NRAC

Naval Research Advisory Committee

Autonomous and Unmanned Systems in the Department of the Navy

September 2017

Table of Contents

Executive Summary

Autonomy, both mobile and "at rest," is poised to revolutionize warfare as completely as steel, gunpowder, electricity, aviation, and computers did in prior generations. If we do not radically change the path we are on, America's adversaries will soon be able to defeat us in several arenas because of their rapidly increasing military capacity combined with their more aggressive fielding of autonomous capability. It follows that if the United States military does not accelerate the development and fielding of autonomous technology, the U.S. will cede military pre-eminence to those countries that do. Numerous reports conducted for the Department of the Navy, the Department of Defense, and our sister Services have endorsed the importance of autonomy and unmanned systems.

The commercial sector leads in the development of autonomous capabilities today. Advances in machine learning technologies, such as Deep Learning, underlie widely reported accomplishments such as the IBM Watson *Jeopardy* and AlphaGo victories over human champions. As the opportunities for financial gain or looming irrelevance become clear, industries as disparate as advertising, finance, automotive, cyber, medicine, aviation, and maritime transport are investing heavily in the development of autonomous and/or unmanned systems. Because industry is global, these investments are not exclusive to the United States.

A key finding of this study is that while the U.S. may currently have the advantage in autonomous systems, our adversaries are catching up. In some technical areas, our adversaries may already be ahead. The gap between defense and industry is growing – which in turn provides another vector for adversaries to leapfrog our autonomous capability. They might, for example, simply buy small U.S. entrepreneurial companies where much of the cutting edge technology is being developed.

A new element of the most advanced unmanned systems and autonomy at rest is software that "learns." In learning, autonomy software is transformed by the data used to train it. The code in many new autonomous systems is continuously rewritten by the data to which it is exposed. Because data is central to the creation and operation of autonomous systems, it is our greatest and most precious asset. Indeed, in industry, data ownership provides the key competitive advantage – think of Facebook. This leads to our first recommendation: **DoN must urgently develop an organizational data plan**. This plan must articulate the value of data to the DoN enterprise and should include processes to maximize that value and assign responsibilities. The key goal of the plan should be to enable the more rapid creation, evolution, and testing of learning systems in all areas of warfare from intelligence to operations. The plan must foster speed without relaxing cyber security.

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. Accordingly, we recommend that **DoN create a world-**

class VV&A research program for autonomous systems by dramatically expanding the work being done by the DoD Autonomy Community of Interest. We note that cyber security, already a grave concern for DoN and all of DoD is an even more urgent concern for autonomous systems and must be integrated at every level in the VV&A effort. The threat is that an adversary might turn our own systems against us.

In our visits to Warfare Centers and Laboratories, we discovered an institutional tendency towards "one size fits all" risk mitigation procedures for obtaining operating clearances which seem to be hindering larger goals. Risk mitigation for a system that costs tens of millions of dollars clearly should be approached differently than it should for a system that costs only a fraction of that. Indeed, relaxed requirements and restrictions on lower cost systems provide a pathway to accelerate innovation and experimentation of new autonomous technologies. This should be exploited. Thus, our third recommendation is that **policies should enable rapid test and evaluation of autonomous systems, and restrictions on their early deployment and use should be commensurate with risk thus entailed**.

While advances in computer and networking technologies have been driven by industry for decades, the DoN procurement system remains structured around Cold War assumptions of the pre-eminence of government-funded technology. Lengthy approval processes that at times reach all the way back to Congress, effectively give more nimble adversaries and organizations a head-start measured in years. This dysfunctional system is eroding our technological edge and puts us at risk of losing a major military engagement in the future Below we offer a series of recommendations designed to make use of powerful new procurement tools, to greater engage Navy and Marine Corps leadership in fielding autonomy capabilities, and to create an innovation ecosystem borrowing lessons learned from industry and the intelligence community. Specifically, we recommend:

- **Use Other Transactional Authority (OTAs) as forcefully directed by Congress, whenever possible.**
- **Demand plans from type commanders to leverage existing autonomy to generate new capability.**
- **Create a "Shark Tank" approach to internal Naval innovation funding within the DoN, and fund with an initial $50M pilot.**
- **Create Naval-oriented incubators in the Silicon Valley and the Boston areas.**
- **Create a DoN-owned venture firm with $30M patterned on In-Q-Tel.**

Autonomy is a fast moving domain, arguably still in its infancy. A heated battle for talent is already well under way, with companies like Uber hiring 40 university researchers from a single center, or with Toyota Research Institute announcing a one billion-dollar investment in autonomy. While DoD and DoN have nurtured many of the research activities that now underlie the explosion in commercial investment in autonomy, the government is currently losing the battle to hire the best minds. Consequently, **we recommend that DoN remove barriers to recruit and retain the best talent.** Today's hiring processes are slow, lack flexibility in minimum formal education requirements; and new hires are compensated poorly compared to industry. This jeopardizes hiring the best candidates. Further, incorporation of autonomy in education and training across the DoN must be expanded. Autonomy is not a stovepipe domain: it has applicability to all aspects of the Naval force.

We emphasize that elements of the DoD enterprise are mobilizing to these challenges. The Third Offset Strategy explicitly leverages advances in learning systems. The new Digital Warfare Office and the emerging Modeling and Simulation Enterprise are important steps towards organizing around the new imperatives.

However, these efforts are tiny compared to the opportunity and the threat. Given the breadth and importance of autonomy to the future of the Department of the Navy, **we recommend that the Secretary of the Navy be the Champion for Autonomy** in order to ensure funding and urgency are provided to enable the Navy and Marine Corps to stay ahead of all potential adversaries.

The NRAC authors of this report are:

Chairman, Dr. James Bellingham, Director for the Center for Marine Robotics at the Woods Hole Oceanographic Institute (WHOI),

Vice-Chair, VADM William Bowes USN (Ret.), former Commander Naval Air Systems Command,

Dr. Ann Gates, Department of Computer Science at The University of Texas at El Paso,

Dr. Ingrid Y. Padilla, Professor in the Department of Civil Engineering and Surveying at the University of Puerto Rico, Mayagüez,

Dr. Jay Walsh, Vice President for Research at Northwestern University,

RADM Charles Young USN (Ret.), former Director, Strategic Systems Programs.

Authors/Subject Matter Experts are:

VADM Michael Connor USN (Ret.), Founder and Chief Executive Officer, ThayerMahan Inc.

Dr. Kevin Fall, former Chief Technology Officer and Deputy Director of the Software Engineering Institute (SEI) at Carnegie Mellon University,

Mr. Scott O'Neil, former Executive Director and Director for Research and Engineering at the Naval Air Warfare Center Weapons Division (NAWCWD) China Lake,

Mr. George Nolfi, Communications Specialist and Filmmaker.

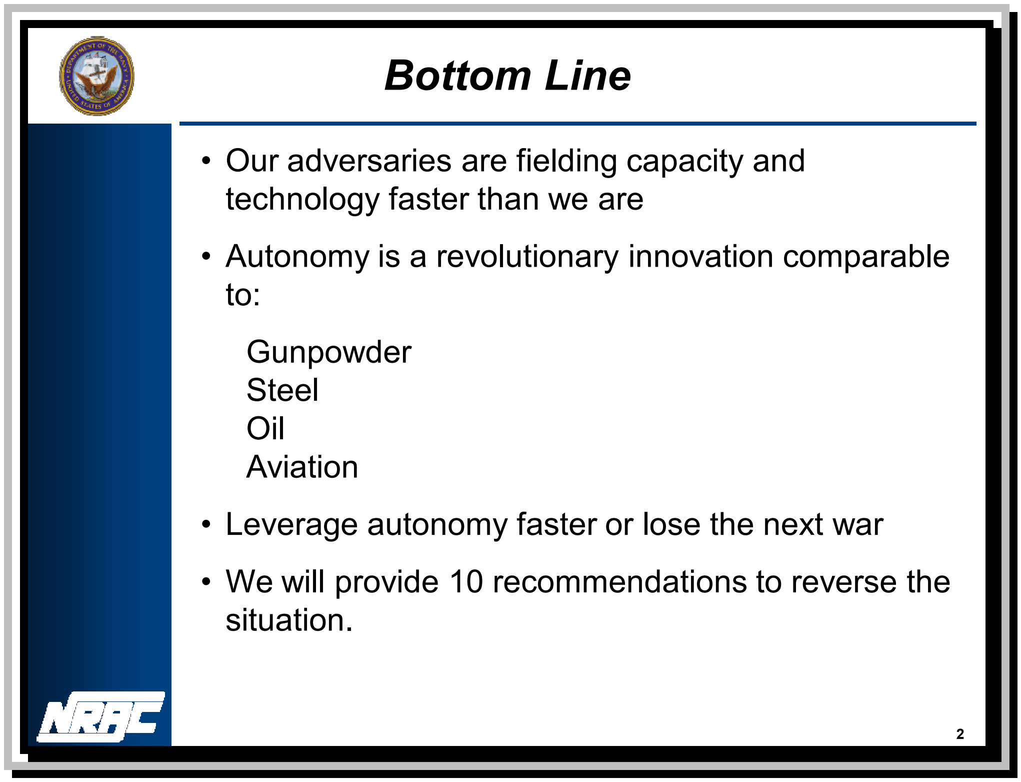

Autonomous technology – a basket of technologies ranging from automation to AI (Artificial Intelligence) in both robotic and "at rest" systems – is poised to revolutionize warfare as completely as steel, gunpowder, electricity, aviation, and computers have before it. It follows that if the United States military does not field "better" autonomous technology (and counter-autonomy) and do so more quickly than its adversaries, the U.S. is in danger of losing a major military conflict in the not-too-distant future.

While this is truly unthinkable to many, our year-long examination of how the Navy and Marine Corps are approaching autonomy suggests that if we do not radically change the path we are on, America's adversaries may soon be able to defeat us in several potential arenas because of their rapidly increasing military capacity combined with their more aggressive fielding of autonomous capability.

The slow pace of innovation in the Department of the Navy relative to the pace of innovation in the private sector, including the private sector in other countries, shows the magnitude of the leadership challenge we face in this area. Potential adversaries, such as China and Iran, are aggressively advancing their civilian autonomy sectors. China already dominates the international drone market, for example, and intends to dominate the artificial intelligence market by 2030.

The Naval Research Advisory Committee was tasked by the Secretary of the Navy in July 2016 to examine unmanned system strategies in the Department of the Navy with a goal of providing recommendations on the application of autonomous and robotic systems, levels of autonomy, learning machines and human-machine teaming.

Autonomy and unmanned systems are rapidly changing the way the world functions – both in civilian and military domains. Autonomy will revolutionize warfighting by drastically affecting the speed, accuracy and persistence/longevity with which sailors and marines detect and evaluate conditions in the battlespace, disseminate information, and execute operations.

The Navy and Marine Corps (and the other Services) have been slow to adapt to the new reality as compared to our adversaries. The Naval research and development sector is not incentivized to operate at the speed of the private sector. It is not sufficiently connected to the new tech economy – especially at the start-up level. Instead, the Naval R&D establishment attempts to recreate the capabilities of "Silicon Valley" through Naval laboratories/warfare centers, systems commands, and a few large defense contractors. This process is administered through a set of slow and outdated budgeting and contracting procedures. It should not be surprising that this process consistently fails to produce innovative autonomy at the pace of the private sector.

The Department must learn to rapidly exploit civilian and adversary military technology advances for defense purposes. This could be described as a "buy, don't make" approach. Failure to make this fundamental change in the way we access technology may well lead to strategic failure should the U.S. engage rapidly-evolving powers such as China in the future.

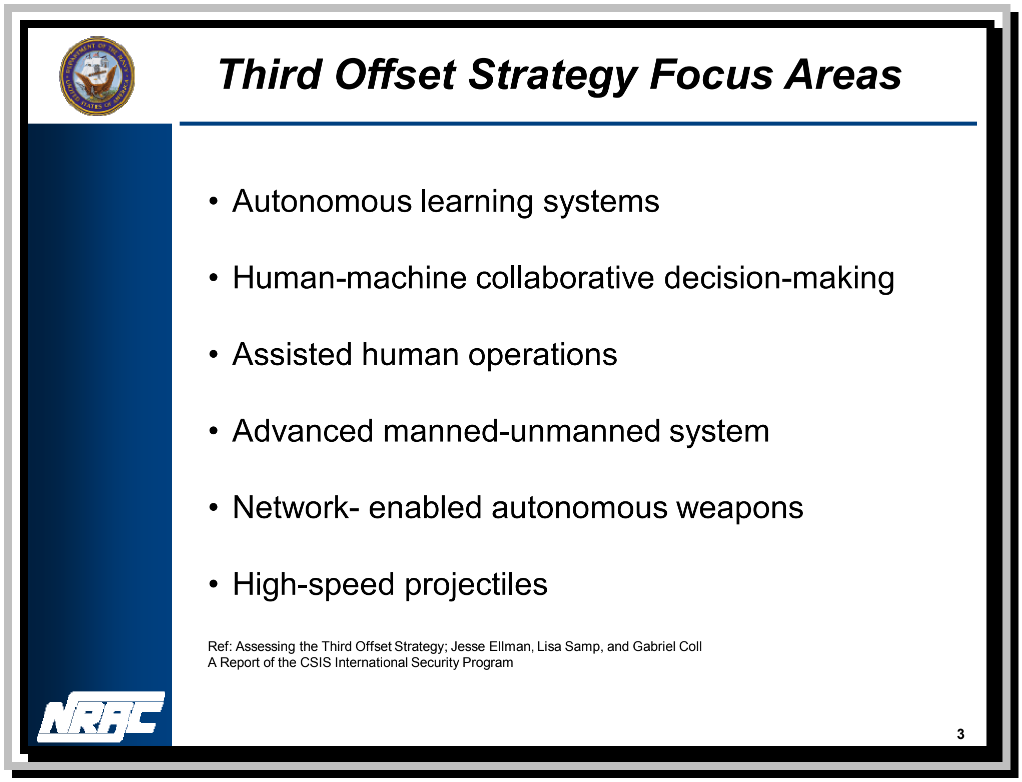

The so-called *Third Offset Strategy* encompasses autonomous learning, collaborative decisions by human-machine teams, advanced manned and unmanned systems, autonomous network-enabled weapons, and other areas.

