to evolve and modify their own behavior over time.⁵¹

The acquisition of *autonomous* weapons entails testing and evaluation procedures distinctly different from conventional systems, even those with an intensive software focus (i.e., a major distinctive element may be appreciated, intuitively, by reflecting on the difference between “conventional” operating instructions for automated hardware systems and AI-derived behavioral-logic necessary to govern autonomous vehicles).⁵²

Determining whether human operators will trust an UxS goes to the heart of the need for the T&E Trust Building Operationalization Regime. Before proceeding further, it is important to explore the definition and nature of trust as it relates to UxS.

Trust is not an innate trait of the system; rather, it is a relative measure of how a human operator (or operators — whose own performance depends, in part, on collaborating in some way with the system — experiences and perceives the behavior of the system; or, better, *how a human operator perceives the behavioral pattern of a system*. Trust is “an attitude that an agent will help achieve an individual’s goals in a situation characterized by uncertainty

50 Expanding on V&V and T&E, DoDD 3000.09 requires that they must “assess system performance, capability, reliability, effectiveness, and suitability under realistic conditions, including *possible adversary actions*, consistent with the *potential consequences of an unintended engagement* or loss of control of the system.” Furthermore, in regard to the human-machine interface (for autonomous and semiautonomous weapon systems), it must be: (1) “Readily *understandable* to trained operators;” (2) “Provide *traceable feedback* on system status;” and (3) “Provide clear procedures for trained operators to *activate and deactivate system functions*.” Department of Defense, *DoD Instruction Number 5000.61: DoD Modeling and Simulation (M&S) Verification, Validation, and Accreditation (VV&A)* (Arlington, VA: Department of Defense), December 9, 2009.

51 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, p. 22.

52 Ilachinski, *AI, Robots, and Swarms*, p. 197.

and vulnerability.”⁵³ Though there are many alternative formulations of this basic description, attempts at a more formal definition, and lists of basic attributes needed, to varying degrees, for the creation of trust in a system, the underlying truth — and, as will be argued, a key reason why autonomy is so hard to certify — is that *trust is an abstraction that cannot be easily mathematized*. No absolute measure of trust exists; rather, it is a relative measure that can be described only in terms of how an a priori level of trust has changed.⁵⁴

Trust is complex and multidimensional. 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.⁵⁵

Operationalization is the process by which Navy must conduct a full spectrum T&E on an UxS spanning algorithmic V&V; robust modeling & simulation assessment to thoroughly discover and explore failure “corner” & “edge” scenarios; human-machine teaming/interoperability assessment; and live testing of representational scenarios in a live environment on a range. Here it is worth considering the challenging complexity of the endeavor.

Before an AI system is put into widespread use, assurance is needed that the system will operate safely and securely, in a controlled manner. Research is needed to address this challenge of creating AI systems that are reliable, dependable, and trustworthy. As with other complex systems, AI systems face important safety and security challenges due to:⁵⁶

- *Complex and uncertain environments*: In many cases, AI systems are designed to operate in complex environments, with a large number of potential states that cannot

53 John D. Lee and Katrina A. See, “Trust in Automation, Designing for Appropriate Reliance,” *Human Factors*, 2004, p. 51.

54 Ilachinski, *AI, Robots, and Swarms*, p. 183.

55 There is no objective absolute measure of “trust.” It is, at best, a relative concept, such that terms such as those that appear in the above list can be used to measure the *relative differences* between (e.g., “Is system A more or less ‘reliable’ than system B for a given set of tasks?”). If, and when, trust is established (by human operator, P, in a given system, S), it will *not* be the result of some “trust threshold” being exceeded (and measured in a mathematically precise manner). Rather, the trust will emerge over time, as P trains and works with S, and eventually determines — partly as a result of objective measures, and partly as a result of a subjective assessment of S’s patterns of behavior — that S can adequately perform the set of tasks assigned to it. Even so, trust is more of a dynamic attribute of an ongoing series of human-machine collaborations than it is a static measure (that, once achieved, remains fixed). Relationships between human operators and robots can evolve and mature just as they do between humans. They can strengthen, weaken, and transform (in unanticipated ways) over periods of time, and the degree of trust that an operator bestows on a machine, in general, depends on specific contexts and mission goals. U.S. Department of Defense, Defense Science Board, *Summer Study on Autonomy* (Arlington, VA: Department of Defense Office of the Under Secretary of Defense of Acquisition, Technology, and Logistics), June 2016, Section 2; and Ilachinski, *AI, Robots, and Swarms*, p. 185-186.

56 Jon Bornstein, “DoD Autonomy Roadmap: Autonomy Community of Interest,” *NDIA 16th Annual Science & Engineering Technology Conference/ Defense Tech Exposition*, March 24-26, 2015.

be exhaustively examined or tested. A system may confront conditions that were never considered during its design.

- *Emergent behavior*: For AI systems that learn after deployment, a system's behavior may be determined largely by periods of learning under unsupervised conditions. Under such conditions, it may be difficult to predict a system's behavior.
- *Goal misspecification*: Due to the difficulty of translating human goals into computer instructions, the goals that are programmed for an AI system may not match the goals that were intended by the programmer.
- *Human-machine interactions*: In many cases, the performance of an AI system is substantially affected by human interactions. In these cases, variation in human responses may affect the safety of the system.⁵⁷

At the end of this comprehensive regime, scientists, engineers, and Fleet Sailors/operators will collectively make a recommendation on whether they have achieved a sufficient level of confidence to trust the UxS for fleet introduction. "Achieving effective interactions between humans and AI systems requires additional R&D to ensure that the system design does not lead to excessive complexity, undertrust, or overtrust. The familiarity of humans with the AI systems can be increased through training and experience, to ensure that the human has a good understanding of the AI system's capabilities and what the AI system can and cannot do."⁵⁸

As part of this confidence assessment, the APO in coordination with Commander, Test & Evaluation Force (COMOPTEVFOR) will issue some type of certification or acceptance recommendation, which includes the degree/scope of operations authorized (e.g. continued testing, exercise operations, deployed operations with restrictions, unrestricted deployed operations, etc.). From this point, the operationalization moves into a final phase in which the Fleet collaborates with the scientists and engineers to adjust/modify the planned UxS concept of operations (CONOPS) for its employment to account for the prototypes actual operating parameters and profile. Along with this step, personnel will also need to develop the training doctrine to support introduction and integration into the Fleet.

57 Jeffrey Mark Bradshaw, Robert Hoffman, Matthew Johnson, and David D. Woods, "The Seven Deadly Myths of 'Autonomous Systems'," *IEEE Intelligent Systems*, no. 3, May 2013, pp. 54-61.

58 National Science & Technology Council, Committee on Networking and Information Technology, *Artificial Intelligence Research and Development Strategic Plan* (Washington, DC: Executive Office of the President), 2016.

CHAPTER 3

Improving Organizational Strategy

The Department of the Navy has articulated the latest version of a signed Strategic Roadmap for Unmanned Systems and an accompanying Unmanned Systems Goals memorandum that continue to develop the strategy and vision for research, development, and procurement of unmanned vehicles.⁵⁹ However, these high-level documents could go further in laying out a practical organizational framework to execute this strategy, achieve these goals, and transform the vision into concrete actions, which will deliver technological advancements. In other words, absent from these official pronouncements are the comprehensive organizational restructurings that are needed to swiftly transition the Navy onto an optimal track. Without this restructuring, the existing bureaucracy may be insufficient to the task, and the current organizational weaknesses may continue to hamper UxS technological progress/ advancement. The U.S. Navy is progressing toward the completion of the Navy Unmanned Campaign Plan, which may address the above items, but as of this writing, this is unknown.

In an ideal situation, it is within this middle organizational tier that the senior naval leadership's strategy roadmap and vision get translated into practical actions for the scientists, engineers, managers, logisticians, and other key supporting personal to pursue. Conversely, this tier articulates back to that senior leadership the barriers, requirements, and other important information they encounter that hinder progress, in addition to highlighting where and how the leadership can assist in overcoming those challenges.

However, the Navy's current UxS R&D ecosystem is not structured along these lines. This middle tier, as a coherent and unifying organizational force, is missing within the UxS R&D

59 The Navy intends to mature this guidance into the forthcoming Navy Unmanned Campaign Plan. Department of the Navy, Department of the Navy Strategic Roadmap for Unmanned Systems (Arlington, VA: Department of the Navy), 2017; and Mandy Mayfield, "SNA NEWS: New Navy Unmanned Systems Plan Forthcoming," National Defense Magazine, January 11, 2021, available at <https://www.nationaldefensemagazine.org/articles/2021/1/11/navy-unmanned-campaign-plan-to-be-released-soon>

ecosystem. Put succinctly, the Navy lacks an operational middle tier that can overcome the inefficiencies described below and more effectively and centrally coordinate the UxS R&D ecosystem's collective efforts.

Lack of awareness of key organizational centers of expertise within the R&D ecosystem

During CSBA's engagements with the ecosystem, there was frequently a lack of awareness among these organizations regarding the core competencies or expertise resident within their peer UxS ecosystem organizations. This gap in organizational competency awareness directly manifests itself in an inability to seek expert assistance from other R&D organizations when technical challenges stymie efforts to advance research or continue to the next logical phase in a project. Furthermore, it engendered a "series" vs. "parallel" mentality regarding R&D advancement where some individual ecosystems ceased work on projects once they hit a technical challenge in which their organization lacked the expertise to continue making progress. In essence, they self-limited their progress and in some instances, indicated that they were waiting for other UxS components to be developed before they could continue their work. For example, very few offices within the Navy and organizations with the ecosystem were aware that the NASA Jet Propulsion Laboratory (JPL) possessed exceptional mission autonomy and T&E expertise, which they had developed over five decades of manned and unmanned, orbital and interplanetary space vehicle design, development, launching, and operations. For the R&D ecosystem to achieve its maximum potential, these organizations must be fully aware of their peers' capabilities and core competencies.

Lack of awareness of complementary R&D efforts and resident expertise at other organizations within the R&D ecosystem

In many cases, very similar R&D efforts were occurring at various ecosystem organizations, yet those centers were unaware of that parallel work at their peer organizations. One example was the UxS Heterogeneous Mesh Network R&D effort, in which an architecture allows for the integrated communication and coordination by unmanned autonomous air, surface, and undersea vehicles collaboratively to accomplish a primary mission, consisting of multiple sub-missions. During engagements, four major research constellations were conducting advanced R&D on this effort, which included significant live demonstrations on exercise ranges, yet their awareness of each other's work ranged from minimal to none. Consequently, there was little to any formal collaboration between these R&D efforts, and they shared few lessons learned and other successes (or failures) to advance this important research effort.

In other examples, when collaboration between and among organizations did occur, it typically originated through personal contacts. These collaborations are mostly informal in nature and are not widely known at the organizational leadership level nor among Navy UxS

leadership, and often cease after key individuals move on. Uncoordinated but overlapping R&D efforts create a competitive environment, albeit one that lacks the benefits of formally defined competition. This results in potentially wasted resources and missed opportunities for the synergistic efforts that collaboration can produce.