The potential risk in the near future to US military dominance in certain arenas is recognized by the Department of Defense and directly led to the generation of the Third Offset Strategy.

The strategy formally recognizes the speed with which potential adversaries have adopted commercial technology in military systems and provides guidance to the services on focus areas for future developments to mitigate the erosion in our superiority. Many of those focus areas overlap with the autonomy topics in this study.

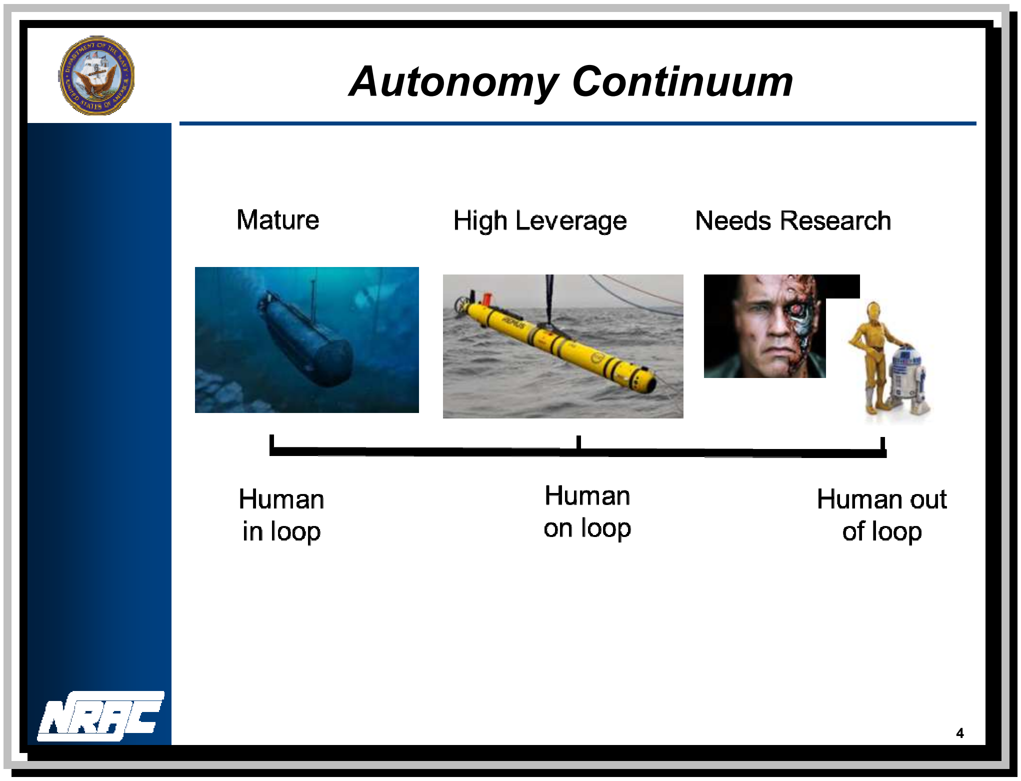
Autonomy Continuum
Mature
High Leverage
Needs Research
Human
in loop
Human
on loop
Human out
of loop
4

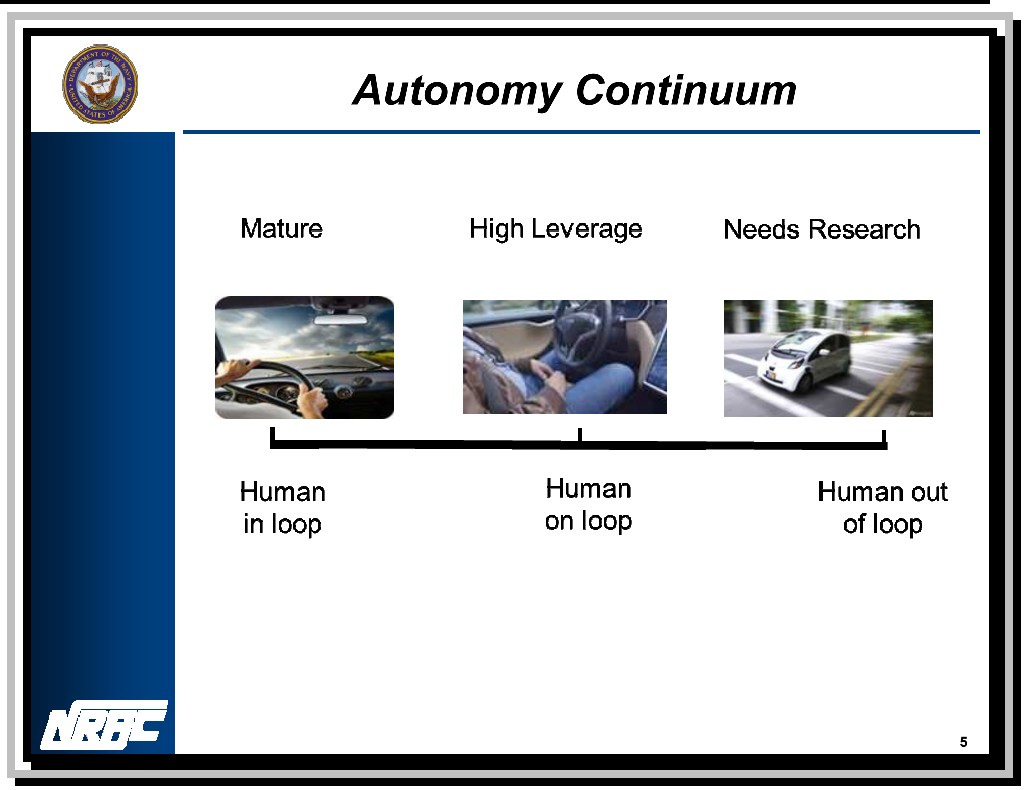
Autonomy Continuum
Mature
High Leverage
Needs Research
Human
in loop
Human
on loop
Human out
of loop
5

Simple autonomous devices are not new. The current "on-the-loop" (i.e.; humans fully present) systems are making way for human "out-of-the-loop" systems that will allow smart learning systems to make their own decisions without human intervention.

Autonomy is not new – we have had systems that operate autonomously for generations. Generally, legacy systems were limited in the range of decisions that they made and actions they could take. The first known industrial robot (circa 1937) could pick up materials and stack them in a pattern. The human was fully "in-the-loop".

Humans remain in-the-loop with many of today's "mature" autonomous systems that are capable of far more complex decisions and actions. Examples of mature systems, which are commonplace and trusted, include an automobile's cruise control and newer collision avoidance systems. The driver is in-the-loop for both systems (i.e., the driver can take over the car operation at any point), but cruise control can independently change the fuel flow for the vehicle to maintain its speed. Collision avoidance systems apply the brakes independent of the actions of the driver. This is true even if the driver falls asleep and maintains pressure on the gas pedal.

One area where there is a significant opportunity to rapidly enhance our autonomous capability is with systems where the human might be described as "on-the-loop". Here the human is more of an observer as the autonomous system makes complex decisions, but he or she can quickly intervene if need be.

We are seeing this on-the-loop application in the commercial world. Tesla automobiles use software that allows for mostly autonomous operation. The car has the appropriate sensors to observe its environment, make decisions, and act. That said, the driver is periodically instructed to take control of the car. The human needs to be at the ready (i.e., on-the-loop). Data indicate that the self-driving ability of cars leads to safer operation than with human-only driven cars. Trust in these on-the-loop systems is growing as each new system is operated, iteratively tested, and improved.

Other examples of autonomous activity include vacuuming our homes with autonomous robots; credit card companies employing autonomy at rest to monitor transactions instantaneously (e.g., looking for anomalous activity and automatically stopping subsequent credit transactions); and pension funds that employ autonomous systems to make investments.

As indicated on the far right of the above charts, systems will eventually place the human "out-of-the-loop" (i.e., the system makes decisions on its own, learns from external stimuli, and develops new processes and actions on its own). We are starting to see very early evidence of human out-of-the-loop systems. It is expected that there will be self-driving taxis – with software developed by MIT spin-off *nuTonomy* – operating on the streets of downtown Singapore in the near future.

We can imagine, but not yet engineer, even more sophisticated out-of-the-loop systems that learn and interact with humans and exhibit human-like behavior (what science fiction has imagined as anything from *C3PO* to the *Terminator*). Considerable research is required to move toward autonomous systems with these capabilities.

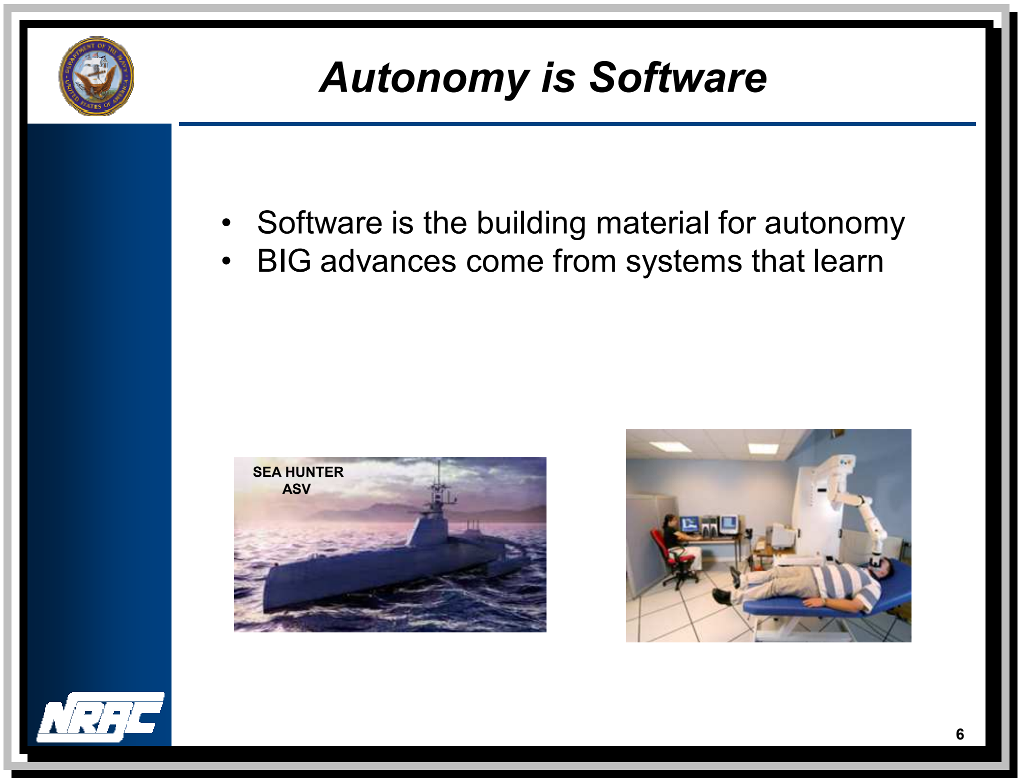

Software is the "brain" or key component of autonomous systems.

Autonomous systems of any significant complexity are built of software. Software is the key building material both for autonomy at rest – where it may operate on conventional computing devices such as laptops or servers – but also for autonomy in motion. Autonomy in motion is closely associated with robotics, where the software is responsible for the higher-order behaviors that humans observe as a robot accomplishes its tasking. Software is the "brain" of autonomous systems.

Whether autonomy software is used in motion or at rest, it acts in terms familiar to us for human behavior – in particular, the OODA loop (observe, orient, decide, and act). Even simple software is effective at quickly taking into account large data sets (observations), computing basic results, and performing actions. However, as the demands upon autonomous operations increase, software necessarily becomes more advanced – and more complex.

Recently, the sub-field of computer science known as Machine Learning has demonstrated remarkable advances in several key and historically difficult problems relevant to autonomy: computer vision, pattern recognition, strategy development (i.e., competitive games), language translation, planning, and "recommender" systems. Virtually all of these have direct applicability to autonomous systems with military application.

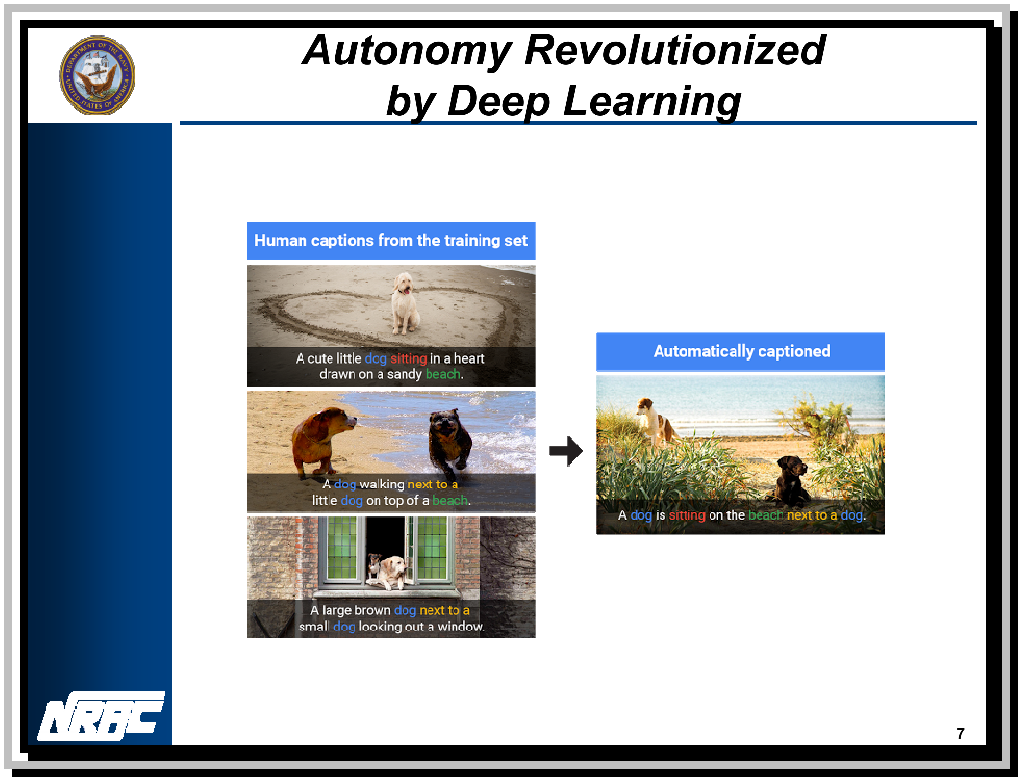

Deep learning, a class of machine learning, is the training of neural networks with data sets. Recent advances in computing power have allowed the practical use of much larger neural network systems.

At the time of the 2012 NRAC Autonomy report, it was accurate to say that a two-year-old child had better ability to understand visual scenes than the best software running on the fastest computers. That statement is no longer true. In the past five years, computer vision has made enormous leaps – with revolutionary implications for both unmanned and autonomous military capabilities.

The science of unmanned systems has been fundamentally limited by the inability of a system to characterize its surroundings with onboard sensors (e.g., as with UUVs). The more complex and unstructured the environment, the more difficult that process of creating situational awareness becomes. Consequently, the ability for computers to abstract key elements of information from complex sensor data is a key enabling technology for robotic systems to successfully operate in hostile environments.

Advances have been made on a number of fronts, but for the purposes of this brief we will focus on deep learning, a technology that is becoming familiar through the success of the IBM Watson computer. In 2012, Watson beat the top human contestants on the *Jeopardy* TV show. More recently, AlphaGo, a Google-developed computer program that plays the board game *Go*, beat the Chinese *Go* master to become the world champion of this complex strategy game. Deep learning has racked up some impressive demonstrations. But how does it work?

Deep learning refers to machine learning algorithms generated through the learning of multiple levels of representation/abstraction. It enables significant improvement in computer vision, object recognition, speech recognition, and language processing/understanding. This leads to the training of massive neural networks with data sets. For example, images that have been described by humans (i.e., annotated) were fed into a Google training algorithm. The resulting trained system was able to identify objects in a photograph, determine what the objects are "doing", and know where they are. The software did this despite encountering partial obscuration of the object and never having seen the exact image before. The trained system had an astounding accuracy of 94%.

It is worth noting that Google has released the code for this system on the web – literally anyone can download it. What they have not released is the data set with which they trained the software. The lessons here are four-fold: 1) computers are approaching human skill levels in the ability to interpret imagery, 2) what makes this work is the large annotated data set, 3) the results are available to everyone and 4) the "data" are the most valuable commodity.

In essence, deep learning enables "robots" to characterize their surroundings – and that is just the beginning…

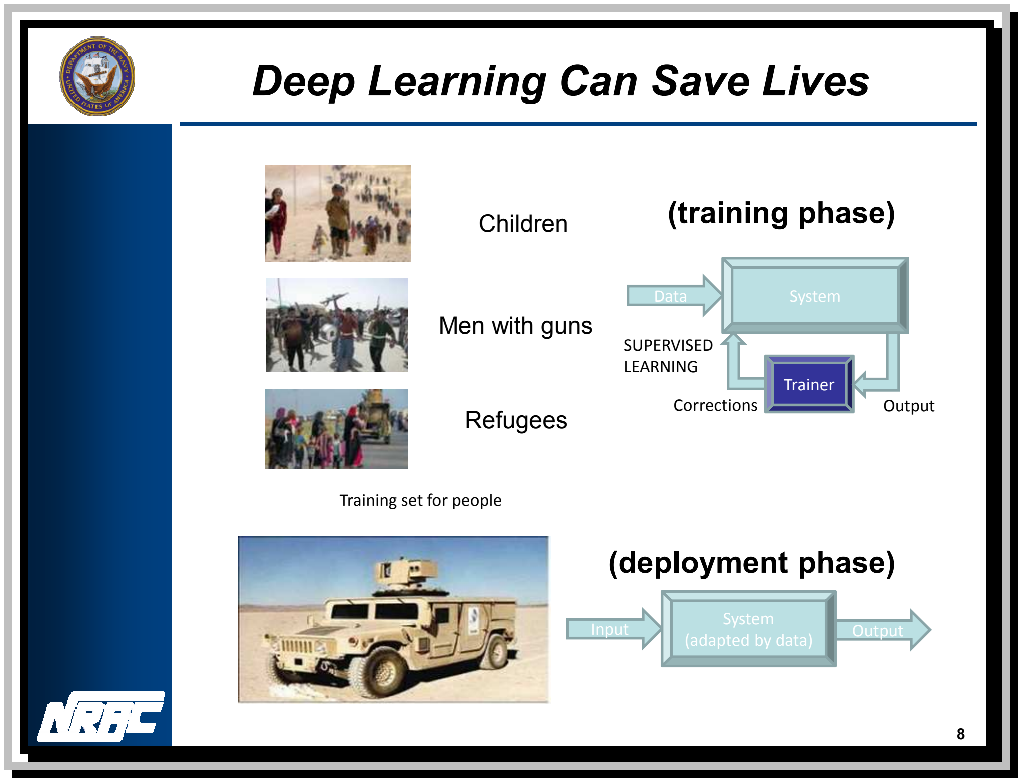

New autonomous systems with deep learning capability could significantly improve Marine Corps battlefield capability when constant vigilance is required for extended times.

A squad of Marines operating in a potentially hostile environment needs to be constantly aware of its surroundings. Marines in a Humvee will be looking for possible insurgents, IEDs, and multiple other threats. In most cases, the detection system of greatest value is the "mark-1 eyeball" connected to the human brain. However, humans get tired and constant vigilance is exhausting. When combined with fear, it can be debilitating over time. This is an area where deep learning can have an enormous positive impact.

Postulate a future with neural nets trained with multi-source battlefield intelligence and imagery in the Marine squad's area of operation. The resulting autonomous image classifier can be mounted with a 360-degree view on the top of the Humvee. We benefit from the digitization of sensors, and their exponentially dropping costs. The ultimate system would embody persistence, precision, and recall.

Our next example considers the opposite extreme, when information is plentiful, but time is limited.

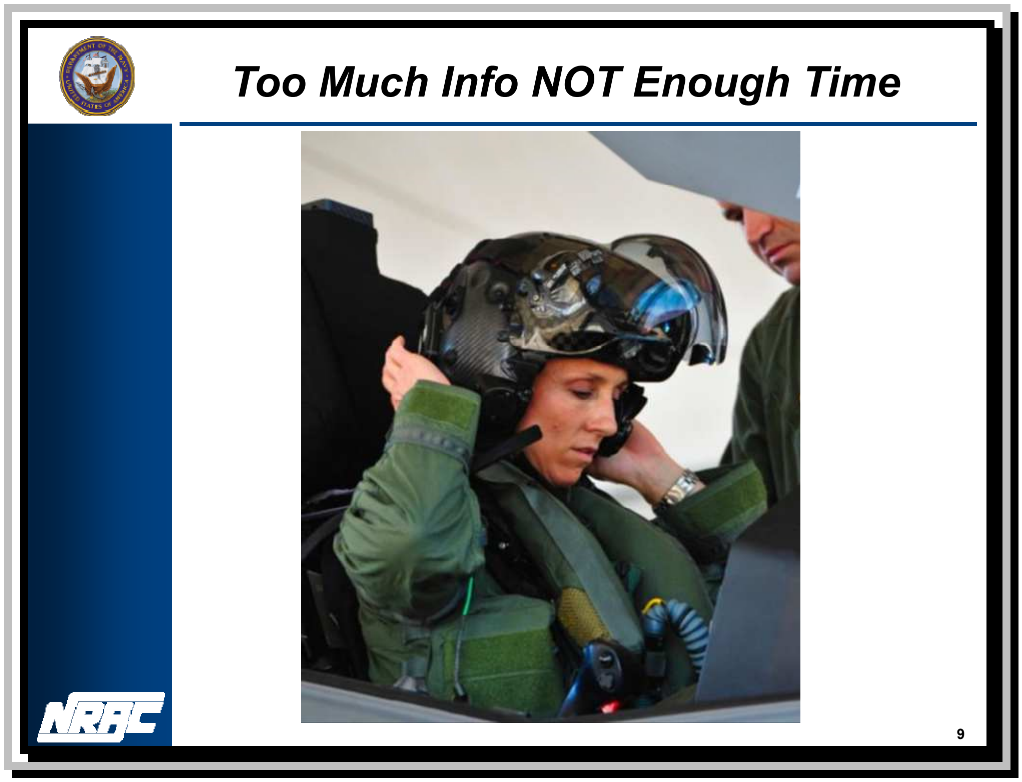

The air is a domain in which a fraction of a second may be the difference in survival. Adding AI to manned aircraft with the ability to take flight control – operating autonomously when the threat warrants – would provide a significant improvement in combat outcomes.

Imagine an F-35 pilot strapping into her aircraft for a mission to attack a defended target. As she approaches the enemy coast line, her electronic warfare (EW) warning system alerts her to five surface-to-air missiles heading towards her position. Operating at machine speed, the AI embedded into the F-35's integrated weapons system analyzes the flight path of each missile, determines which have an actual lock-on and are a threat to her jet, and then determines the best actions for the F-35 to take: in electronic jamming, chaff dispensing, and evasive maneuvers to avoid being hit. The F-35's AI completes the OODA loop almost instantaneously, takes the optimal electronic warfare system actions, and temporarily takes flight control from the pilot, allowing her F-35 to successfully evade all the missiles. Flight control is then returned to the pilot to continue the mission and strike the assigned target.

A current example of an AI application is the Air Force installed Auto G-CAS (ground collision avoidance system) in their Block 40 F-16s. This is an AI capability linked to the flight control system. In a direct parallel to the hypothetical above, when the F-16's system senses when "controlled flight into terrain" i.e., CFIT, might occur without immediate action, the AI-enhanced flight control system takes control of the airplane and executes as violent a maneuver as necessary within structural limits to avoid ground collision. It then returns control of the airplane to the pilot. In the first year after Auto G-CAS was installed (Nov 2015) the USAF documented four saves from Auto G-CAS. The Air Force has lost 6 to 7% of every fighter type from CFIT, and today these mishaps account for over 75% of fatalities. (Auto G-CAS is not incorporated in F-15, F-18, F22 or F-35).

Airbus commercial airliners have similar capabilities installed to prevent mid-air collisions. Because AI can complete the OODA loop faster than is humanly possible, AI-enabled systems can be used to counter threats previously too overwhelming for a human to deal with, prevent CFIT, avoid mid-air collisions, and perform countless other enhancements.

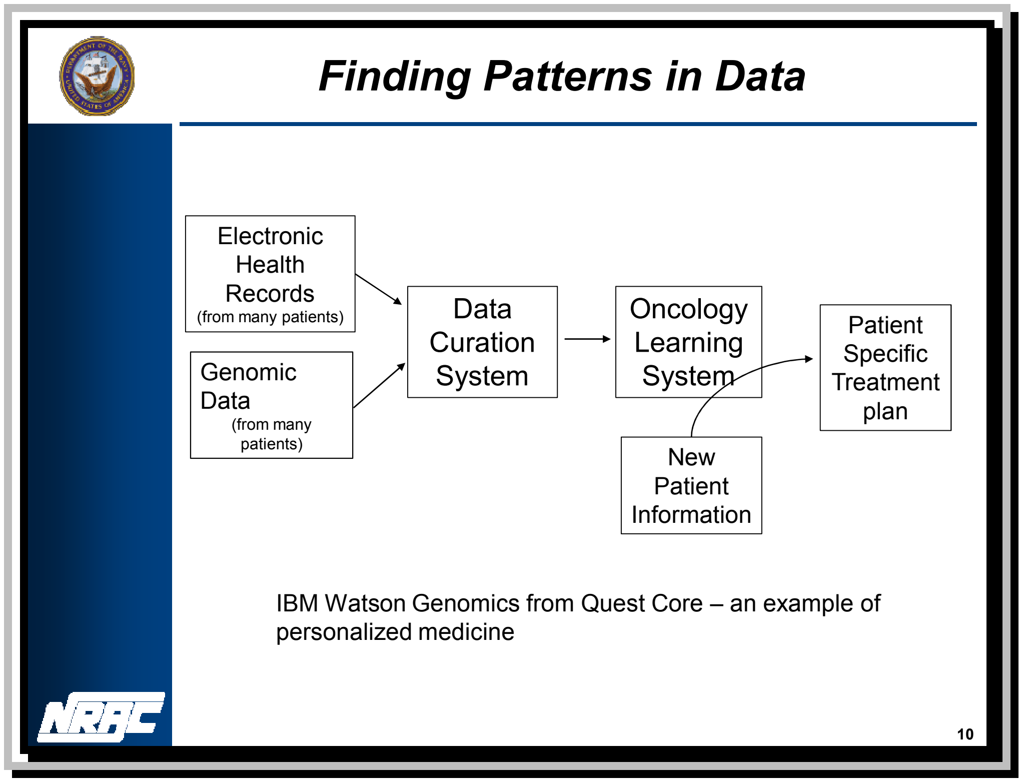

Sometimes there are patterns that even highly trained humans cannot detect without computer assistance. Computer-aided patient treatment is an example of data manipulation by AI learning systems.