Few acquisition programs have onramps designed to support the transition of promising UxS demonstrators and prototypes into applicable programs. This outcome further constrains the opportunity to reevaluate previously shelved prototypes when the surrounding technologies or conditions advance to a point that makes them viable. During research engagements with the UxS R&D ecosystem, scientists expressed concerns about S&T prototypes developed at their organizations to demonstrate future capabilities but were orphaned without ever even being evaluated for utilization. To be clear, these weren't occurrences in which a Navy entity assessed the product or capability and decided although functional, the demonstrated item did not hold value for further exploration or transition into a future system. Instead, after completion of the system development and expiration of the associated R&D funding, these systems simply were shelved, relegated to being forgotten in the archives. The Science & Technology (S&T) community frequently refers to this issue as failing to cross the Technology Readiness Level (TRL) "Valley of Death." In 2015, a Government Accountability Office (GAO) report explicitly called attention to this issue:

DoD has long noted the existence of a chasm between its science and technology community and its acquisition community that impedes technology transition from consistently occurring. This chasm, often referred to by department insiders as 'the valley of death,' exists because the acquisition community often requires a higher level of technology maturity than the science and technology community is willing to fund and develop.⁶⁰

In particular, the Valley of Death issue revolves around the ability to transition technology demonstrations from lower TRLs (e.g. three or four) to high TRLs (e.g. five, six, or seven). The GAO report on DARPA explains that the agency's primary mission of discovering and nurturing innovative new technologies through investment to demonstrate their feasibility and potential application often conflicts with a secondary priority to successfully transition them to other DoD organizations, such as military service research agencies, laboratories, and acquisition programs of record, for further adaptation and refinement for warfighter operations.⁶¹ Given the immature technology readiness level of some of these DARPA projects at the conclusion of their demonstrations, potential military service partners are not willing to adopt the new technology without further maturation and that needed R&D frequently goes unfunded due to competing budgeting priorities, creating a major impediment to transitioning new technology.⁶²

60 Government Accountability Office, *Defense Advanced Research Projects Agency*, 2015, p. 7.

61 Government Accountability Office, *Defense Advanced Research Projects Agency*, 2015, pp. 7,11.

62 Ibid.

Although the Navy has sought to address the technology transition problem through the Office of Navy Research (ONR) Future Naval Capabilities (FNC) Program, establishing a cross-organization Autonomy Portfolio Manager, and implementing rapid acquisition/prototyping programs, the Valley of Death issue continues to hamper funded projects also.⁶³ The transition problem becomes an acute issue for UxS R&D because of the lack of integration in the overall development of those systems and vehicles. Since DARPA and ONR technology development funding is often assigned to organizations within the R&D ecosystem in discrete, isolated streams, the prototype components that are ultimately produced are frequently not tethered to existing UxS main projects or long-term capability pursuits. Consequently, they cannot be evaluated against those overarching projects (if evaluated at all) to assess whether the demonstrated capability is worthy of continued development/refinement for spiraling into future UxS prototypes. Put another way, technology R&D seed funding is allocated to ecosystem organizations along discrete LOEs and is subsequently judged for future worthiness within that myopic category as opposed to a more holistic, high-level, and broad overview that a UxS prototype vehicle perspective could better provide. As discussed in the Mission Autonomy section, this very phenomenon was encountered where a mission autonomy technology demonstration system was neither transitioned for further prototyping nor had funding continued to improve on its capabilities, despite there being no other known mission autonomy system that was its equal or better.

On some occasions, the failure of technology demonstrations to transition is the result of cultural resistance from the service or a community within the service, which doesn't advocate for that particular technology for a variety of reasons. As GAO observed with DARPA innovations, "The introduction of DARPA's radically innovative technologies can disrupt the status quo for military programs, budgets, and warfighting doctrine, which can drive cultural opposition within the military services. DARPA officials stated that the agency's research sometimes leads to the identification of technologies and capabilities that military

63 The Future Naval Capabilities (FNC) program, initiated by the Department of the Navy in 2002, is a science and technology (S&T) process designed to develop and transition cutting-edge technologies to acquisition programs of record within a three-year timeframe. The program delivers these technologies as FNCs for integration into platforms, weapons, sensors or specifications to improve Navy and Marine Corps warfighting and support capabilities. The program for FY18 and out was restructured to accelerate both the selection to commencement and the S&T development timelines. FNCs typically begin at a point at which component validation in a laboratory/relevant environment has occurred (Technology Readiness Level, or TRL 4/5). FNCs are subsequently matured to the point of a demonstrated model or prototype in a relevant/operational environment (TRL 6/7). Once the FNC is demonstrated, the acquisition sponsor takes responsibility for conducting any additional research, development, test and evaluation (RDT&E) necessary to engineer and integrate the technology into an acquisition program of record, or other program, ultimately leading to the deployment of the new capability into the fleet or force. Each Future Naval Capability (FNC) Enabling Capability (EC) is an aggregate of science and technology that is aligned to an identified warfighting gap or warfighting capability, and it can deliver a distinct, measurable improvement that contributes to closing the corresponding warfighting gap. FNCs manage the development of 6.2 (applied research) and 6.3 (advanced technology development) efforts to mature technology for transition to naval acquisition programs. Department of the Navy, "Future Naval Capabilities (FNC)" (Arlington, VA: Office of Naval Research), available at <https://www.onr.navy.mil/en/Science-Technology/Departments/Code-35/All-Programs/air-warfare-and-naval-applications-352/future-naval-capabilities>

service officials do not initially want or think their services will need, although these technologies can eventually provide important military capabilities.”⁶⁴

There is another consequence beyond the immediate impact of missed opportunities to harness the fruits of R&D efforts of orphaned technology demonstrations. As currently configured and organized, the Navy UxS R&D ecosystem is not a learning organization that has awareness and can leverage all of the knowledge gained across UxS R&D activities. As a result, orphaned projects are sometimes forgotten only to be resurrected and reduplicated down the line when other organizations are tasked and funded to pursue the development of similar technology. Several ecosystem organizations recounted stories of this type of occurrence. Unfortunately, the outcome of this behavior is the waste of precious resources (people, time, money), which are not available in abundance.

In summary, the lack of awareness of ongoing R&D efforts throughout the ecosystem (especially similar endeavors), is indicative of multiple, decentralized R&D authorities, which do not synchronize their efforts and research mandates due to the lack of formal authorization and accounting structure.

Bureaucratic, Administrative, and Risk-Aversion Impediments to R&D

UxS R&D efforts are hampered by a significant level of bureaucratic and administrative barriers and impediments that often take the form of risk management measures. There are of course some measures, which make sense for safety, financial, and other reasonable purposes. However, there exists a notable number of policies, regulations, and other measures that have created levels of risk aversion, which stifle the research and innovation efforts of scientists, engineers, and other support personnel involved in UxS R&D within the ecosystem. The NRAC report identified this as a significant issue, which it addressed throughout the document, starting with the following comment:

The current bureaucratic approach, centered on exacting thoroughness and minimizing risk above all else, is proving to be increasingly unresponsive. We must transition to a culture of performance where results and accountability matter...The aversion to taking programmatic risk at the expense of strategic capability is something that is reinforced by incentives and disincentives at many levels. We have layers of accountability at the program manager level when they ‘overreach’ and ‘fail,’ but there is no accountability at the senior level when they fail to advance our capabilities fast enough to meet the strategic threat.⁶⁵

During research engagements with the ecosystem, high levels of frustration and complaints were expressed by personnel when dealing with measures involving battery operations (Lithium receiving the greatest complaints), UAS testing, interpretations of DoDD

64 Government Accountability Office, *Defense Advanced Research Projects Agency* 2015, p.19.

65 Department of Defense, *Summary of the 2018 National Defense Strategy of The United States of America* (Arlington, VA: Department of Defense), 2018, pp. 10, 28.

3000.09H as it relates to types of UxS R&D permitted, and Interim Authority To Tests (IATTs).⁶⁶ Researchers across the ecosystem possessed a deep desire to deliver advanced cutting edge R&D technological advancements for the military, but expressed exceptional levels of dissatisfaction with policies, regulations, and measures that hobbled their efforts without seemingly adhering to logic or common sense in their drafting and/or application. Ecosystem researchers often decried that some measures were crafted without direct input or feedback from the subject matter experts who had the requisite knowledge in the associated field. The NRAC report further discusses:

Virtually everywhere we did our fact-finding, people working on Navy unmanned/autonomous systems pointed to another factor they deemed even more consequential in slowing the movement of this new capability to the fleet — an unwillingness to take what might be called “bureaucratic risk” or to project a realistic cost/benefit analysis for the testing that will be required...Our research strongly suggests that the Department of the Navy and those private entities that regularly do business with it suffer from a pervasive culture of minimizing risk at the program management level at the expense of risk at the strategic level. Others like Amazon and Space-X are willing to accept reasonable risk where warranted...It is not an exaggeration to say that people we spoke with whose job encompassed developing and pushing new technology out to the warfighter felt their attempt to do so encountered debilitating barriers that strangled initiative and demoralized them. We recognize these are very strong words, but, in truth, we would use stronger words if we could find them.⁶⁷

Others complained about tedious or overly bureaucratic review/approval processes and regimes, which added substantial administrative burden to successfully traverse. In the interim, critical and innovative research would often come to a halt while negotiating the approval processes, which some criticized seemed to serve themselves more than the R&D customers pursuing the innovation efforts. Taken in totality, the barriers/impediments that engender the risk aversion are akin to overgrown or out-of-control weeds that can starve or inhibit plant growth within a garden.

66 Lithium batteries were often cited by scientists and engineers as being overly subjected to administrative safety controls that stifled innovation efforts. Whereas the commercial sector worked with engineers to attain reasonable requirements that maintained safety while allowed R&D efforts to continue apace, the NRAC deemed the Navy did not do so. Additionally, scientists and engineers reported their UAS R&D efforts were hamstrung by flight clearance safety policies that were designed for manned aircraft R&D and T&E that should not apply to unmanned aerial systems. Furthermore, there is concern that well-meaning regulations contained within the DoDD 3000.09H, which governs the policy that control Lethal Autonomous Weapons Systems (LAWS) development, could hinder the R&D of military AI technology. U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, p. 28-29; and Daniel S. Hoadley and Nathan J. Lucas, *Artificial Intelligence and National Security* (Washington, DC: Congressional Research Service), 2018, p. 6.