Over the past several years, medicine has accelerated its move toward the long-sought goal of computer-aided diagnosis and treatment. For generations, medical data (signs, symptoms, lab data, treatment plans, outcomes, etc.) have been tied up in paper records. In the past two decades electronic health records (EHRs) have changed how patient information is collected, managed, stored, and made available to health care practitioners. The ease with which medical data can be delivered into the EHRs has greatly increased the data that are stored in ways that are readily processed by computers.

Traditional data in electronic health records are now being combined with genomic data and fed into a machine learning system. That system can now look at the data from a new patient and determine from the patient's profile what the optimal treatment will be. The chart above shows a curated data system as a key component. 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. A key element of the curated data set is the outcomes that resulted from a treatment regime. Thus, if a new patient comes to the clinic his or her information is transmitted to the "learning system" which then uses the patterns from the past to suggest a treatment plan.

The improvements realized are dramatic. In one program, when heart failure data were analyzed, the improvements led to an 83% decrease in mortality 24-months following the implementation of guideline-directed therapies.

For medicine, the goal is to win the war against disease. An important path to that goal is the generation of data, the storage of those data, the curation and thus ready access to those data, and the use of those data in the development of artificial intelligence-based systems.

These lessons are relevant to the Department of the Navy. Data is the ultimate resource for learning algorithms.

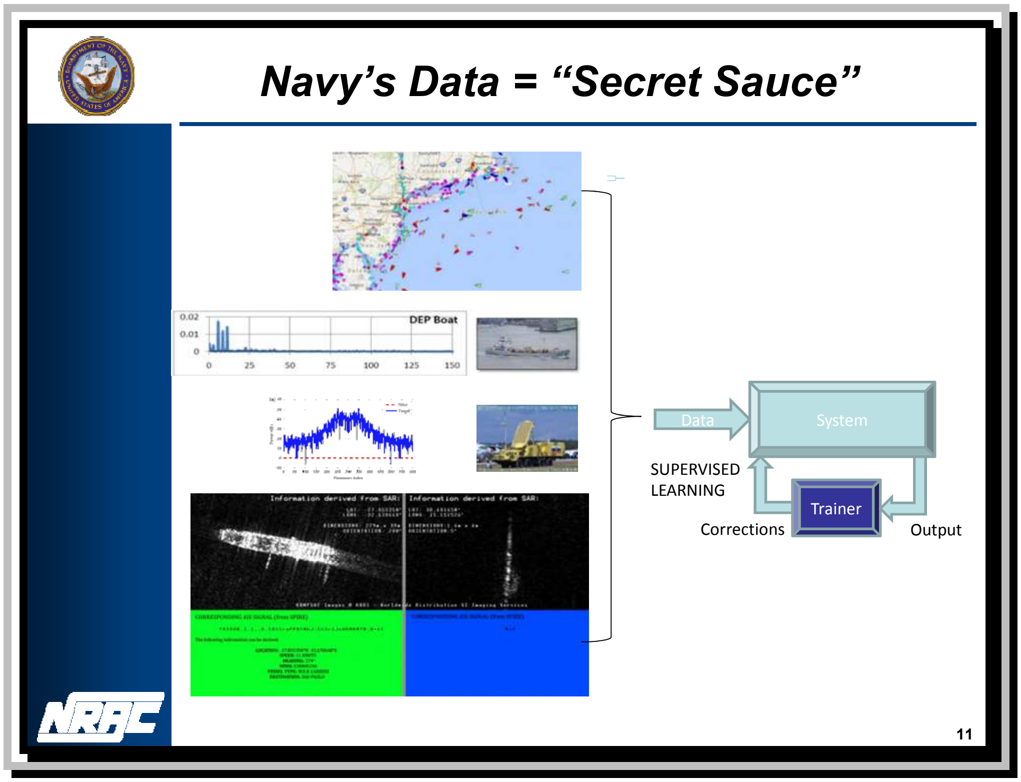

In industry, it is access to data rather than access to algorithms that increasingly provides the competitive edge in the machine learning field. Google's Tensor Flow and Amazon Machine Learning are both open source frameworks made available by industry giants. However, to train the machine learning 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.**

All digital data and related software can be vulnerable to theft and manipulation, meaning that cyber tools are required to detect and prevent hacking threats. Ultimately, these tools must approach 100% reliability if we want to justify the use of AI-based systems. Significant resources will need to be allocated for cyber security.

The Navy and Marine Corps must create plans for capturing, assembling, describing, protecting, and making use of its data. Sources such as social media, consumer activity, and traffic systems may provide key data points when attempting to infer an adversary's intentions.

While the military has many sources of data that are unique and usually secure, it must not ignore civilian data sources that might become unavailable or corrupted by hackers.

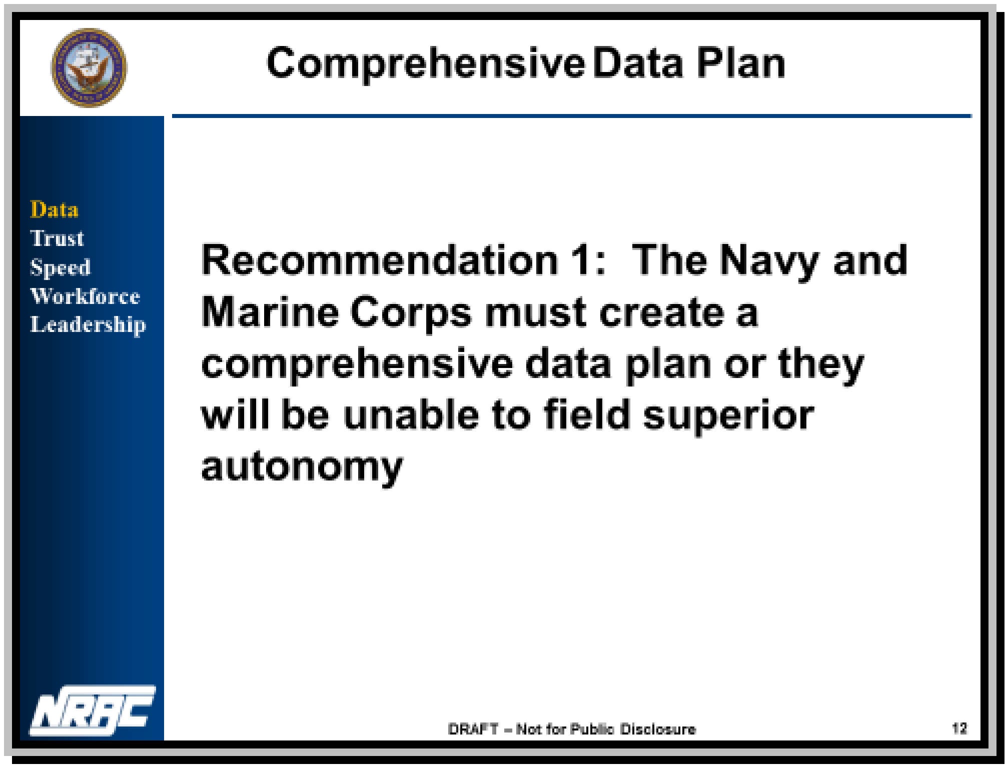

Data is the ultimate "component" for AI systems and must be controlled. **Recommendation 1**: The Navy and Marine Corps must create a comprehensive data plan or they will not be able to field superior autonomy.

Warfare is increasingly a data intensive enterprise – and will become exponentially more so – as autonomous systems proliferate. The uses of data are legion. It is often the "product" in the sense that an autonomous system might provide key information such as tracking or targeting information to a human operator or other autonomous system. It is also a record of "what happened", which provides the ability to learn from past experiences. This is particularly important in the case of machine learning systems, which may be trained from curated data sets constructed from actual field activities. Modeling and simulation systems will be validated in part against such field data. Prediction is at the core of warfare, whether it is of the operational environment, or action of an adversary. In either case, data is central to the trustworthiness of the prediction. As we develop machine learning systems capable of working with larger and more diverse sets of data, machine prediction may provide critical aid to commanders. In short, just as the commercial world is being transformed by data, the sources and applications of data will proliferate in ways that profoundly impact how war is fought.

We recommend that the Department develop a data plan. The plan must have a clear articulation of the value of data to the Naval enterprise – in all its forms. It should include an assessment of the capabilities data should provide, but also the vulnerabilities they create. The plan should frame interactions with the research community and industry, and provide transition paths to enable fleet capability. Examples of diverse issues that should be addressed include: what is the appropriate strategy to best engage creativity in the research communities and in the commercial world? How should data rights be handled in procurement? What institutional mechanisms should be required to ensure that the natural inclination to hold data tightly is balanced by a rational but aggressive push to gain data supremacy? How should the DoN leverage investments within the larger DoD? The Department's understanding of the role of data will evolve over time, consequently the data plan will need to be treated as a living document.

In the subsequent slides, we discuss other elements essential for the fielding of advanced autonomous systems.

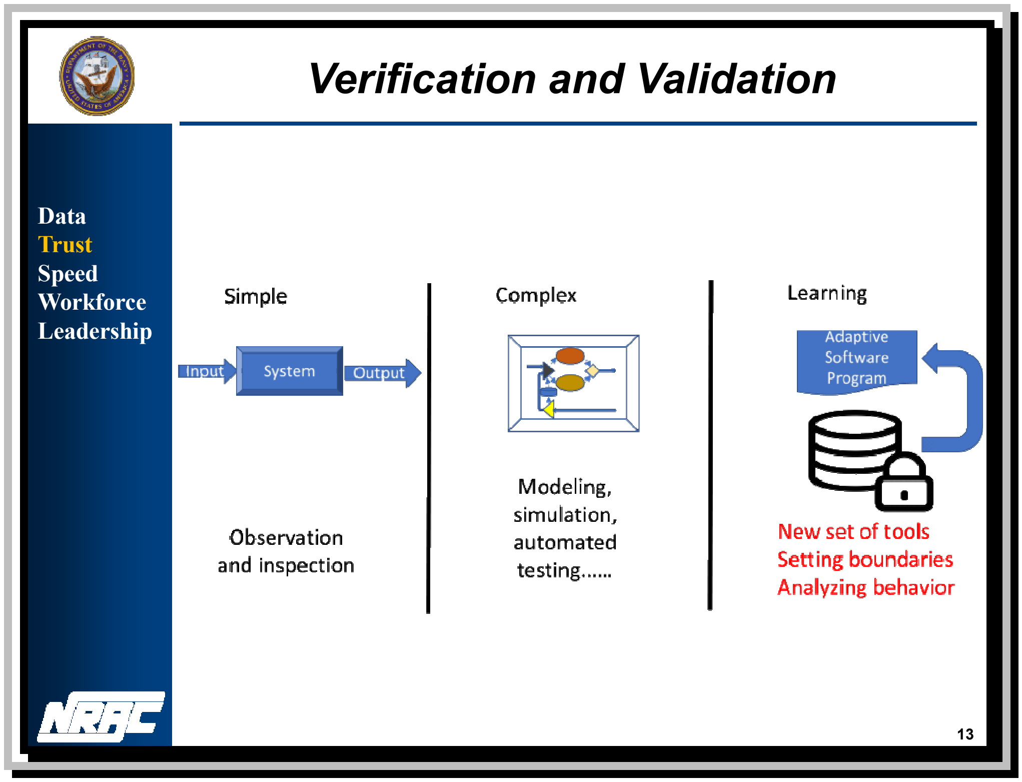

Verification and Validation
Data
Trust
Speed
Workforce
Leadership
Simple
Input
System
Output
Observation and inspection
Complex
Modeling, simulation, automated testing……
Learning
Adaptive Software Program
New set of tools
Setting boundaries
Analyzing behavior
13