67 U.S. Naval Research Advisory Committee, *Autonomous and Unmanned Systems in the Department of the Navy*, p.28.

The R&D ecosystem is unable to nimbly leverage emerging technologies

The Navy cannot nimbly pivot itself and the R&D ecosystem to leverage cutting-edge technologies emerging from industry and academia and evaluate them for further development to support incorporation into future systems or prototypes. Although DARPA and ONR have a very successful history of incentivizing and nurturing nascent technologies through their R&D funding programs, both organizations' mission and mandate don't necessarily provide the agility to rapidly collaborate with the innovators of the new technology, evaluate its potential value to UxS prototyping, and surge resources and focus among the R&D ecosystem to continue advanced research and development.

The Naval Research Advisory Council's (NRAC) report on Autonomous and Unmanned Systems in the Department of the Navy (DoN) captured these organizational shortcomings aptly. It identified several key issues of concern for the DoN to include: (1) Naval R&D sector not incentivized to operate at private sector speed; (2) not sufficiently connected to new technology; and (3) attempts to recreate Silicon Valley capabilities within Naval R&D sector that fails to produce cutting edge autonomy advances like the civilian sector.⁶⁸ In summary, the NRAC contends the DoN is slow to adapt to the modern paradigm of the civilian sector leading the development of innovative AI and autonomy advances (which are relevant to future warfighting capability), and if the DoN does not learn to adapt to this shift, it risks strategic failure in potential future conflicts with peer/near-peer competitors like China.⁶⁹ Since this report, the Navy has established NavalX and Tech Bridges to realize improved collaboration with the commercial sector, and these innovations are a positive step in the right direction.⁷⁰ Additionally, the Advanced Naval Technology Exercise (ANTX) series exercises have provided opportunities for industry, academia, and government R&D organizations "to demonstrate emerging technologies and engineering innovations in operationally relevant environments and scenarios."⁷¹ The Navy efforts should strengthen and expand on these efforts to more concretely and rapidly leverage the commercial and academic sectors' AI and autonomous/unmanned systems expertise.

68 Ibid., p. 2.

69 Ibid.

70 Megan Eckstein, "NavalX Innovation Support Office Opening 5 Regional 'Tech Bridge' Hubs," *USNI News*, September 3, 2019; Megan Eckstein, "Navy Adds 6 More Tech Bridges to Expand Network of Innovators," *USNI News*, May 12, 2020; and NavalX, *Tech Bridges 2020 Annual Report* (Arlington, VA: Office of the Secretary of the Navy), 2020, available at <https://www.secnav.navy.mil/agility/Pages/techbridges.aspx>

71 Each ANTX exercise, or series of exercise(s), is focused on Navy and Marine Corps mission priorities and gaps. The ANTX exercises also allow warfighters to access the utility of the technologies and provide real-time feedback to the innovators. The operational and technical assessments from the ANTX provide risk assessment for larger Fleet/USMC exercises and inform investment decisions. Naval Research and Development Establishment, *Advanced Naval Technology Exercise (ANTX)* (Arlington, VA: Office of the Secretary of the Navy), available at <https://www.secnav.navy.mil/agility/Pages/ANTX.aspx>

The Artificial Intelligence specialty, Machine Learning, provides an instructive example of the inflexibility and inability of the Navy's current UxS R&D framework to quickly pivot to the pursuit of new technologies. Machine Learning has been around for decades, but R&D progress has accelerated only within the last 5-10 years, in large part enabled by advancements in both hardware (e.g. graphical processing units which deliver dramatically faster speeds) and software algorithms. However, despite being a decade into this technological revolution, the Navy has only begun to investigate how it can harness it to enhance its operational and warfighting capabilities.

Several recent breakthroughs in Machine Learning exemplify the strong potential for its applicability to future operational and warfighting capabilities — the successes of Google Deepmind's AlphaGo/AlphaGo Zero and AlphaStar systems, DARPA's AlphaDogfight competition, and Carnegie Mellon University's Libratus system. These systems both have particular applicability to naval unmanned autonomous vehicles in a way similar to how Machine Learning systems comprise the central component in many driverless vehicles currently being tested by Google Waymo, Uber, Tesla, Nissan, and others. In 2016, AlphaGo defeated the Go world champion, a shocking event that no one had predicted for at least a decade in this ancient, notoriously complex game.⁷² "During the games, AlphaGo played a handful of highly inventive winning moves...so surprising they overturned hundreds of years of received wisdom, and have since been examined extensively by players of all levels... AlphaGo somehow taught the world completely new knowledge about perhaps the most studied and contemplated game in history."⁷³

Go is considered to be a game of such complexity that for several millennia its study has been considered essential to understanding the nature of strategy in conflict. Starting at a novice level, AlphaGo's developers trained it on a compendium of past Go game records (kept for centuries) by practitioners. Using a combination of human-supervised learning and its own reinforced learning, it developed grand champion expertise that allowed it to win against the best human Go player in the world.⁷⁴

Marking another milestone in AI, in 2016, Carnegie Mellon University's Libratus AI system defeated a group of Texas Hold'em Poker champions for the first time in history. This

72 The game of Go originated in China 3,000 years ago. The rules of the game are simple: players take turns to place black or white stones on a board, trying to capture the opponent's stones or surround empty space to make points of territory. As simple as the rules are, Go is a game of profound complexity. There are an astonishing 10 to the power of 170 possible board configurations — more than the number of atoms in the known universe — making Go a googol (10^{100}) times more complex than Chess.

73 Deep Mind, "The Story of AlphaGo so far," <https://deepmind.com/research/case-studies/alphago-the-story-so-far>

74 In *supervised learning*, each training data element is explicitly labeled as an input-output pair, where the output is the "correct" desired value that one wishes the system to learn to associate with a given input (thereby learning the general rules by which to associate input-output pairs not in the original training set), and the "output" represents a "supervisory signal"). In *unsupervised learning*, the system attempts to discover hidden structure in data on its own — i.e., no reward signals are given to "nudge" the system as it processes the training data. Ilachinski, *AI, Robots, and Swarms*, p. 49.

achievement was notable in that Texas Hold'em is a game in which players have missing or incomplete data about the other players' cards and intentions.⁷⁵ Unlike in Go, where both players have perfect information (because they can observe the game board), Texas Hold'em participants have to contend with uncertainty and misleading behavior from other players seeking to maximize their probabilities to win. As one article explained, Libratus can perform strategic reasoning with only imperfect information at a level surpassing humans in which it can now successfully bluff and beat them. Carnegie Mellon's head of the Computer Science Department offered that Libratus' success underscores the potential of its applications in any area in which information is incomplete and opponents spread misinformation to include business, negotiation, military strategy, and cybersecurity.⁷⁶

Most recently, Google's DeepMind achieved another significant milestone when its AI, AlphaStar, attained Grandmaster level in StarCraft II, a real-time strategy video game, by using multi-agent reinforced learning, scoring victories that ranked it over 99.8% of all active players.⁷⁷ This achievement is noteworthy because of StarCraft II's substantial complexity, heavy reliance on strategic thinking to succeed, decision-making with imperfect information, and 10^{26} actions per turn, which far exceeds that of Go.⁷⁸

In August 2020, DARPA's Air Combat Evolution (ACE) program hosted the AlphaDogfight AI competition, a culmination of a yearlong effort in which eight teams from industry, labs, and other research organizations were invited to field AI pilots (operating an F-16 flight simulator) to compete against each other in a tournament to determine the best program.⁷⁹ The winner competed against a human pilot flying an F-16 flight simulator, emerging victorious over the human pilot, 5-0 in individual dogfight competitions.⁸⁰ Heron Systems developed the successful AI pilot by implementing an algorithmic agent using deep

75 Poker requires reasoning and intelligence that has proven difficult for machines to imitate And no-limit Texas Hold'em is especially challenging because an opponent could essentially bet any amount. Will Knight, "China's AI Awakening," *MIT Technology Review*, October 10, 2017.

76 Katherine Arcement, "In A Pittsburgh Casino, Artificial Intelligence Is Beating Poker Champions For The First Time," *Washington Post*, January 29, 2017.

77 Deep Mind Alpha Star Team, "Alpha Star: Grandmaster Level in Starcraft II Using Multi-Agent Reinforcement Learning," *Deep Mind Blog*, October 30, 2019, <https://deepmind.com/blog/article/AlphaStar-Grandmaster-level-in-StarCraft-II-using-multi-agent-reinforcement-learning>

78 Dan Garisto, "Google AI Beats Top Human Players at Strategy Game *Starcraft II*," *Nature*, October 30, 2019, available at [https://www.nature.com/articles/d41586-019-03298-6#:~:text=An%20artificial%20intelligence%20\(AI\)%20known,of%20the%20region's%2090%2C000%20players](https://www.nature.com/articles/d41586-019-03298-6#:~:text=An%20artificial%20intelligence%20(AI)%20known,of%20the%20region's%2090%2C000%20players).

79 Defense Advanced Research Projects Agency, "AlphaDogfight Trials Foreshadow Future of Human-Machine Symbiosis," August 26, 2020, available at <https://www.darpa.mil/news-events/2020-08-26>; and Patrick Tucker, "An AI Just Beat a Human F-16 Pilot In a Dogfight — Again," *Defense One*, August 20, 2020, available at <https://www.defenseone.com/technology/2020/08/ai-just-beat-human-f-16-pilot-dogfight-again/167872/>

80 Patrick Tucker, "An AI Just Beat a Human F-16 Pilot In a Dogfight — Again."

reinforcement learning and then training it to a high level of advanced aerial combat proficiency that ultimately allowed it to defeat the human pilot.⁸¹

The AI breakthroughs by AlphaGo, AlphaDogfight, Libratus, and AlphaStar are potent examples of boundary-breaking R&D the Navy should be pursuing either directly or through its UxS R&D ecosystem. They have direct applicability to the development of advanced warfighting autonomy systems for manned systems (platforms, weapons, and C2 systems from strategic to tactical levels) and unmanned vehicles. These technological advances do not necessarily translate into immediate capability deployment into systems today; they still require additional research to attain that milestone. However, the Navy should pursue such R&D advances via assignment to its R&D ecosystem for continued parallel development and collaboration with industry & academia while current, traditional systems (which are closer to fielding) are still being advanced. Unfortunately, despite industry and academia leading the way in certain fields, the Navy is not organizationally structured in a way necessary to nimbly leverage these breakthroughs.

Better collaboration is needed between Fleet operators and the R&D ecosystem

There is a chasm or disconnect between what the R&D ecosystem believes it needs to develop and what the Fleet assesses they require to operate. During engagement visits with the UxS R&D ecosystem, many of these organizations complained about their lack of access to the Fleet. Although reassuring that they recognized their primary customer was the Fleet, this insufficient access manifested itself in the ecosystem's lack of understanding about the warfighting and operational capability requirements/needs of the Fleet as well as continued interactions with operational concept develops like the Naval Warfare Development Command (NWDC) and OPNAV N5. Without clear guidance, they sometimes embarked upon R&D activities, which could be considered "solutions in search of problems." On other occasions, advanced capabilities they demonstrated (unconsciously) assumed the Fleet operators would reinvent how they conduct warfare operations to incorporate the new technology. They did not necessarily address warfighter needs to include employability within the current Concept of Operations (CONOPS).

On the Fleet operator side, associated Fleet staffs and commands did not always recognize the value of their proactive engagement with the UxS R&D ecosystem. In some instances, established and regular collaboration with the ecosystem was considered a low priority, wasteful of precious time and typically "someone else's" job. They were unable to see the long-term view and strategic benefit of providing early input and guidance to shape R&D technological advancements they would need in the future. Indeed, the time demands on Fleet operators are formidable, so it is understandable why they would perceive long-term

81 Heron Systems, Inc., "Alpha Dogfight Trials," available at <https://heronsystems.com/work/alpha-dogfight-trials/>; and Patrick Tucker, "An AI Just Beat a Human F-16 Pilot In a Dogfight — Again."

R&D efforts to be less important than current maintenance, training, tactical development, and readiness challenges. However, in the absence of a mutual partnership, this disconnect creates a vicious cycle in which ecosystem organizations continue producing capabilities that the Fleet finds unappetizing and unfulfilling to their operational needs, causing both sides increasing frustration with little progress (thus slowly eroding technological competitive advantage). The dissonance is at its sharpest during T&E or other demonstration events where the Fleet operators do not embrace prototypes and advanced capabilities developed by the ecosystem.

CHAPTER 4

Organizing to Improve Navy UxS Development

The lacking level of orchestration across the R&D ecosystem is indicative of a key organization/framework that is missing. The Navy's current UxS R&D bureaucracy (which includes the ecosystem) could benefit from a more streamlined, clear, and centralized organizational construct. Within the Navy, responsibility for various aspects of UxS development involves a multitude of stakeholders. Due to the numerous offices, organizations, commands, and staffs with overlapping authorities and interests (which sometimes compete or conflict), the leadership must generate an incredible amount of momentum to overcome the bureaucratic inertia and revector the Navy in a new direction. The inefficiencies derived from this bureaucracy impede scientists, engineers, and leaders in the ecosystem from accelerating at a much faster pace of R&D, prototyping, and testing.

However, given the lack of a centralizing strategy and organization to optimize the entire UxS R&D ecosystem effort (Navy, FFRDCs, UARCs, industry, academia, etc.), the current effort, despite improved coordination and collaboration action implemented by the Navy, may not be up fully up to the task to deliver on the mission. The successes the Navy has produced to date (some of which were previously discussed) are remarkable and are made even more so by the challenges of the existing bureaucratic thickets that organizations muscled their way through to achieve the alignment among various organizations needed to deliver progress. Using the garden metaphor again in which the scientific and engineering innovation are the flowers and vegetables, the aforementioned bureaucracy operates as the overgrown weeds, which starve the crop. This bureaucratic friction traps significant amounts of potential energy (innovation, resources, advancements, progress), sapping the pace of Navy UxS developmental progress.

As discussed in detail in earlier passages, the bureaucracy and culture of risk aversion have produced an UxS R&D ecosystem that is hampered from achieving its maximum performance. This is further aggravated by the overall lack of top-level orchestration, organization,

and collaboration across the ecosystem of which the disjointed progress of UxS is indicative. Although the Navy leadership has begun to address this, the bureaucracy of the R&D ecosystem impedes their effort at achieving that progress. However, despite the leadership's significant efforts, the Navy's UxS R&D ecosystem may not be effectively organized to win, especially compared to the nearest peer/near-peer challengers.

Faced with a similar dilemma, most corporations simply reorganize the relevant internal activities within the organization to get the needed realignment to achieve the desired synergy. The Navy should consider executing a targeted reorganization and realignment within the relevant activities to achieve its own synergy in the R&D, prototyping, testing, and operationalization of unmanned autonomous vehicles and AI. As discussed earlier, the Navy leadership has implemented numerous changes to begin addressing this issue. However, more changes are required if it desires to remain the preeminent maritime force through harnessing autonomy advances to deliver UxS and vehicles.

When looking to solve difficult problems of the present and future, it is often useful to look to the past for guidance as something similar will likely have been experienced. The Navy has grappled in its past with how to harness cutting-edge research and produce transformative operational capability. In this case, three historical naval examples are instructive: naval nuclear power, submarine-launched ballistic missiles, and the Aegis Weapon System. In all three cases, the Navy created robust cross-functional, interdisciplinary organizations consisting of personnel from the military, government civilian service, industry, and academia that were given a broad, strong mandate and authority to research, develop, prototype, and deliver transformational strategic capabilities.

Outside of the U.S. Navy, other similar historical examples underscore the benefit of establishing such an organization. These include the Manhattan Project and the NASA Mercury & Apollo Project missions. Similar to the Nuclear Power, Polaris, and AEGIS Weapon System projects, these projects included robust cross-functional, interdisciplinary organizations from the military, government civilian service, industry, and academia that were given a broad, strong mandate and authority to research, develop, prototype, and deliver transformational strategic capabilities.

In the 1950s, Admiral Hyman G. Rickover established a cross-functional, interdisciplinary organization that delivered the revolutionary development of compact nuclear power plants to power submarines (and later to surface ships and aircraft carriers). This organization consisted of not only military personnel, but also personnel from government civil service, academia, and industry.⁸² Naval Nuclear power was a transformative leap in warfighting advantage for the United States, especially as it related to the strengthening of strategic deterrence enabled the creation of nuclear-powered ballistic missile submarines and power

82 Francis Duncan, *Rickover: The Struggle for Excellence* (Annapolis, MD: Naval Institute Press), 2001.

projection through the near-unlimited range of aircraft carriers that no longer needed frequent refueling.

During the late 1950s to early 1960s, then Rear Admiral William “Red” Raborn was directed by the Chief of Naval Operations, Admiral Arleigh Burke, to establish his own cross-functional, interdisciplinary organization to develop a submarine-launched (nuclear) ballistic missile (SLBM) that would form the basis of the Polaris project.

The Navy established a special projects office directed by Admiral Raborn to develop what eventually became the Polaris missile program. At the start of the program, existing nuclear warheads were too large and heavy to fit on submarines, but Raborn’s project office overcame numerous complex technical challenges to develop a more compact nuclear weapon design that scaled to submarines and utilized a safer propellant.⁸³

The Polaris naval project, SLBMs, and nuclear-powered ballistic missile submarines (SSBNs) became the most survivable leg of the strategic triad, a role they still fulfill to this day. Furthermore, the SLBM was hailed as an American triumph in the cold-war arms race and as a boon to the country’s allies, whose acceptance of American nuclear missiles on their soil was a constant source of friction with the Soviet Union.⁸⁴



From left to right: Rear Admiral William “Red” Raborn, Admiral Hyman G. Rickover, and Rear Admiral Wayne E. Meyer. All photos courtesy of the U.S. Navy.

In the 1970s through early 1980s, Rear Admiral Wayne E. Meyer led his own cross-functional, interdisciplinary organization, the Advanced Surface Missile System (ASMS) office (1970-83) that sought to combine the advanced capabilities of cutting edge phased-array radars; the networking of the most powerful computers at the time to calculate complex fire control solutions and process large amounts of inputs and information; lethally effective surface-to-air missiles and the latest designs in optimized human-computer interaction

83 Amaani Lyle, “Work Chronicles History at Fleet Ballistic Missile Program Event,” *DoD News Features*, October 3, 2015.

84 Michael Wines, “William F. Raborn Is Dead at 84; Led Production of the Polaris Missile,” *New York Times*, March 13, 1990.

and automation to maximize watch stander performance.⁸⁵ This office was responsible for the development of the Aegis Weapon System. The ASMS office eventually transitioned into the Aegis Shipbuilding Project Office, which he led from 1976-1983.⁸⁶ Vice Admiral (Retired) James H. Doyle emphasized the revolutionary impact that the Aegis Weapon System delivered through Rear Admiral Meyer's leadership and vision during a 2001 seminar panel.

This monumental endeavor required over a decade of effort and ultimately produced the Aegis Weapon System, considered to be the first fully integrated radar, data, fire control, and weapons system that revolutionized warfighting at sea when the *Ticonderoga*-class cruisers (and later the *Arleigh Burke*-class destroyers) were delivered to the Fleet.

The Aegis Weapon System in 1983 represented a true revolution in shipboard air defense. Based on an enormous investment in time, resources, and management focus, Aegis was the first truly integrated ship-based system. It brought together radar and sensor detection, tracking, and missile interception into a coherent, well-integrated weapon system. This system was a staggering engineering achievement and the culmination of nearly 40-years of Navy experience in confronting increasingly dangerous air defense challenges, beginning with kamikaze attacks in the waning months of World War II and extending to Soviet Backfire bombers in the 1970s and 1980s.⁸⁷

Over thirty years later, the Aegis Weapon System remains the warfighting tactical system mainstay of the U.S. Navy's surface combatant fleet.

85 Navy Marks Passing of Rear Admiral Wayne E. Meyer," *Navy News Service*, September 1, 2009; and US Naval Institute, "Meyer, Wayne E., Rear Adm., USN (Ret.) (1926–2009)," available at <https://www.usni.org/press/oral-histories/meyer-wayne>

86 Ibid.

87 Robert Holzer and Scott C. Truver, "Not Your 'Father's Aegis'," *Center for International Maritime Security*, November 21, 2014, available at <https://cimsec.org/not-fathers-aegis/>

CHAPTER 5

Toward an Autonomy Project Office

To best deliver UxS, autonomy, and AI capabilities to the Fleet, the U.S. Navy should apply the concept and framework of these past successes to its current situation. The Navy should establish a cross-functional, interdisciplinary Autonomy Project Office (APO) to pursue the research, development, prototyping, testing & evaluation, and operationalization of transformative, revolutionary technologies for UxS.⁸⁸ The main mission of the APO would be to unify oversight, authority, and direction of all the research and development lines of effort across the R&D ecosystem for UxS, and continuously spiraling the latest technological advances into an assembly line of autonomous unmanned air, surface and undersea vehicle prototypes that can be demonstrated for warfighting value, utility, and viability. Contracting mechanisms, such as Other Transaction Authorities (OTAs) can be employed to nimbly and swiftly fund ecosystem R&D and prototype vehicle/system construction, which has historically made this authority very attractive to innovative technology companies, universities, and other R&D organizations.⁸⁹ Those vehicles deemed particularly effective can be either pursued in low production rate using OTA to fulfill imminent warfighter needs or obtained through the regular acquisition process as major programs of record.

88 An APO vice an Artificial Intelligence Project Office is recommended because autonomy and autonomous systems are the useful enabling technology, practical military application or instantiation of the AI scientific discipline. Also considered, but rejected was the title Unmanned Vehicle Project Office because although this research study focused on that effort, the APO as discussed in the following passages is structured to support the R&D of AI technological advancements that can be incorporated into manned platforms (e.g. ships, aircraft, submarines), weapon systems and theater/operational battle management command and control decision support systems.