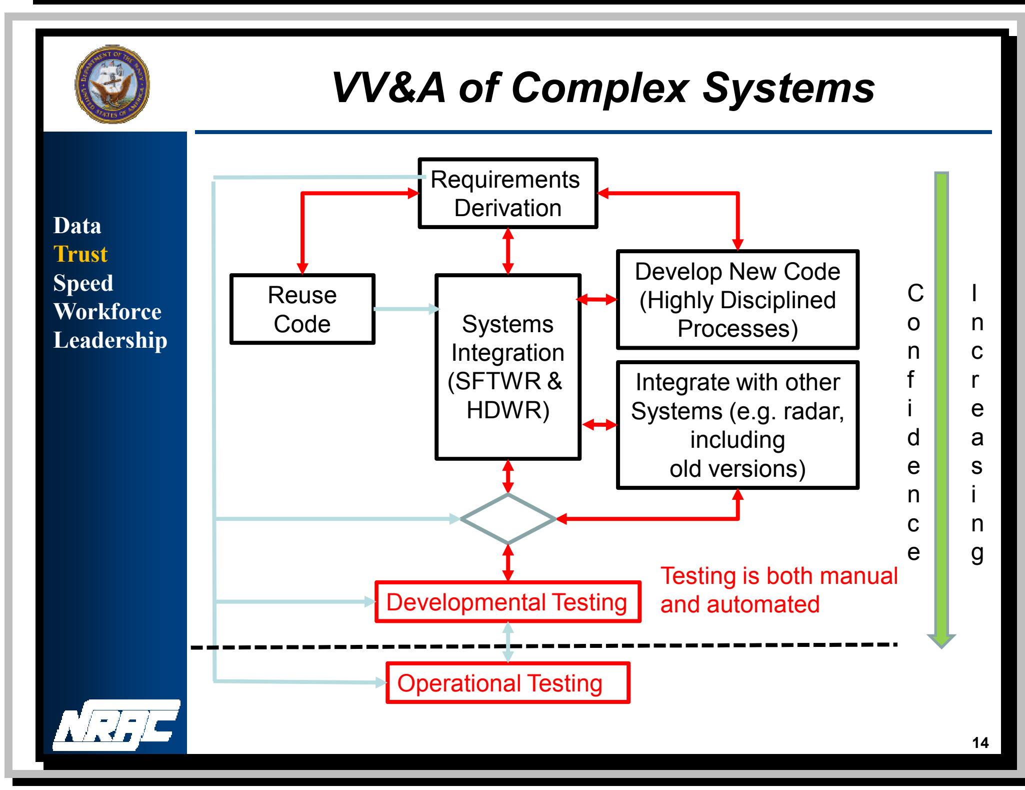

VV&A of Complex Systems
Data
Trust
Speed
Workforce
Leadership
Requirements Derivation
Reuse Code
Systems Integration (SFTWR & HDWR)
Develop New Code (Highly Disciplined Processes)
Integrate with other Systems (e.g. radar, including old versions)
Developmental Testing
Operational Testing
Testing is both manual and automated
Confidence
Increasing
14

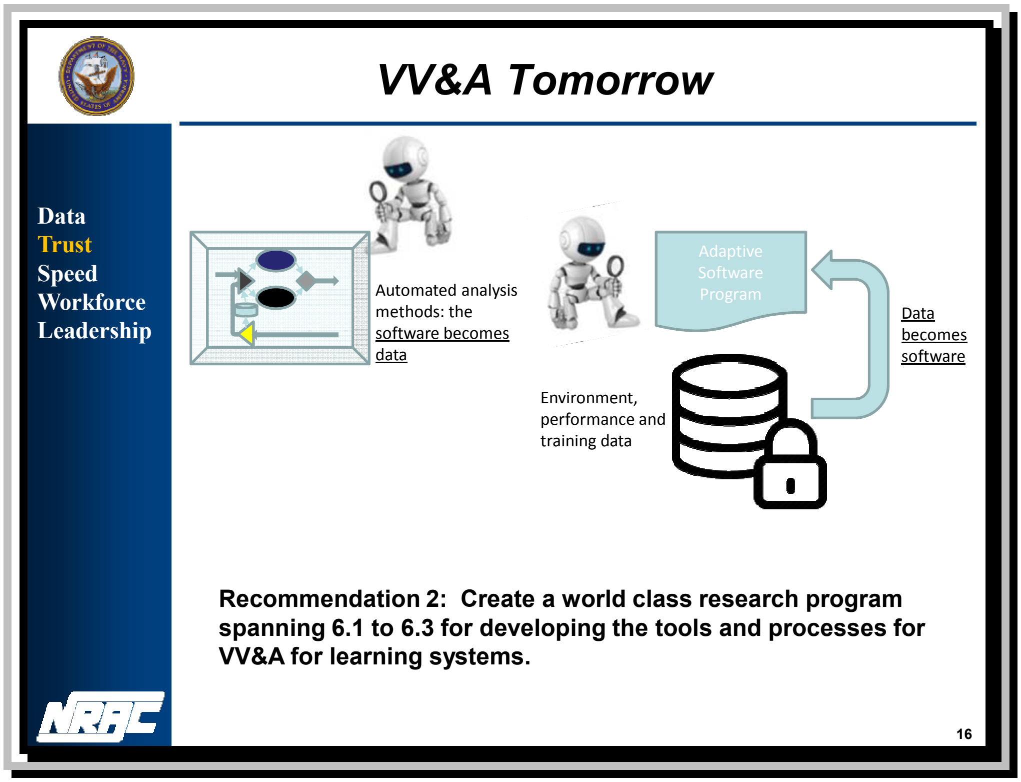

These three charts focus on verification, validation, and accreditation (VV&A). Verification refers to the correctness of the implemented system and its associated data, validation refers to the accuracy of the representation and intended use of the system, and accreditation is the official certification that the system and its associated data are acceptable for use. Traditional VV&A processes are not applicable for systems that incorporate learning algorithms – the basis for **Recommendation 2:** Create a world-class research program in 6.1 to 6.3 for the VV&A for autonomous learning systems.

In the future, where autonomous learning systems will supplant human operations in certain realms, VV&A becomes even more critical – as well as more difficult to achieve.

Early digital systems were simple and their performance could be readily proven by running simple tests, observing behavior, and inspecting the results. As digital systems grew in hardware and software complexity, and cybersecurity became a threat impossible to ignore, the methodology for the VV&A of system performance also evolved.

Software for military and other "high consequence" systems involves special risks. Software must offer a reasonable degree of trustworthiness, and it is of little value to invest significant analysis and testing resources in a new software module if it is executed on a software infrastructure that has serious known security vulnerabilities (e.g., Microsoft Windows XP).

During the latter part of the Cold War, learning to trust complex or highly integrated military systems of software and hardware (e.g., systems for strategic deterrence) demanded a new, composite approach. This approach touched the design process and required highly structured software development processes and reuse of proven software for specific recurring functions. It also affected the VV&A process. Generally, evolving VV&A methods and tools lagged the development of the systems to be evaluated, but through deliberate investment in research (and a growing industrial software base), they were developed and eventually employed.

Today's VV&A includes early testing in software design and development, and rigorous automated testing in more advanced stages. It employs modeling and simulation of operationally realistic environments to stress the design envelope and validate performance. It also includes a cadre of tests for resistance to cybersecurity attacks and resilience (i.e., the ability to recover after compromise). Eventually, after gaining sufficient confidence, the system moves from the laboratory to field-testing in actual, even hostile, environments.

The goal – from the outset – is to design the software correctly and to test early to find and fix problems, which reduces costs and shortens schedules. As the system matures, the testing allows for more realism and counters more advanced threats. Typically, known inputs are matched with known system responses under a wide range of operating circumstances. This methodology, while disciplined and sophisticated, is still relatively linear and deterministic in its nature and will not work for autonomous learning systems whose behavior *adapts* based on its specific experiences. Like conventional (non-learning) systems, the "brain" of an autonomous learning system is its software. Software that incorporates machine learning (i.e., software having the ability to mutate based on exposure to inputs without being explicitly programmed) is difficult to analyze using traditional methods – and not intuitive to programmers. These systems typically begin in an untrained "empty" state and evolve, by changing their own behavior, over time.

With machine learning systems, a new approach and VV&A tool set is required. The VV&A methods and tools for learning systems need deliberate investment. These methods and tools are not sufficiently mature today. A significant step has been taken by the DoD Autonomy COI (Community of Interest) that is focused on a V&V strategy, including the Test and Evaluation of autonomous learning systems. But, it also requires attention to cybersecurity aspects of autonomy which is currently being handled in a different COI. The autonomy COI has representation from each of the services and OSD – with the Office of Naval Research and Naval Research Laboratory scientists participating.

While this report does not focus on cyber security, it would be incomplete without a clear acknowledgement of the over-riding importance of security as it applies to our autonomous systems. A successful cyber-attack by an adversary could potentially let that adversary use our autonomous systems against us. However, throttling our autonomy efforts because of this fear creates the risk that our adversary's autonomy capabilities surpass ours. The challenge will be to assure cyber security while at the same time accelerating progress in autonomy.

Autonomy computer code will operate as part of a larger system, including human interfaces, data systems, and other elements that provide vectors for cyber-attack. The system is only as strong as its weakest link, thus the use of systems with vulnerable operating systems (e.g., Windows XP) may defeat heroic measures elsewhere. Security of autonomous systems is not just about securing software on the platform. It also includes every other element of the network that interacts with – and enables – that software. The highly networked nature of military systems effectively amplifies the consequences of a loss of cyber integrity.

There are unique challenges associated with the security of autonomous systems that incorporate machine learning. Autonomous software code is influenced and changed by incoming data.

As we invest in efforts to secure all aspects of our national autonomy investment, we should simultaneously invest in offensive techniques to attack vulnerable aspects of adversary systems.

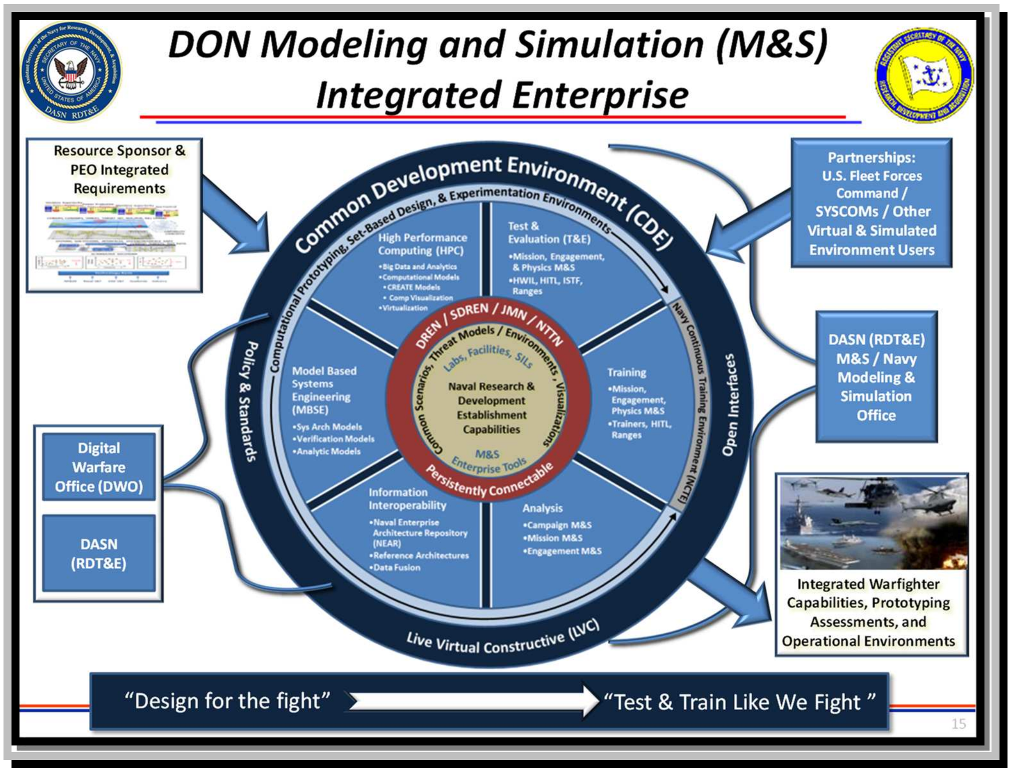

The Department of the Navy is making significant progress in the use of modeling and simulation for the design, development and testing of new systems as well as for operator training. Live, Virtual, Constructive (LVC) training is used throughout Navy and Marine Corps aviation. During the NRAC visit to NAVAIR, the DoN modeling and simulation (M&S) executive presented an overview of the M&S integrated enterprise. It is used across the life cycle of new weapon systems – from mission analysis, requirements development, software/hardware development, during test and evaluation, – and finally throughout the training cycle for aviation deployments. The joint simulation environment is a collaborative effort with the U.S. Air Force – another positive attribute of this ambitious and needed effort.

The M&S executive works closely with the Navy's Digital Warfare Office (DWO) reporting to N2N6. The DWO has the responsibilities that are essential for the successful introduction of autonomy across the DoN. The modeling and simulation and LVC efforts are needed to:

- advance and efficiently integrate leading edge M&S capabilities in response to complex emerging threats and future warfare environments;
- provide a comprehensive set of M&S capabilities and tools to support analysis, engineering, prototyping, testing, experimentation, and training to assess and address mission capability gaps and enhance interoperability;

- deliver and validate scenarios, simulations, and threat environments, at the appropriate level of fidelity, which can be reused throughout program life-cycles;
- promote consistent application of M&S "best practices" and associated technologies to improve support across the DoN enterprise; and,
- provide reusable connected test beds across the DoD using open architecture best practices.

The NRAC commends the DoN for the vision and organization plans for the DWO and the M&S executive.

.

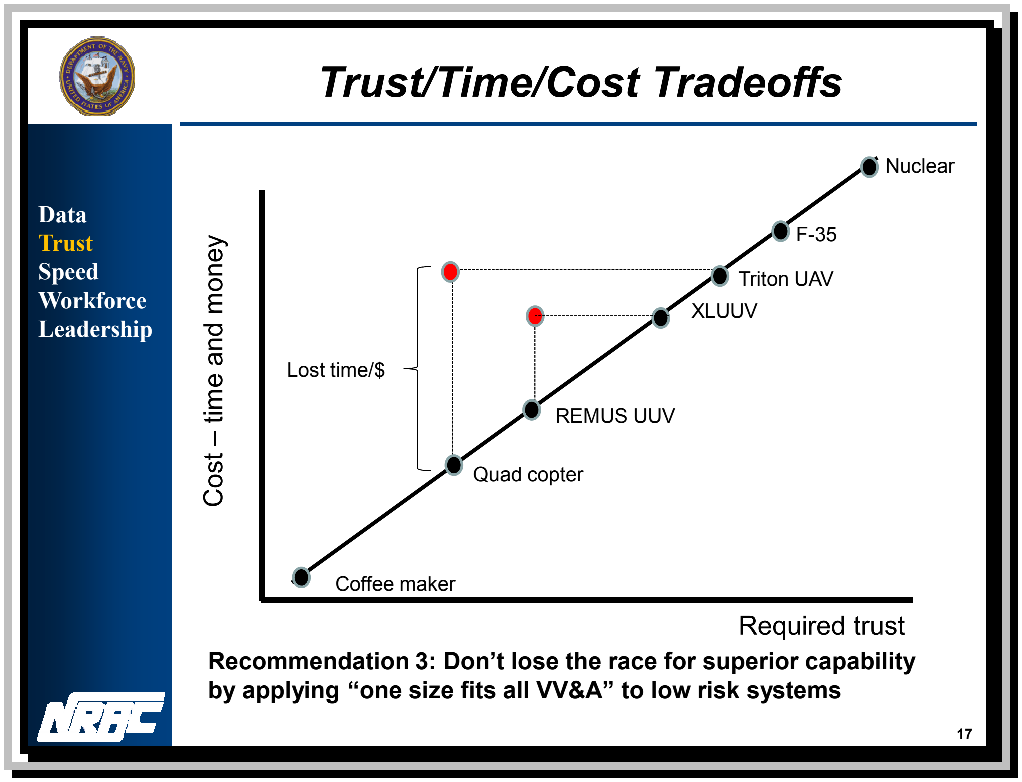

Recommendation 3: Don't lose the race for superior capability by applying "one size fits all VV&A" to low risk systems.

"Verification, Validation, and Accreditation" (VV&A) is ultimately about mitigating risk. The question it seeks to answer: can I trust this system to work as planned? If the system in question is a nuclear weapon or a nuclear reactor, we rightly demand that it work as advertised virtually all the time. In addition, we demand multiple layers of fail-safes, so that if it does fail, it does not fail in a catastrophic way. Because the risk of a nuclear system failing involves the possibility of large loss of life (and of military capability) we are willing to expend a great deal of resources to ensure trust in the technology. Obviously, it would be absurd to spend the same amount of resources to ensure the proper functioning of a coffee maker.

Likewise – though slightly less absurd – we should not expend the same amount of resources on the VV&A of a $1M REMUS UUV as we do on a $60M Orca Extra Large Unmanned Undersea Vehicle (XLUUV). We "pay" to gain trust in a system in the dollar cost of the VV&A process, but we also pay for this process in a manner that might be less in the forefront of our thinking – time. The extra time it takes us to field a new capability due to the extensiveness of VV&A processes costs us in unrealized operational capability.

Our adversaries – in the interest of fielding new capability quickly – are almost certainly willing to tolerate less trust in new autonomous systems than we are. The legal and ethical implications of autonomous systems are serious concerns for the United States – certainly more than they are for our principal adversaries. But 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 into the fleet – an unwillingness to take "bureaucratic risk". 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. Private companies like Amazon and Space-X are willing to accept reasonable risk where warranted. It is not an exaggeration to say that those who develop new warfighter technology felt that their attempts encountered debilitating bureaucratic 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.

That said, we do not intend to point the finger at any one group of people, any one element of the bureaucracy, or any one set of regulations or laws. 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 leader level when they fail to advance our capabilities fast enough to meet the strategic threat. The willingness to take program risk to achieve large gains in system capability is dampened by the Nunn-McCurdy Act. A well-intentioned measure to encourage realistic budgeting by the Department of Defense, this act tends to drive risk-aversion in order to avoid having a Nunn-McCurdy breach. The concerns addressed by Nunn-McCurdy must be balanced with the goal to ensure that U.S. technology advances at a pace faster than our adversaries.

Program failures below Nunn-McCurdy thresholds often trigger the creation of policies to prevent reoccurrence. Those who implement such policies tend to be overzealous because the accountability of subsequent failure is shifted to them. This ultimately furthers the risk-averse culture. Such policies are well intended, but their implementation serves to disconnect intent from practice. A current example is the effort to force the conduct of shock trials on the USS Ford (CVN-78). This test will add at least another year to the Ford achieving full operational capability. The value of such expensive testing might not do well in a thorough cost/benefit analysis using the myriad of data that are available based on previous shock trials of other CVs that have the same hull shape.

Another example: the Advanced SEAL Delivery Vehicle 1 (ASDS-1) prototype sustained a Lithium-ion battery fire while ashore that resulted in the loss of the prototype vessel. This failure is often cited when imposing restrictions for unmanned systems carried aboard ships and submarines today – despite the many subsequent improvements in battery technology and the fact that many new unmanned systems use relatively small batteries. Meanwhile, after some similar setbacks with lithium-ion batteries, other industries (e.g., electric automobiles, residential energy storage, commercial aircraft) have improved the technology, added some mitigations, and moved on. Today, thousands of homes have lithium-ion Tesla Powerwalls and all 600 Boeing 787s flying around the world have lithium-ion batteries installed throughout the aircraft. The difference? Tesla and Boeing assign engineers to solve the problem. The Navy takes administrative actions to eliminate (but NOT solve) the problem, without a process that forces all parties to consider the resulting lost opportunity.

The Department of the Navy has created strong incentives that focus on the risks associated with a system failing in the field or an acquisition program failing to achieve schedule and budget. Given our adversaries' rapid capacity and capability gains, we strongly urge the DON to change this incentive structure.

We now turn to the question of what specific strategies the Navy and Marine Corps might employ – and employ quickly without a multibillion-dollar cost – to jump-start both rapid innovation in autonomy and the rapid fielding of autonomous systems.

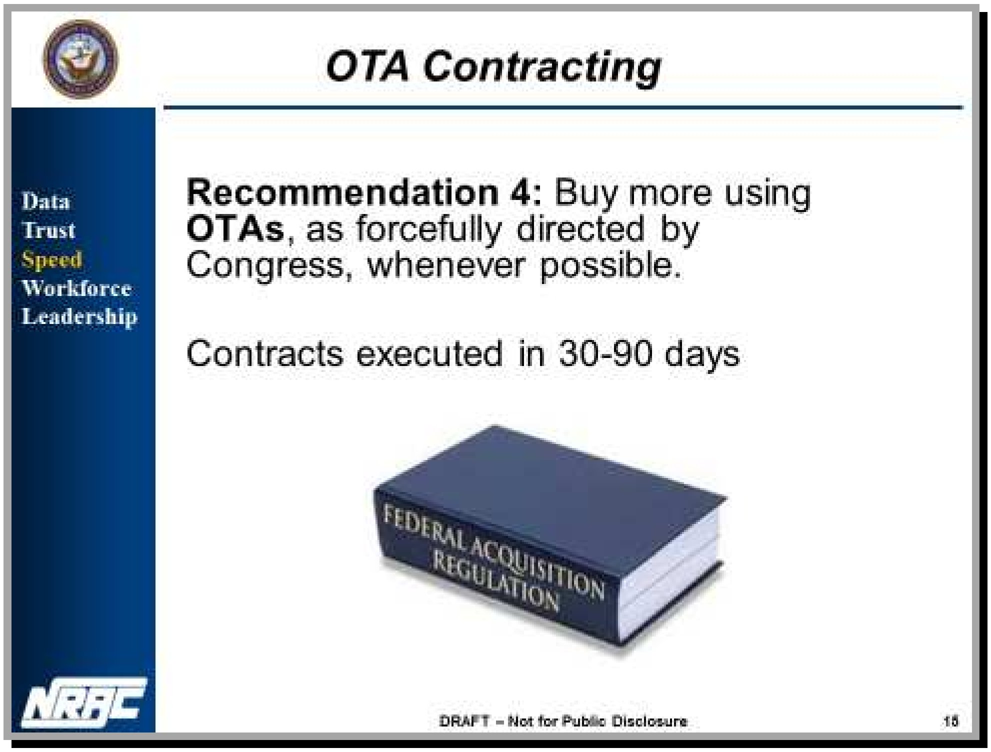

The speed of advanced system acquisition that relies on commercial technology must be rapid to ensure that the DoN paces the threat. **Recommendation 4**: Buy using OTAs (Other Transaction Authority), as forcefully directed by Congress, whenever possible.

The future pace of the advancement of Naval capability will be closely linked to the pace at which the Navy and Marine Corps mainstream autonomous capabilities. Autonomy techniques pioneered in the commercial sector promise to aid in rapid classification of surface, sub-surface, and air contacts provided by radar, sonar, and optical sensors; correlation of heterogeneous information sources; evaluation of courses of action; and precision delivery of effects.

The Navy, like the rest of the DoD, struggles to incorporate commercial technology at the pace at which it is developed. This is important, because the pace of commercial development probably defines the pace at which this technology becomes available to potential adversaries. In the future, failure to improve at the speed of the private sector implies failure to improve at the speed of the enemy. It would appear that Chinese commercial entities are in the "passing lane" if not already at the forefront of developing advanced autonomy, including the emerging field of artificial intelligence. Other competitors, such as Iran, have demonstrated the ability to be fast followers of commercial technology.

Unfortunately, the DoN is usually a slow follower. Part of the problem lies in the fact that Navy, like many large industrial organizations, chooses not to see the potential of technologies that might fundamentally affect established lines of operation. A major part of the problem involves developers in Naval labs/warfare centers and systems commands that feel hamstrung by a contracting process that is slow, overly cautious, and completely disconnected from the sense of urgency required to sustain our Naval dominance.

The NRAC recommends the following specific measures to accelerate the pace of Navy and Marine Corps innovation in general – and autonomy innovation – in particular:

Maximize the use of "Other Transaction Authority" (OTA), as directed by the Senate's 2018 National Defense Authorization Act. (S. 1519 Sec. 873). OTA allows rapid contracting for prototyping and rapid conversion of prototype projects to programs of record when desired. Despite the availability of OTA for years, today's Navy contracting establishment is not aggressively pursuing this option. Many people in the systems commands provided examples of projects that experienced delays on the order of 18 months in the completion of intermediate steps while waiting for their contracting office to procure essential components. In some cases, sequential contracting hurdles caused programs years of delay. In the opinion of the engineers and scientists we spoke with, their supporting contract offices simply did not share their commitment to rapid delivery of capability.

Inventory of autonomous systems does not guarantee their proper utilization in the Fleet, therefore senior leaders should be tasked with innovation in this area.

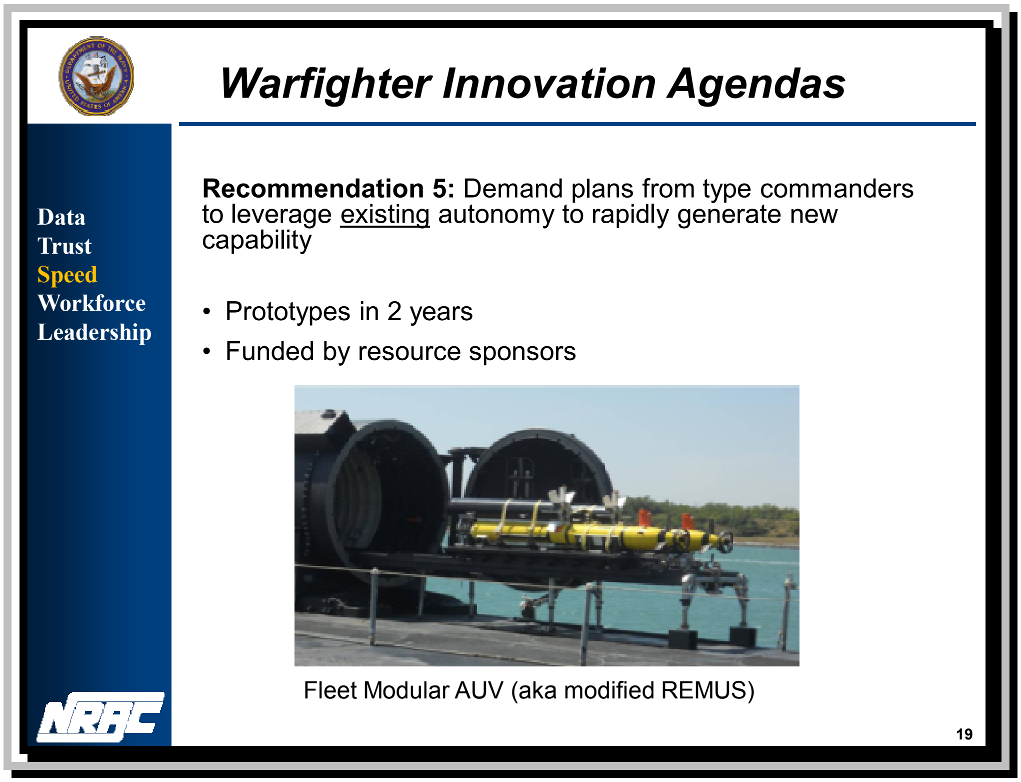

Recommendation 5: Demand plans (i.e., warfighter innovation agendas) from type commanders to leverage existing autonomy to rapidly generate new capability.

Warfighters have unique insight into the problems that they must overcome to prevail in the battlespace. Navy Type Commanders (i.e., warfare community leaders), are responsible to man, train, and equip their communities for their wartime mission. Part of the "equip" function involves providing the community with up-to-date tactics, techniques, and procedures. Type commanders understand the limits of their current capabilities, and in many cases, have the best insight in where to invest on the margins to optimally improve their warfighting impact. These commanders should be tasked to pursue innovation agendas to identify and prototype new ideas, including those involving autonomy. The Submarine Force provided an example of the desired type of action by creating the *Undersea Rapid Capability Initiative* that conceived, designed, tested, and deployed operational demonstrators in less than 24 months. The Fleet Modular AUV shown above is one example of a deployable capability produced under that program.

Recommendation 6: Create a "Shark Tank" to jump-start internal Navy and Marine Corps innovation funding within the DoN.