89 An OTA is a special vehicle used by federal agencies for obtaining or advancing R&D or prototypes. Government procurement regulations and certain procurement statutes do not apply to OTs, and accordingly give agencies the flexibility necessary to development agreements tailored to a particular transaction. Elaine L. Halchin, *Other Transaction (OT) Authority* (Washington, DC: Congressional Research Service), 2011, pp. 1,7.

Of note, the APO should focus on developing hybrid UxS Human-Supervised/Human-on-the-Loop (HOTL) UxS prototypes as defined in DoDD 3000.09.⁹⁰ In such a hybrid blended UxS, human operators will have defined and pre-selected a limited subset of actions and/or behaviors that are reserved for human approval before the UxS will execute them. These actions could include changing the primary mission objective, breaking off a patrol to follow a specific contact, or engaging an adversary platform classified as hostile based on the current rules of engagement. All other actions and behaviors not delineated within this subset would be available for the UxS to perform without human intervention.

Technologically and philosophically, the U.S. Navy is not yet close to fielding fully autonomous vehicles and systems. Although R&D should certainly continue striving, it is still a far distance away from the true instantiation of such autonomy. Conversely, the military's current inventory of vehicles is proliferated with Semi-Autonomous/Human-in-the-Loop systems (e.g. Predator, Global Hawk, etc.) that have been employed extensively on the battlefield and throughout operational theaters for over two decades.⁹¹ Although these systems have transformed how the U.S. military monitors, shapes, and engages the theater, they also come with inherent disadvantages such as large manpower footprints for control, delayed response, and limited flexibility to adjust to dynamic circumstances. Pursuing HOTL unmanned vehicles will negate some of these disadvantages, unleash the advantages that greater autonomy avails to these systems while allowing the human operators to maintain control over the critical and potentially lethal actions/behaviors of the vehicle. In a sense, HOTL is very similar to what human military forces operate under so this construct would be familiar to those operating such UxS. Full, unfettered autonomy (for humans) does not exist within the military C2 decision-making hierarchy. Instead, for a certain subset of actions, higher headquarters retains the requirement for a unit's commander or commanding officer to seek permission or approval to conduct those actions. For the remainder of available actions/decisions, the unit commander is granted the autonomy to conduct them.

90 Human-Supervised/Human-on-the-Loop is a level of autonomy in which an autonomous weapon system that is designed to provide human operators with the ability to intervene and terminate engagements including in the event of a weapon system failure, before unacceptable levels of damage occur. Department of Defense, *Department of Defense Directive Number 3000.09*, pp. 13-15.


91 Semi-Autonomous/Human-in-the-Loop is a system where once the weapons system is activated, it engages only those targets that have been selected by a human operator. Ilachinski, *AI, Robots, and Swarms*, p. 6. Department of Defense, *Department of Defense Directive Number 3000.09*, pp. 14.

Autonomy Project Office Organizational Design

To accomplish its primary mandate, the APO does not need to duplicate the excellent innovation and existing research, development, science, engineering, systems integration, and testing & evaluation organizations that already exist within the UxS R&D ecosystem. Doing so would be a waste of resources and talent that the Navy can ill afford. Instead, the APO will exist to better facilitate orchestration, collaboration among them to produce unity of effort/direction, and delivery of UxS prototypes. It should be established as a cross-functional interdisciplinary organization consisting of scientists, engineers, experts, and managers from all geographic corners of the UxS R&D ecosystem. Commensurate with its accountability and responsibility, the APO should be invested with the authority to direct the ecosystem in all matters pertaining to UxS, autonomous and AI research, development, prototyping, and testing & evaluation, to include funding assignment. In pursuing its mission to accelerate prototyping, the APO's leadership and staff will need to demonstrate a deft touch, whereby they facilitate the appropriate amount of orchestration, collaboration, and synchronization among the ecosystem to achieve increased unity of effort/direction without stifling innovation.

The NAVSEASYSOM program office, PMS-406 Unmanned Maritime Systems, which is located within the renamed PEO Unmanned and Small Surface Combatants, would serve as an excellent initial core staff and foundation to the establishment of the APO. Although a small office, for over a decade, it has served as ground zero for the prototyping and acquisition of unmanned surface and undersea vehicles and has developed a deep level of expertise and experience with the UxS R&D Ecosystem, industry, ONR, DARPA, OPNAV, and the Secretariat. The office's successes with the XLUUV RFP, Sea Hunter MDUSV, Knifefish, and follow-on Future Surface Combatant Unmanned Surface Vehicle (FSC-USV), Ghost Fleet SCO project, and several other projects underscore the importance of integrating them into the APO.⁹² The UxS centric components of NAVAIRSYSCOM's PMA-281 (Strike Planning and Execution Systems), PMA-268 (Unmanned Carrier Aviation), PMA-262 (Persistent Maritime UAS), PMA-266 (Multi-mission Tactical UAS), and PMA-263 (Small Tactical UAS) program offices could also leverage their successes and find increased opportunity through realignment into the APO for advancing the UAS technological and warfighting capabilities on current and future platforms through closer marriage with both the AI/mission autonomy and the cutting edge UAS R&D occurring within the ecosystem. There are other organizations and offices involved in UxS related work that could either be integrated directly into the APO or aligned with it to support APO's mission and mandate.

92 Lee Hudson, "Pentagon to host industry day for SCO Ghost Fleet project," *Inside Defense*, September 25, 2017.

FIGURE 1: AUTONOMY PROJECT OFFICE PROPOSED ORGANIZATION

Source: Author

To be successful, this cross-functional, interdisciplinary organization should be organized along the previously discussed leadership structure and LOEs as follows:

APO Director. The APO Director billet should be designated a vice-admiral billeted position to ensure the director is vested with sufficient positional and rank authority to accomplish his or her demanding responsibilities. This will allow the director to coordinate, collaborate and operate across the critical three-star Deputy CNO positions, Systems Commanders, and Type Commanders that comprise a sizable portion of the Navy's senior flag leadership. For organizations outside of the Navy that reside within the R&D ecosystem, the APO Director should be assigned the requisite authority to issue contracts and other memoranda to support the employment of these organizations to pursue UxS R&D LOEs. For organizations and commands within the Navy that do not fall under the APO's direct administrative control (ADCON), it is recommended the APO be invested with Direct Liaison Authorized (DIRLAUTH) and provisional/situational ADCON to support the efficient conduct of APO R&D, prototyping and T&E activities. To ensure project office stability, unity of effort/focus, and sustained efficacy, the APO Director should be assigned to the position for a minimum eight- to ten-year duration. Similar to the Naval Reactors four-star flag position, this ostensibly means the director will be removed from the typical flag billet rotation of approximately every two years.

Admirals Rickover and Meyer and naval aviation pioneer, Admiral Moffett benefitted from long-term assignments to their positions during the development of nuclear-powered submarines & surface ships; the Aegis Weapon System integrated into cruisers and

destroyers; and carrier naval aviation; respectively, which allowed them to provide consistent and strategic leadership across their organizations. Ultimately, they were able to articulate a vision and goal for their organizations, inculcate that vision into the workforce and successfully implement it, delivering successful prototypes that became the basis for revolutionary and transformative capabilities.

The APO Director should have dual-reporting responsibility to both the Assistant Secretary of the Navy for Research, Development & Acquisition (ASN RD&A) and the Vice Chief of Naval Operations (VCNO). By placing the APO external to OPNAV and the SYSCOMs, such an organizational relationship will help firewall the APO from the bureaucratic, budgetary, and Navy cultural turf fights and impediments that could unduly detract from the APO's primary mission. The four-star fleet commanders should have DIRLAUTH to the APO Director to ensure their operational capabilities/requirements requests can be directly transmitted. To further imbue independence, the APO should have a set-aside budget of its own.

The APO Director will be charged with forging a brand new, Navy-wide, cross-functional, interdisciplinary organization consisting of a mosaic of commands, staffs, organizations, scientists/engineers, and other experts, which is a daunting task. The closest equivalence in size to this endeavor in recent memory was the establishing of the Naval Surface & Mine Warfighting Development Center (SMWDC), Undersea Warfighting Development Center (UWDC), and Information Warfighting Development Center (IWDC), but even these organizations are small in scale as compared to what is being proposed with the APO. The F-35 Joint Program Office, which contained six program offices within it, is an example albeit a joint one.

TABLE 1: APO DIVISION TITLES & ROLES (PART 1)

APO Division Title	APO Division Role
Navigation Control	Lead, orchestrate, guide and direct the research, development and prototyping of advanced navigation systems and algorithms for integration into unmanned air, surface and undersea systems.
Mission Autonomy	Lead, orchestrate, guide and direct the research, development and prototyping of mission autonomy systems that will control the operation of UxS, integrate the various subsystem components (e.g. sensors, communications, payload, energy endurance/propulsion, etc.) and facilitate the human-machine teaming interface to accomplish assigned operational tasks and missions.
Prototype Design / Systems Integration & Assembly	Lead, orchestrate, guide and direct the R&D efforts and actual production/manufacturing of fully-instantiated prototypes to include integration of approved LOE prototype systems. This division will also conduct the operational and functional checks to ensure the prototype is ready for transfer to the Test & Evaluation Division for rigorous assessment.
Testing & Evaluation	Lead, orchestrate, guide and direct the R&D, analysis and formal requirements design to support conducting the comprehensive verification, validation, accreditation, certification and trust-building assessments of unmanned autonomous vehicles, autonomous systems, AI components/systems, Machine/Deep Learning systems, etc. These T&E efforts will occur not only in labs and centers, but also will include live demonstration and testing on aviation and maritime ranges as applicable.
Fleet Liaison	Facilitate direct communication and collaboration with the APO's primary customers, the four-star operational Fleet Commanders, by ensuring key operational staffs and organizations can provide input/guidance to the APO on desired UxS prototype warfighting requirements, capabilities, needs, envisioned missions and tactical deployment. This division will assist the APO in ensuring the LOE R&D within the constellation and the production of prototypes are geared toward the Fleet's operational and tactical needs.

Source: Author

APO Deputy Directors. Two Senior Executive Service (SES) or equivalent Deputy Directors will report to the APO Director. The Deputy Directors should be recognized subject matter experts with broad expertise in the areas in which they are charged to oversee as well as experienced leaders and large project managers.

The Deputy Director for UxS Delivery will lead the overall R&D, systems integration, and prototype manufacturing/production. This director will be assigned leadership over the following divisions: Mission Autonomy, Navigation Control, Communications, Sensors, Payloads, Kinetics, Propulsion & Energy Endurance, Cyber Security & System Hardening, and Prototype Design & Systems Integration. The Deputy Director for UxS will robustly engage with the UxS R&D Ecosystem to leverage technological advances, expertise, and experience resident therein.