The Department of the Navy needs to initiate a method of funding good ideas and programs quickly. To do so, NRAC proposes allocation of ~$50M per year to an internally administered "Shark Tank" style innovation funding operation. The intent is to create the decision speed of a venture capital operation within the Naval establishment. The "sharks" would consist of a combination of successful DARPA/ONR program managers, systems command engineers, fleet representatives, and non-Navy venture capitalists. The non-Navy/Marine Corps elements are essential to ensure that the process does not get captured by the legacy establishment. It is essential that a funding decision be made by the panel within a few days of receipt of a proposal (a week or two at most). It is also essential that innovation funding decisions not be subject to review by the Naval establishment. Several of those interviewed by the NRAC panel contended that existing rapid innovation processes were co-opted by senior leaders to fill holes in legacy programs rather than fund true innovation.

The Office of Naval Research has administered a number of programs in recent years intended to accelerate the development of new ideas and capabilities. An example of the problems associated with existing programs is the *Chief of Naval Operations Speed to Fleet Program*, governed by OPNAVIST 3050.26 dated 9 April 2015. Under this program, innovative programs compete over the course of a year for innovation funding in accordance with the process outlined below.

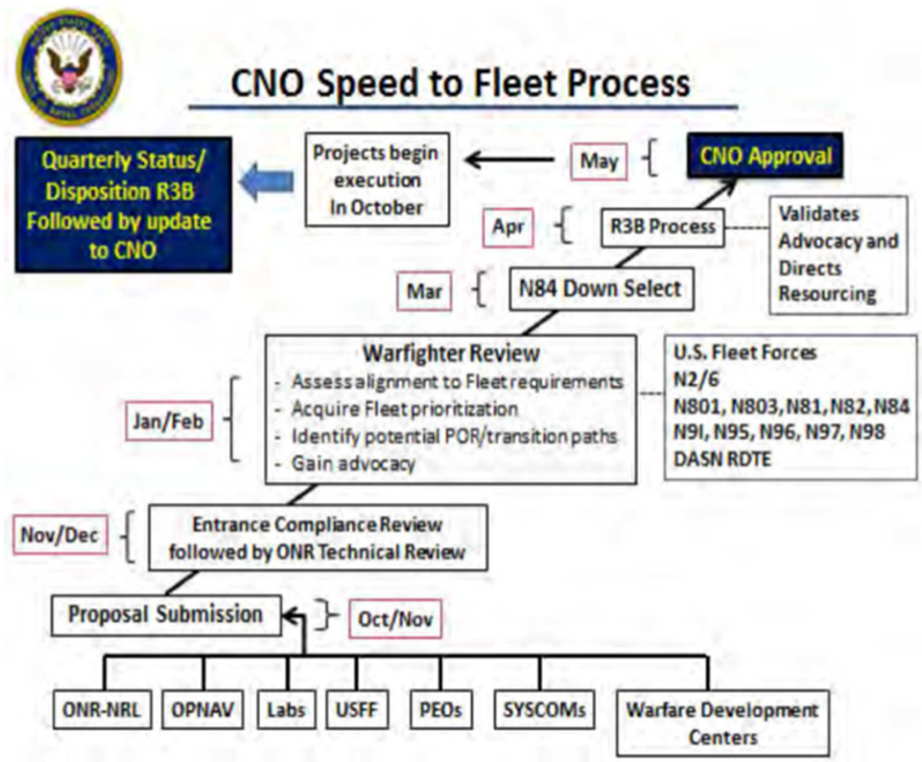

The process executes "speed to fleet" over the course of a fiscal year. By definition, this process adds a year to any development initiative that theoretically requires speed. Should the proposal be selected, it is at the cost of losing a year. Should the proposal not be selected, it has also lost a year, as the advocate now has to look to other means to develop a capability.

The most insidious aspect of these programs is that they allow the Navy and Marine Corps bureaucracy to identify themselves as moving quickly – when, in fact, they are not. The lifecycle of a cutting-edge autonomy program is about two years. A built-in one-year delay in starting a program becomes the first of many bureaucratic delays that compound almost ensuring system obsolescence at initial delivery.

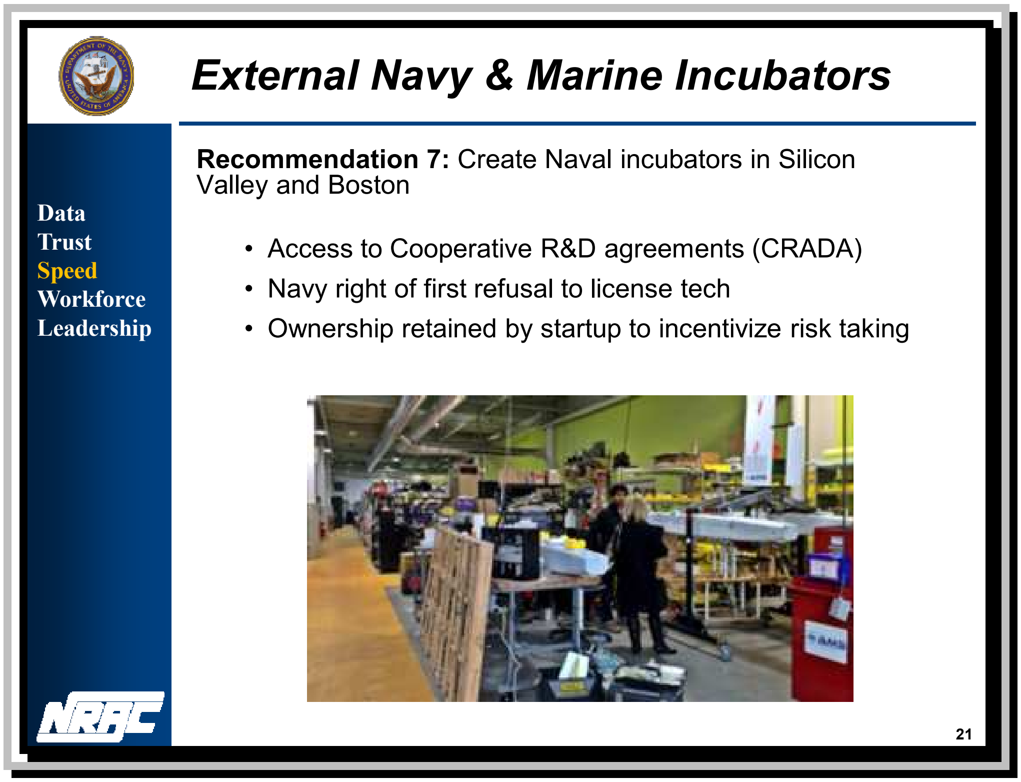

Recommendation 7: Create Naval-oriented incubators in Silicon Valley and the Boston area.

It can be argued that most of today's best scientists and engineers in the autonomy field live and work in a few concentrated regions, such as within California and Massachusetts. They are part of a $4 trillion world-wide information technology industry. Their ability to stay at the cutting edge of their field depends to a large degree on continuing to live and work in the most dynamic and productive environments. In order to leverage this talent base and encourage them into the Naval autonomy community, DoN needs to establish a presence in these locations.

The startups who work in Navy-sponsored incubators could benefit from reduced rent, access to their peers who support the Navy, and access to Cooperative R&D Agreements (CRADA), while retaining commercial rights to their intellectual property. Creating a startup environment in a Naval incubator is an essential step toward an overall goal of creating a pool of aggressive small companies that can rapidly deliver capability to the Navy in the same way that similar companies rapidly deliver capability to the private sector. Another direct benefit of Naval incubators would be priority access to license rights for the technology developed there.

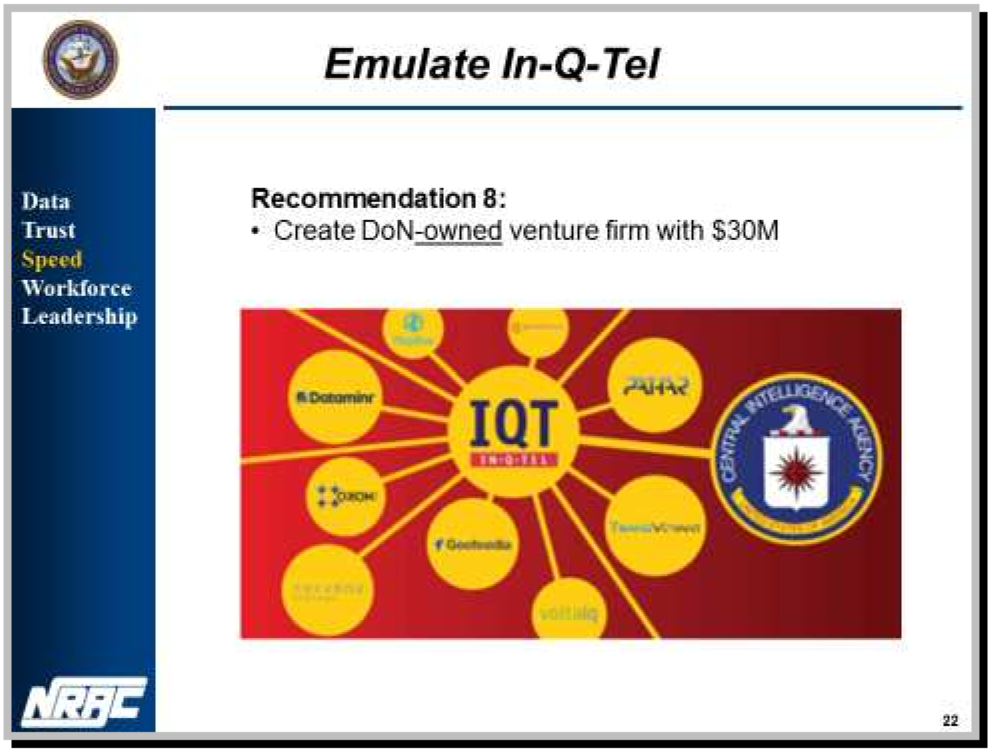

Recommendation 8: Create a DoN-owned venture firm with $30M.

It is possible for the Department of the Navy to stimulate interest in solving Navy and Marine Corps problems if it invests in start-up companies that have solutions. The CIA has demonstrated how this is accomplished. In the late 1990s, the CIA created a not-for-profit venture capital firm, In-Q-Tel, to invest in information technology start-ups. In-Q-Tel invests in small companies that are developing technologies of interest to the Intelligence Community. A number of promising applications including voice recognition, translation, map visualization, and malware screening were advanced to the benefit of both the CIA and the funded technology firms.

While the CIA has been successful in this area, the DoN and DoD have significant challenges interacting with cutting-edge technology firms. Senior defense leaders frequently visit the Silicon Valley and Boston areas to survey the sector and learn more about innovation practices, but seldom execute contracts with companies that do not align hiring, accounting, and security practices with Department standards. Tech firms derisively refer to Defense Department visitors as “tech tourists” because of the paucity of purchases relative to visits.

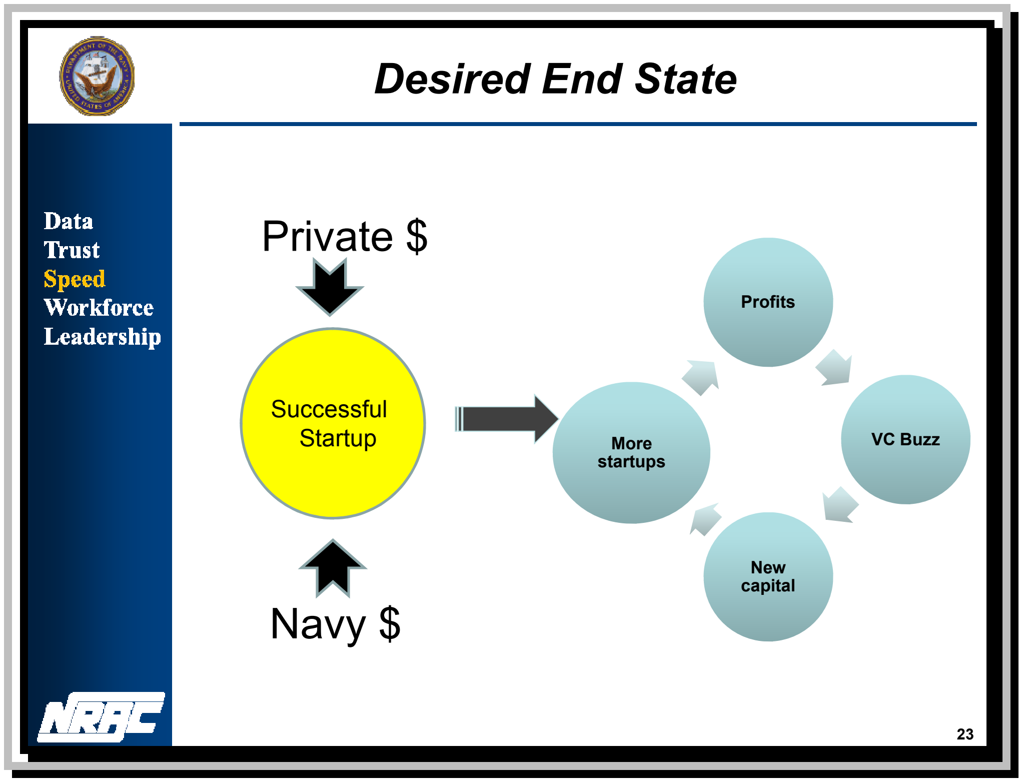

Creation of a Naval-oriented technology start-up sector – driven by private investment – would stimulate research funding for start-up entrepreneurs. Start-up successes, where Navy is the primary customer, could have a positive effect on private venture capital focused on Naval markets. Private venture capital is the most agile source of technology funding. The resulting ecosystem would be a virtuous cycle reinforcing itself through a feedback loop.