The Deputy Director for UxS Operational Development & Evaluation will lead the overall Test & Evaluation of the UxS/unmanned autonomous vehicles, Human-Machine Teaming R&D/concept development, and liaison engagement with the key Fleet and Navy Staffs. This director will be assigned leadership over the following divisions: Human-Machine Teaming,

Fleet Liaison, and Test & Evaluation. The Deputy Director for UxS Operational Development & Evaluation will robustly engage with the UxS R&D Ecosystem to leverage technological advances, expertise, and experience resident therein for T&E and Human-Machine Teaming activities. The Deputy Director will also liaison significantly with 4-star Fleet Commander staffs, Type Commanders, OPNAV N9, System Commanders, Operational Test & Evaluation staff, Navy Warfare Development Command, and the four Warfighting Development Centers to maximize two-way communication between these commands and the APO.

APO Division Leadership. The selection of the Division Directors who will be in charge of the LOE Divisions is equally vital to the choosing of the APO Director and Deputy Directors. It is within these divisions that the critical work the APO must accomplish will occur, namely coordination, collaboration, and orchestration of LOE R&D across the ecosystem to maximally harness its innovation. To achieve the vision of a cross-functional, interdisciplinary organization, this mandate must begin at the deputy director division lead level. Division Directors should be selected from distinguished leaders from industry, academia, UARCs/FFRDCs, Navy & DoD, and labs and warfare centers who are subject matter experts in the LOE R&D areas in which they would be assigned to lead. Additionally, Navy civilian and military personnel would be chosen for positions within the APO based on their expertise, experience, and leadership whether it be in UxS-related disciplines, supporting fields (e.g. logistics, acquisition, etc.), or warfighting/tactical areas (related to CONOPS development, tactical application, etc.).

Similar to the APO Director and Deputy Directors, the Division Directors should possess exceptional leadership, management, and communication skills, especially as it relates to working on large projects and with science and engineering experts from a multitude of diverse organizations. Additionally, the Division Directors should have a deep level of expertise and experience within their LOE area as their primary responsibility is to effectively communicate with the ecosystem organizations within their assigned LOE to leverage the latest technological advancements and breakthroughs for R&D pursuit in designing UxS.

Mission Autonomy Division. The role of the Mission Autonomy Division is to lead, orchestrate, guide, and direct the research, development, and prototyping of mission autonomy systems that will control the operation of UxS, integrate the various subsystem components (e.g. sensors, communications, payload, energy endurance/propulsion, etc.) and facilitate the human-machine teaming interface to accomplish assigned operational tasks and missions. This division's work is primarily underpinned by the Artificial Intelligence discipline to include all of its subcategories such as Machine Learning, Deep Learning, Reinforced Learning, Multi-Agent Systems, Cognitive Computing, Data Science, Data Engineering, etc. and includes the collection, labeling, and curation of data for training and testing of autonomous systems.

Navigation Control Division. The role of the Navigation Controls Division is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced navigation systems and algorithms for integration into unmanned air, surface, and undersea systems.

Prototype Design / Systems Integration & Assembly Division. The role of the Prototype Design/Systems Integration & Assembly Division is to lead, orchestrate, guide, and direct the R&D efforts and actual production/manufacturing of fully-instantiated prototypes to include the integration of approved LOE prototype systems. This division will also conduct the operational and functional checks to ensure the prototype is ready for transfer to the Test & Evaluation Division for rigorous assessment.

Test & Evaluation Division. The role of the T&E Division is to lead, orchestrate, guide, and direct the R&D, analysis, and formal requirements design to support conducting the comprehensive verification, validation, accreditation, certification, and trust-building assessments of unmanned autonomous vehicles, autonomous systems, AI components/systems, Machine/Deep Learning systems, etc. These T&E efforts will occur not only in labs and centers, but also will include live demonstration and testing on aviation and maritime ranges as applicable. The T&E Division will coordinate with Commander Operational Test & Evaluation Force (COMOPTEVFOR) and the OSD Director, Operational Test & Evaluation (DOT&E).

Fleet Liaison Division. The role of the Fleet Liaison Division is to facilitate two-way direct communication and collaboration with the APO's primary customers, the four-star operational Fleet Commanders, by ensuring key operational staffs and organizations can provide input/guidance to the APO on desired UxS prototype warfighting requirements, capabilities, needs, envisioned missions and tactical deployment. This division will assist the APO in ensuring the LOE R&D within the ecosystem and the production of prototypes are geared toward the Fleet's operational and tactical needs. It will also play a significant role in the R&D efforts for the Human-Machine Teaming LOE and the Testing & Evaluation Division's efforts, thereby keeping APO's efforts relevant to the Fleet. Furthermore, this division will facilitate communications that allow the Fleet to provide lessons learned and other recommendations/guidance from operations, exercises, demonstrations, and other events that will inform APO activity.

TABLE 2: APO DIVISION TITLES & ROLES (PART 2)

APO Division Title	APO Division Role
Payloads	Lead, orchestrate, guide and direct the research, development and prototyping of advanced payloads for integration in to unmanned air, surface and undersea systems. These payloads will be designed to enhance system capabilities or support accomplishment of assigned missions.
Kinetics	Lead, orchestrate, guide and direct the research, development and prototyping of advanced weapon systems for integration into unmanned air, surface and undersea systems.
Human-Machine Teaming	Lead, orchestrate, guide and direct the research, development and prototyping of the advanced Human-Computer Integration design and systems that will enhance the usability and utility of AI/autonomous systems and unmanned air, surface and undersea systems to the human operators in the field.
Communications	Lead, orchestrate, guide and direct the research, development and prototyping of advanced communications systems for integration into unmanned air, surface and undersea systems.
Sensors	Lead, orchestrate, guide and direct the research, development and prototyping of advanced sensor systems for integration into unmanned air, surface and undersea systems.
Propulsion/Energy Endurance	Lead, orchestrate, guide and direct the research, development and prototyping of advanced propulsion, fuel, battery and other energy endurance technology (e.g. aluminum water combustion, lithium polymer, etc.) for integration into unmanned air, surface and undersea systems.
Cyber Security & System Hardening	Lead, orchestrate, guide and direct the research, development and prototyping of advanced cyber security and physical security to facilitate and strengthen anti-tamper, anti-hacking, anti-exploitation hardening systems for integration into unmanned air, surface and undersea systems.

Source: Author

Propulsion/Energy Endurance Division. The role of the Propulsion/Energy Endurance Division is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced propulsion, fuel, battery, and other energy endurance technology (e.g. aluminum water combustion, lithium polymer) for integration into unmanned air, surface, and undersea systems.

Communications Division. The Communications Division's role is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced communications systems for integration into unmanned air, surface, and undersea systems.

Sensors Division. The Sensors Division's role is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced sensor systems for integration into unmanned air, surface, and undersea systems.

Cybersecurity & System Hardening Division. The Cybersecurity & System Hardening Division's role is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced cybersecurity and physical security to facilitate and strengthen

anti-tamper, anti-hacking, anti-exploitation hardening systems for integration into unmanned air, surface and undersea systems.

Human-Machine Teaming Division. The Human-Machine Teaming Division Communications Division's role is to lead, orchestrate, guide, and direct the research, development, and prototyping of the advanced Human-Computer Integration design and systems that will enhance the usability and utility of AI/autonomous systems and unmanned air, surface and undersea systems to the human operators in the field.

Payloads Division. The Payloads Division's role is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced payloads for integration into unmanned air, surface, and undersea systems. These payloads will be designed to enhance system capabilities or support the accomplishment of UxS missions assigned by the Fleets.

Kinetics Division. The Kinetics Division's role is to lead, orchestrate, guide, and direct the research, development, and prototyping of advanced weapon systems for integration into unmanned air, surface, and undersea systems.

CHAPTER 6

The Case for an Autonomy Project Office: The Advantages

There are some definitive advantages to establishing an APO that will accrue to the Navy's efforts to prototype and operationalize unmanned autonomous vehicles and artificial intelligence (to include all of its various subordinate disciplines like machine/deep learning, cognitive computing, etc.). In 2016 the Defense Innovation Board, chaired by Alphabet CEO Eric Schmidt, recommended the creation of an AI and Machine Learning Center of Excellence inside DoD "to spur innovation and transformational change."⁹³ An organization of this type could also create a single focal point for Congress to consult on defense-related AI issues.⁹⁴ This recommendation was ultimately implemented as the DoD Joint Artificial Intelligence Center (JAIC). The above statement and its stated advantages apply on a smaller scale within the U.S. Navy in the following ways.

Clear, coherent authority and unity of effort.

First and foremost, the creation of an APO will allow for the consolidation of a multitude of dispersed, decentralized UxS efforts into a single organization that can provide unity of direction, unity of effort, clear organizational C2, centralization, and authority. This consolidation will clear out the bureaucratic thickets and underbrush resident in the current organizational construct, thus enhancing the Navy's ability to harness the incredible innovation and energy of the UxS R&D ecosystem. By centralizing UxS R&D efforts within the APO, it will possess the authority and mandate to tackle the Hyper-Focus/Starvation and Lack of Coherent Organization Strategy issues that diminish R&D productivity within the ecosystem, unhelpfully squander its finite resources and unwittingly stifle innovation from its critical scientists and engineers.

93 Aaron Mehta, "Defense Innovation Board Lays Out First Concepts," *Defense News*, October 5, 2016.

94 Daniel S. Hoadley and Nathan J. Lucas, *Artificial Intelligence and National Security*, p. 6.

Solidify established UxS gains and make them permanent.

Establishing the APO allows the SECNAV and CNO to solidly affirm the UxS, autonomy, and AI mandate they have imparted throughout the Navy into a more permanent instantiation. Without that permanence anchored by an APO, the SECNAV and CNO's focus and vision can be nullified (either through benign or purposeful neglect) by future leadership. Given the herculean level of effort the leadership dedicated to making modest progress under the current bureaucracy, it will only require a little neglect to begin dissipating the innovation momentum built up to overcome the Navy's anti-UxS organizational inertia. Those gains need to be consolidated and made permanent in an APO.

Bypass institutional resistance and barriers to UxS development

Pockets of institutional and cultural resistance continue to stymie UxS development. In some cases, UxS were perceived to challenge existing norms and bureaucracies within the Navy, which potentially produced organizational reactions ranging from indifference (to UxS potential value) to recalcitrance/antipathy. Depending on the source, positions, and seniority of these beliefs, UxS R&D can be stymied and deprioritized (with other justifications given) to protect the status quo.⁹⁵ APO can overcome institutional resistance and barriers to UxS as it will be largely immune to such parochial concerns, allowing the organization to immerse itself in producing UxS, unimpeded by these distractions. Also, freed from budget competition with competing communities that could lead to bureaucratic & budgetary turf wars, the APO can pursue production of the very best prototypes of UxS warfighting capabilities.

An APO can help eliminate impediments to R&D technology transition for prototyping

In an earlier section, this report asserted the current UxS R&D ecosystem lacked a process to evaluate capability demonstrations for worthiness to transition to further development or prototyping. This resulted in S&T capabilities that were orphaned without evaluation for utilization. This problem has been recognized and studied for several decades, and “in 2007, DoD reported that this gap can only be bridged through cooperative efforts and investments from both communities, such as early and frequent collaboration among the developer, acquirer, and user.”⁹⁶ The APO can address this issue, but before exploring how, it is helpful to refer back to some of the findings from the GAO report on DARPA technology transitions. That report highlighted four factors deemed critical to successfully cutting edge technologies into funded military programs: (1) military or commercial demand for the technology; (2) linkage to an area where DARPA has sustained interest; (3) active collaboration

⁹⁵ Kenneth M. Jensen, *The Cultural Challenges Navy Unmanned Air Systems Pose to Naval Aviation*, (Carlisle, PA: U.S. Army War College, 2017), pp. 9-17.

⁹⁶ U.S. Department of Defense, *Department of Defense Report to Congress on Technology Transition* (Arlington, VA: U.S. Department of Defense, 2007).

with potential transition partners; and (4) achievement of a clearly defined technical goal.⁹⁷ Additionally, the report noted transition success for DARPA R&D was further strengthened by one or more of the following agreements existing between it and (1) a military service, (2) a DoD research agency or laboratory, and/or (3) another warfighter representative that identified a military capability gap or requirement that the technology fulfilled.

As proposed in the mission and organizational design, the APO could demonstrate all four of the above transition success factors in the S&T R&D capability programs that it pursues as part of the development and prototyping of UxS and its component LOEs. Furthermore, S&T R&D activities funded by DARPA, ONR, SCO, and other R&D funding organizations, which yield promising results of interest to APO, but not of value to these other organizations, can also be shifted over to the APO for continued R&D and prototyping.

An APO can provide AI R&D innovation and technology for manned platforms & C5I systems

Previously, this report stated the APO could provide practical AI-infused military applications or solutions in support of manned platforms, weapons systems, and battle management command and control (C2) decision support systems. The Mission Autonomy division will focus its efforts on supporting, nurturing, and harnessing the artificial intelligence scientific discipline toward the R&D of technological capabilities that can be adopted for military systems. From this perch, the division will be led and staffed by expert, accomplished AI engineers and practitioners from across industry and academia and thus will have direct access to cutting-edge breakthroughs.

Systems commands (e.g. NAVAIRSYSCOM, NAVSEASYSYSCOM), Type Commanders (e.g. COMNAVSURFOR, COMNAVSUBFOR, COMNAVAIRFOR, COMNAVIFOR), Warfighting Development Centers (NAWDC, UWDC, SMWDC, IWDC), and the Navy Staff will have input to the Mission Autonomy division as well as the other divisions in support of recommending, articulating, developing, pursuing and budgeting for future naval capabilities and technology requirements. Their ability to harness the intellectual advice and insight from key R&D subject matter experts who are familiar with military/naval organizations and issues will significantly enhance the products developed by these commands and staffs.

An APO can more nimbly pivot towards harnessing revolutionary technological breakthroughs.

The successes of Google Deepmind's AlphaGo (plus AlphaGo Zero), AlphaStar, DARPA's AlphaDogfight, and Carnegie Mellon University's Liberator Texas Hold 'em Poker algorithms were cited as technological breakthroughs that would be difficult for the Navy to harness

97 U.S. Government Accountability Office, *Defense Advanced Research Projects Agency*, pp. 12-13, 15-16.

given its current organization.⁹⁸ AlphaGo and AlphaStar are Machine/Deep Learning systems, which provided the developers some level of transparency as they could monitor game performance learning progress and provide adjustment (i.e. essentially tutoring) along the way. Liberato delivered a first-of-its-kind breakthrough in which the system could win against professional poker champions with incomplete, missing, or incorrect (deceiving, bluffing) data. AlphaDogfight demonstrated an algorithm could defeat a human pilot in simulated fighter-to-fighter aerial combat. These successes highlight significant relevance to the R&D of mission autonomy systems for UxS as well as other military manned system platforms/applications.

Aluminum-based power systems are another potential future breakthrough that could revolutionize battery endurance and fuel systems for underwater use. Aluminum as an energy source has been demonstrated to achieve 2-10 times more energy than other battery technologies.⁹⁹ Imagine a scenario where these cutting-edge advancements from industry and academia can be promptly and rapidly identified by the Navy and then pursued for further UxS R&D and application to future Navy systems. The APO, as proposed and organizationally designed, could quickly identify ground-breaking new technologies, coordinate with the source organizations to leverage the expertise, and then spiral their continued pursuit into the UxS R&D ecosystem, shepherded under the auspice of the appropriate division directors and their staff. While more mature technologies are assigned for integration into upcoming UxS prototype assembly and T&E, these breakthrough technologies can in parallel continue to be evaluated, shepherded, and nurtured for spiraling into future prototypes. Additionally, as was discussed in the technology transition section, R&D cutting edge proof of concept capabilities and technologies sponsored by DARPA, ONR, SCO, and other research funding organizations, which they no longer desire to pursue, but have relevant applications to UxS, can be transferred over to the APO to avoid orphaning and maximize R&D efficiency and return on investment.

98 For example, the AlphaGo and AlphaGo Zero breakthroughs occurred in nearly three years ago, but we found no evidence this underlying innovation that spurred the development of these algorithms is being explored. Even the discipline through which they were developed, Machine Learning, have barely penetrated into Navy R&D efforts, despite having gained a resurgence 8-10 years ago due to hardware advancements such as GPUs.

99 There are three types of Aluminum-based power systems in development. **Type #1: Aluminum as a Heat Source.** This earliest type of system uses the Al-H₂O reaction to generate heat, which in turn is used to drive a heat engine (e.g. a Stirling cycle engine) and/or thermoelectric generators. **Type #2: Aluminum as a Hydrogen Source.** The next type of system uses the Al-H₂O reaction as a hydrogen gas source, combined with a complementary system for oxygen storage. The two gases are used to drive a proton exchange membrane (PEM) fuel cell. **Type #3: Direct Electricity Generation.** This most recent and most advanced type of aluminum power system generates electricity directly from the Al-H₂O reaction. This is done by splitting the reaction into two half-reactions to create a fuel cell. "AL-H₂O Aluminum Water Energy Models," L3Harris, available at <https://www.l3harris.com/all-capabilities/al-h2o-aluminum-water-energy-modules>; "Tech Notes: Aluminum as a Fuel," Lincoln Laboratory, Massachusetts Institute of Technology, available at <https://archive.ll.mit.edu/publications/technotes/Aluminum-as-fuel.html>

An APO can significantly increase the pace of experimentation and advance.

The unity of effort and prototyping mandate that the APO can bring could produce tangible benefits to increasing the pace of experimentation. The assortment of bureaucratic and administrative barriers discussed previously, which hamper the innovation of the ecosystem's science and engineers, caused a corollary effect on the boundary-pushing innovation that produces technological breakthroughs. This risk aversion is intellectually stifling to the ecosystem. With an eye toward prudent safety and ethics, the APO can help pare back unnecessary or onerous policies/regulations to allow experimentation across all LOEs to progress more openly. Furthermore, just like formalizing collaboration across the UxS R&D ecosystem will enhance technological advancement, the APO can also accelerate innovation and advancement through the formalization of competitive experimentation. Within academia and the commercial sector, technological competitions of all varieties have rapidly accelerated the state of the art by unbridling the innovation and intelligence of the best and brightest scientists, engineers, students, and practitioners. Examples include DARPA Driverless Vehicle Grand Challenges, DARPA Robotics Challenge, DARPA Grand Cyber Challenge, Hackathons, Advanced Naval Technology Exercises (ANTXs), NavalX Tech Bridges, and most recently DARPA AlphaDogfight. In particular, the DARPA Grand Challenge series directly spurred the driverless vehicle innovation that has since exploded into a sizable, advanced technology industry occupied by such corporate giants as Google's Waymo, Tesla, Uber, and many others.¹⁰⁰ Lastly, these events also would provide a venue for the APO to discover future talent who could be hired, nurtured, mentored, and educated as part of the critical talent development and management process needed to retain this intellectual expertise.

100 The DARPA Urban Challenge was held on November 3, 2007, at the former George AFB in Victorville, Calif. Building on the success of the 2004 and 2005 Grand Challenges, this event required teams to build an autonomous vehicle capable of driving in traffic, performing complex maneuvers such as merging, passing, parking and negotiating intersections. This event was truly groundbreaking as the first time autonomous vehicles have interacted with both manned and unmanned vehicle traffic in an urban environment. A total prize purse of \$3.5 million was offered for the three fastest and safest vehicles that traversed a 60-mile urban course in moving traffic in less than 6 hours. Six vehicles completed the course, the fastest at an average speed of approximately 13 miles per hour. This achievement marks a significant landmark in the development of autonomous vehicle technology and represents a major advancement toward achieving the Congressional goal that by 2015 one-third of the Armed Forces' operational ground combat vehicles be unmanned. Defense Advanced Research Projects Agency, *ACTUV 'Sea Hunter' Prototype Transitions to Office of Naval Research for Further Development*, January 30, 2018, available at <https://www.darpa.mil/news-events/2018-01-30a>; and Alex Davies, "Inside the Race That Jump-Started The Self-Driving Car," *Wired Magazine*, November 10, 2017.

An APO can deliver on effective Operationalization of UxS to the Fleet through Unity of Effort in DOTMLPF development.

Earlier, this report explained how the disconnect between the UxS R&D ecosystem and the Fleet operators results in capability technological developments or prototypes that do not appeal or address the needs of the Fleet. This engenders frustration on all sides while wasting precious resources. To address this issue, the proposed APO organizational design contains a Fleet Liaison division that will provide coordination and collaboration access to the ecosystem during all phases of UxS research, development, prototyping, test & evaluation, and operationalization. Fleet Liaisons will most importantly be drawn from the four-star fleet commands, Naval Warfighting Development Center, and the warfighting development centers to ensure that their respective operational and tactical requirements, including Concept of Operations (CONOPS), are clearly articulated to ecosystem organizations. However, Fleet Liaisons from SYSCOMs, TYCOMs, OPNAV N9 resource sponsors, and other Navy Staff are also important to ensure those staffs' insights, concerns, and recommendations are incorporated into ecosystem R&D considerations. Lastly, the four-star fleet commanders should have DIRLAUTH with the APO Director to solidify the importance of direct feedback and communications among these leaders.

The APO collaboration with the Fleet can produce substantive benefits to the unity of effort/focus necessary for the successful operationalization of UxS. As discussed previously, operationalization includes the requirements to develop CONOPS and personnel training. APO and Fleet coordination can ensure the full Joint Capabilities Integration & Development System (JCIDS) Doctrine, Organization Training, Materiel, Leadership, Policy, and Education, Personnel, Facilities and Policy (DOTMLPF-P) is considered during every phase of the UxS ecosystems R&D, prototyping and T&E of future autonomous vehicle capabilities.¹⁰¹ If the DOTMLPF-P development analysis is conducted in stride with the developmental process, the APO will be poised to avoid the creation of systems in a vacuum that have little Fleet operational value or may flounder while waiting for missing supporting infrastructure to be established. Following this process incorporates Fleet input at all stages, which will ultimately ease the transition and operationalization of UxS to the Fleet.

¹⁰¹ Chairman of the Joint Chiefs of Staff, *CJCS Instruction Number 3170.01: Joint Capabilities Integration and Development System (JCIDS)* (Arlington, VA: Joint Chiefs of Staff, January 23, 2015); and Patrick Wills, "Joint Capabilities Integration and Development System (JCIDS)," *Defense Acquisition University*, February 22, 2017.



Photo courtesy of the U.S. Naval Institute.

An APO can serve as a service component AI/Autonomy office to the DoD's Joint AI Center

Under the Navy's current, dispersed UxS and AI/autonomy R&D structure, the DoD Joint AI Center (JAIC) would encounter a difficult task to determine which organization or office to coordinate activity with.¹⁰² Coordination of this effort is a challenging prospect within the Navy itself. However, the APO would naturally create a bridge to allow facilitation between the JAIC and the Navy. This relationship could function similarly to the one that the Defense Security Cooperation Agency (DSCA, the DoD umbrella organization for security assistance activity) and the Navy International Programs Office (NIPA, DoN's security assistance organization) have shared for several decades where DSCA passes maritime-related foreign military assistance equipment and training purchase requests to NIPA for contracting. In a JAIC-APO relationship, AI and autonomy-related R&D projects with a naval or maritime theme that are approved by JAIC can be passed to APO for execution and oversight. This construct could also extend to DARPA and SCO autonomy-related R&D projects in which APO can serve as the executive agent with those organizations and interlocutor/interface with the Navy's UxS R&D ecosystem to assign project responsibility. For example, SCO & DARPA projects such as Ghost Fleet, Cross-Domain Maritime Surveillance and Targeting (CDMaST), and Positioning System for Deep Ocean Navigation (POSYDoN) could be transferred to the APO for oversight execution, continued R&D, prototyping, testing, and eventual transition.¹⁰³

102 Paul McLeary, "Pentagon, Intel Agencies Set Up New AI Joint Office," *Breaking Defense*, April 13, 2018.

103 Kris Osborn, "Navy Pursues Drone 'Ghost Fleet' Concept," *Defense Systems*, February 2, 2017; Lee Hudson, "Pentagon to host industry day for SCO Ghost Fleet project," *Inside Defense*, September 25, 2017; Aaron Kofford, "Cross-Domain Maritime Surveillance and Targeting (CDMaST)," *Defense Advanced Research Projects Agency*, June 24, 2018, available at <https://www.darpa.mil/program/cross-domain-maritime-surveillance-and-targeting>; and Defense Advanced Research Projects Agency, *Positioning System for Deep Ocean Navigation (POSYDoN)*, June 24, 2018, available at <https://www.darpa.mil/program/positioning-system-for-deep-ocean-navigation>

CHAPTER 7

Conclusion

The U.S. Navy has achieved some UxS successes. Despite these achievements, it should establish the APO to better deliver on UxS and AI-based technology and avoid further erosion in the technological advantages that underpin U.S. maritime superiority. Specifically, an APO focused on delivering UxS prototypes can reduce bureaucratic friction that is degrading the ability to rapidly invent, innovate, and prototype, while accelerating progress across the existing UxS R&D ecosystem.

The Navy's current UxS R&D construct creates internal organizational friction that results in hampered progress. It is within this challenging environment that the U.S. Navy is competing to deliver unmanned autonomous vehicles. The slim technological advantage the United States currently enjoys over peer/near-peer adversaries in AI and other key UxS supporting technologies cannot be easily maintained under the existing bureaucracy, resulting in a steady erosion in our race to develop advanced military UxS. Consequently, the resultant friction is degrading our prized ability to rapidly invent, innovate and prototype. These disadvantages make UxS R&D successes difficult to accomplish without applying herculean effort and tremendous senior leadership engagement.

With the surge of R&D in UxS and AI by peer competitors in the pursuit for military purposes, now is the time for the U.S. Navy to comprehensively reorganize its efforts to ensure its continued technical advantage and outpacing of potential adversaries. To achieve this, the Navy must embark on a major transformation of a scope similar to how it pursued historical revolutionary technological development successes like naval nuclear reactors, submarine-launched ballistic missiles, and the Aegis Weapon System program. In all three cases, the Navy created robust, cross-functional, interdisciplinary organizations consisting of personnel from the military, government civilian service, industry, and academia that were given a broad, strong mandate and authority to research, develop, prototype, and operationalize transformational strategic capabilities. This reorganization is best executed by establishing a dedicated Autonomy Project Office (APO) focused on the advancement and delivery of UxS operational prototypes for experimentation, testing, and ultimately

operationalization and acquisition into the Fleet. By establishing an APO, the Navy can build unity of direction and effort across a wide spectrum of organizations, staffs, and commands that play varying roles in this endeavor. The main mission of the APO should be to unify oversight, authority, and direction of all the R&D lines of effort across the UxS ecosystem to continuously spiral the latest technological advances into an assembly line of autonomous unmanned air, surface, and undersea vehicle prototypes that can be demonstrated for warfighting value. An APO with the delivery of UxS prototypes as its primary mandate, can significantly reduce bureaucratic friction, accelerate momentum and truly harness the tremendous innovation and talent of our engineers, scientists, and sailors, quickening the Navy's technological progress.

The APO does not need to duplicate the excellent innovation and existing R&D, science, engineering, systems integration, and T&E organizations that already exist within the UxS R&D ecosystem. Instead, the APO will exist to better facilitate orchestration, collaboration among them to produce the unity of effort/direction needed to more rapidly deliver UxS prototypes. It must consist of scientists, engineers, researchers, experts, and managers from all corners of the UxS R&D ecosystem including FFRDCs, UARCs, industry, academia, and Navy labs and warfare centers. Additionally, it must be established with direct links to the fleet and its uniformed operators to ensure tactical relevance is strongly connected to APO activity.

Commensurate with its accountability and responsibility, the APO must be invested with the authority to direct the ecosystem in all matters pertaining to UxS, autonomous and AI R&D, prototyping, and T&E to include funding assignment. The APO's leadership and staff must demonstrate a deft, light touch whereby they facilitate the appropriate amount of orchestration among the ecosystem to achieve increased alignment without simultaneously inducing strangulation of innovation. There are some definitive advantages to establishing an APO that will accrue to the Navy's efforts to prototype and operationalize unmanned autonomous vehicles.

The Navy should establish the APO to better deliver on UxS and AI-based technology and avoid further eroding the slim technological advantage upon which our maritime superiority depends. An APO with the delivery of UxS prototypes as its mandate can reduce friction, accelerate R&D momentum and fully leverage the talent of our UxS researchers within the ecosystem.


Appendix

TABLE 3: UNMANNED SYSTEMS R&D LINES OF EFFORT DEFINITIONS

APO Division Title	APO Division Role
Navigation Control	R&D of software & hardware programs / modules that can autonomously execute precision navigation utilizing sensors and onboard inertial systems.
Mission Autonomy	R&D of software & hardware programs / modules that can autonomously execute key mission actions/behaviors such as patrolling, anti-submarine warfare, tracking, etc. using all data and sensors available to the system. It is derived from the field of Artificial Intelligence.
Testing & Evaluation	R&D of the T&E Trust Building Operationalization Regime for UxS consisting of Verification, Validation, Accreditation, modeling & simulation analysis and comprehensive field assessment. This T&E will be distinctly different than the conventional systems process because it is driven by the Mission Autonomy systems, derived from artificial intelligence-infused algorithms, which will be embedded within UxS.
Propulsion/Energy Endurance	R&D of battery, fuel, renewable, chemical, ocean-based energy endurance technology and associated propulsion systems for UxS.
Communications	R&D of software & hardware electromagnetic and acoustic spectrum-based communication systems for UxS.
Payloads	R&D of equipment and systems designed for deployment from or operation onboard UxS that support mission objectives.
Kinetics	R&D of weapon systems and associated equipment for employment from UxS.
Sensors	R&D of software & hardware electromagnetic and acoustic spectrum-based sensor systems for UxS.
Human-Machine Teaming	R&D designed to optimize human-machine interface / interoperability for maximizing the operational utility and employment of UxS.

Source: Author

FIGURE 2: ENGAGEMENT ACTIVITIES WITH UXS R&D ECOSYSTEM



LIST OF ACRONYMS

AFB	Air Force Base
AGI	Artificial General Intelligence
AI	Artificial Intelligence
AIRTEVRON	Air Test & Evaluation Squadron
APO	Autonomous Project Office
ASMS	Advanced Surface Missile System
CNA	Center for Naval Analyses
CNO	Chief of Naval Operations
COI	Community of Interest
COMOPTEVFOR	Commander Test & Evaluation Force
CONOPS	Concept of Operations
CSBA	Center for Strategic and Budgetary Assessments
DARPA	Defense Advanced Research Projects Agency
DDG	guided missile destroyer
DoD	U.S. Department of Defense
DoDD	Department of Defense Directive
DoN	U.S. Department of the Navy
DOTMLPF	Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, and Facilities
DOTMLPF-P	Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy
FSC-USV	Future Surface Combatant Unmanned Surface Vehicle
FFRDC	Federally Funded Research and Development Center
FNC	Future Naval Capabilities
GAO	Government Accountability Office
HOTL	human-on-the-loop
JAIC	Joint Artificial Intelligence Center
JPL	Jet Propulsion Laboratory
IATT	Interim Authority to Test
LOE	line of effort
MDUSV	Medium-displacement unmanned surface vessel
ML	Machine Learning
NASA	National Aeronautics and Space Administration
NISE	Naval Innovation Science & Engineering
NRAC	Naval Research Advisory Committee

NSCAI	National Security Commission on Artificial Intelligence
ONR	Office of Naval Research
OTA	Other Transaction Authority
PEO	Program Executive Office
RFP	request for proposal
R&D	research and development
SLBM	Submarine-launched ballistic missile
SME	Subject-matter expert
SSBM	Nuclear-powered ballistic missile submarine
SUBDEVRON	Submarine Development Squadron
SURFDEVRON	Surface Development Squadron
S&T	science and technology
T&E	testing and evaluation
TRL	Technology readiness level
UARC	University-Affiliated Research Center
UAV	unmanned aerial vehicle
USN	U.S. Navy
USV	unmanned surface vehicle
UUVRON	unmanned undersea squadron
UxS	unmanned autonomous system
VV&A	Verification, validation, and accreditation
XLUUV	extra-large unmanned undersea vessel



Center for Strategic and Budgetary Assessments

1667 K Street, NW, Suite 900

Washington, DC 20006

Tel. 202.331.7990 • Fax 202.331.8019

www.csbaonline.org