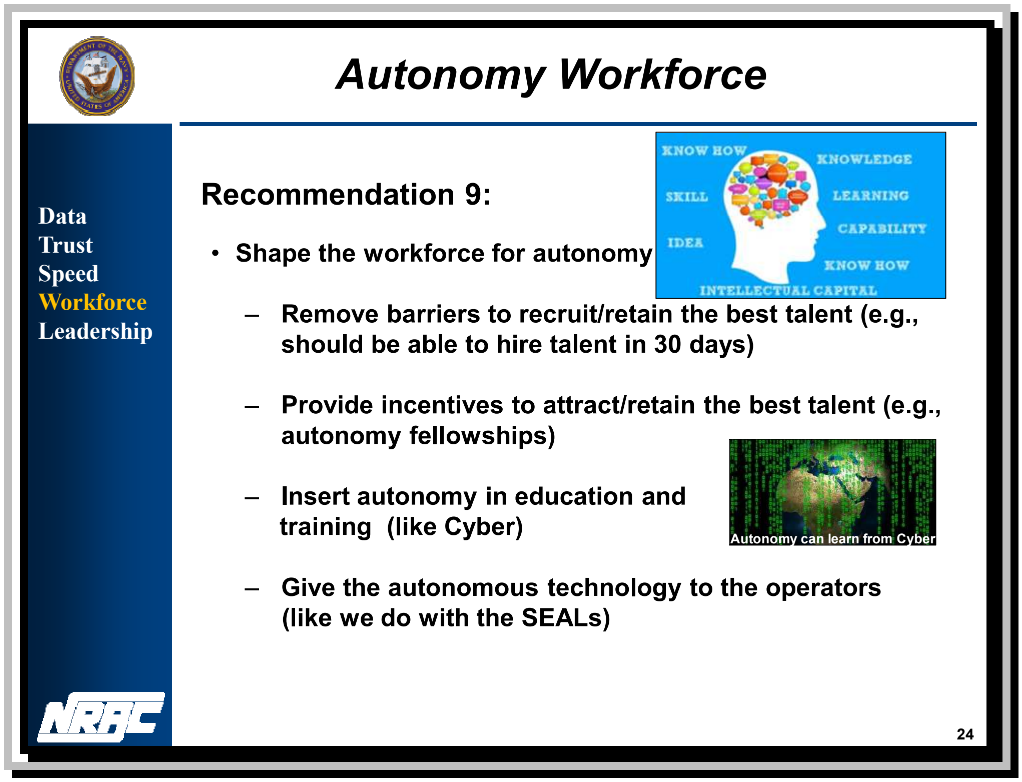

Recommendation 9: Shape the workforce.

The future of warfare – most decidedly – will be built around autonomy operations, and we must incorporate autonomy knowledge into our human capital at all levels.

Early advances in warfighter autonomy were driven by our nuclear-based deterrence requirements; later, military needs led to autonomy upgrades for precision-guided munitions; and intelligence, surveillance, and reconnaissance grids. However, most of today's autonomy advances are driven by the private sector. Working with and recruiting from the private sector will be essential for our defense establishment in the future.

For the Navy to successfully operate in this new environment, the NRAC recommends the removal of barriers to recruit and retain the best talent. A hiring procedure that lacks flexibility in minimum formal education requirements, elongates the time to make hiring decisions, and offers minimal compensation in the early years of service, will significantly jeopardize hiring the best candidates. Many of today's bright men and women are learning and acquiring autonomy-related skills through the internet and other non-traditional ways. They find that with their skill set, they do not need college degrees to earn good compensation in their technology field. According to Forbes Magazine, tech companies are hiring people without Bachelor's degrees and paying high starting salaries. Also, typical civilian companies can hire and place workers within 30 days. Talented foreign-born individuals should also become part of the perspective talent pool if current security requirements can be thoughtfully amended. The Navy and Marine Corps, where appropriate, must attempt to offer total compensation (salary plus benefits) comparable to the private sector.

Incentives should be offered to talented individuals early in their formative years. Offering internships is an opportunity to expose young researchers to the challenges of DoN, and to start relationships that ultimately follow these individuals through their careers. Also, autonomy-related research proposals by academic institutions should include internships at defense R&D facilities and/or autonomy-related startups. This will enhance the transfer of technology and foster recruitment of talent. Building connections with faculty involved in autonomy education and research is critical to identify, attract, and recruit the best talent in academic institutions.

Autonomy is not a stovepipe domain: it has applicability to all aspects of the Naval force. Therefore, it is imperative to insert autonomy in education and training for sailors and Marines. A critical mass of autonomy specialists is needed to foster progress and adoption of autonomous applications. This will require widespread, pervasive changes within the Naval education and training establishment. Autonomy should be inserted in Navy "A" schools, Marine Corps training, NROTC, and Naval graduate schools. Similar efforts within the Cyber domain led to the establishment of the *Cyber Security Studies* program at the U.S. Naval Academy. As with the cyber enterprise, autonomy education should incorporate multidisciplinary elements in formal classroom settings, as well as hands-on experience in operational units. For selected active duty individuals, internships should be offered at systems commands, research activities and industries that are developing autonomy. Navy and joint service war-games should include autonomy applications.

It should be a high priority to give autonomy applications to the operators as soon as it is prudent to do so. This has proven to be highly successful with the Naval Special Warfare Command. A SEAL team can provide valuable feedback to developers in real-time, and begin the development of training, tactics, and procedures.

Recommendation 10: the Secretary of the Navy must be the Champion for Autonomy.

Modern day transformative changes in the Navy such as nuclear powered submarines, the AEGIS Combat System, and Submarine Launched Ballistic Missiles were focused on defined systems that transformed Naval warfare capabilities. In each case, they were led by transformative leaders who gained support not only from DoD leadership – but just as importantly – from the Congress.

Unlike the focused applications and singular leaders of “warfare changing” technology noted above, other significant transformations in warfare (noted in the DoD Second Offset Strategy – including Intelligence, Surveillance, Reconnaissance, stealth, precision-guided munitions, space-based communications and navigation) applied across many systems and warfare areas. These were similar to the earlier developments in gunpowder, steam, steel, and the electromagnetic spectrum. These technologies did not lend themselves to be led by a single transformative leader.

The coming revolution of autonomy will be similar to the technologies of the Second Offset and should not be led by a single autonomy domain leader. It will require multiple Navy and Marine Corps leaders to pull autonomy technologies into the battle force.

That said, it is still necessary that an Autonomy Champion exist to provide the required resources, funding and personnel. Therefore, the NRAC panel strongly recommends that the Secretary of the Navy be that Champion.

Because the pacing threat today is in the Pacific, we recommend that the Commander of the Pacific Fleet be the leader implementing autonomy technology and systems for the Fleet. To do this, PACFLT will need a dedicated and focused autonomy leader on the staff. For the Marine Corps, autonomy pull and implementation should be led by the Commanding General of the Marine Corps Combat Development Command working with Marine Corps operational commanders.

Direct support to the autonomy leaders from the Warfare Type Commanders and the Navy's Digital Warfare Office will be mandatory. These supporting Navy commands and those in the Marine Corps must be appropriately resourced in terms of staff and funding.

Of equal importance to having the right leaders to imbed autonomous systems within the fleet, increased and stable funding resources must be made available throughout the research, development and deployment of these systems. Key to the effort of pushing and pulling autonomy to the fleet will be type commanders, systems commands, the Naval Warfare Development Command and others in bridging the "technical to the tactical".

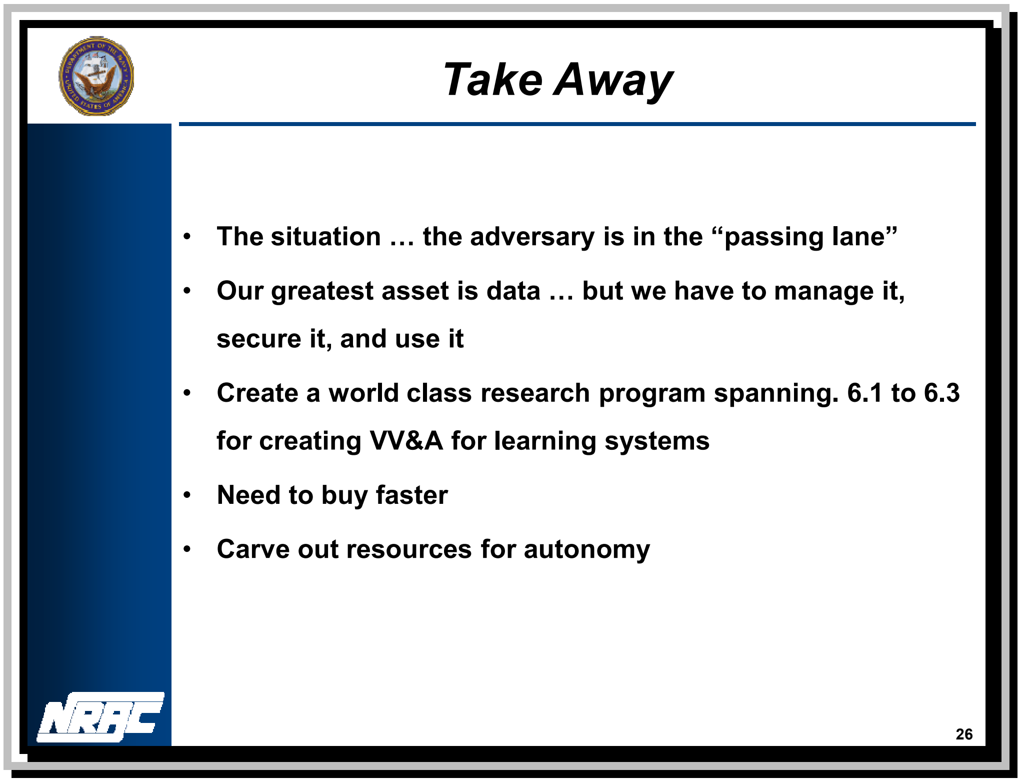

Autonomous systems will transform warfare. In the past, some of the enabling autonomy technologies were seeded by research investments from the U.S. government. Today the commercial world has taken the lead in the development of autonomous capabilities.

Adversaries are catching up quickly. In some cases, they maybe already be ahead of us. Data is our greatest asset. The DoN must internalize this message, and organize around it.

The challenge of assuring that our autonomous systems will do what we want them to do, and remain secure, requires immediate and sustained research. New learning systems do not fit into today's paradigms for VV&A. We need a world class research program to develop trusted VV&A of autonomous systems.

Experimentation, prototyping, and procurement are all bottlenecks in the current system. Inability to quickly iterate on the first two inhibit the concept exploration essential for defining well-structured Programs of Record as well as connecting the fleet's needs with the research community's cutting edge technology. Speed of procurement is essential when technology is changing rapidly, as it is with autonomy.

Introducing autonomy throughout the Department of the Navy will be transformational. It must be resourced adequately. The transformation requires leadership of the Secretary of the Navy.

Recommendations (pg 1 of 2)

R1: Create a comprehensive data plan to field superior autonomy
- Develop, institute and maintain a single, unified, Navy-wide plan for the use of corporate data
- Deem corporate data as a vital asset for strategic and operational decision-masking
- Define data governance roles (e.g., strategic leadership, data stewardship, education, data solutions architect, data quality lead, and data administrator)
- Establish policy to coordinate across the different naval units and address issues such as data redundancy
- Establish standards to support the fusion and interoperability of data from different sources
- Insure data quality for accurate training of learning systems
- Include provenance and metadata to provide trust
- Protect data
- Consider the ingestion of unstructured data (e.g., text, social media data, and images)
- Consider software as data and establish trusted software modules in the plan

R2: Create a world class research program spanning 6.1 to 6.3 for developing the tools and processes for VV&A for learning systems.

R3: Revise policies to enable rapid test and evaluation of unmanned and autonomous systems, commensurate with risk.

R4: Buy much more using OTAs, as forcefully directed by Congress, whenever possible.

R5: Demand plans from type commanders to leverage <u>existing</u> autonomy to rapidly generate new capability.

R6: Establish an internal DoN "Shark Tank"
- Decision negotiated and made 'on the spot'
- $50M pilot
- Delegate spend authority 'sharks' (venture partners)

R7: Create Naval incubators in Silicon Valley and Boston
- Access to Cooperative R&D agreements (CRADA)
- Navy right of first refusal to license tech
- Ownership retained by startup to incentivize risk taking

NRAC

27

Recommendations (pg 2 of 2)

R8: Create DoN-<u>owned</u> venture firm with $30M

R9: Shape the workforce for autonomy
- Remove barriers to recruit/retain the best talent
  - → Waive educational requirements to hire autonomy talent
  - → Speed up the hiring process (e.g., hire talent in 30 days)
  - → Offer total compensation comparable to the private sector
- Provide incentives to attract/retain the best talent
  - → Award **Autonomy Fellowships** funded by ONR
  - → Offer paid internships (10 weeks) at Naval Laboratories
  - → Require inclusion of internships at Naval Labs and/or autonomy-related startups in all research proposals
  - → Sponsor faculty to work in Naval Laboratories
  - → Provide strategic opportunities for non-national talent (n=100) to stay in the USA and join the Naval workforce
  - → Allow autonomy talent that leaves the Naval workforce to return at a later time without jeopardizing their career development
- Insert autonomy in education and training (like Cyber)
  - → Insert across the Navy and Marine Corps, including the Navy enlisted "A" school, Marine Corps MOS schools, Naval Academy, ROTC, and Naval graduate schools
  - → Integrate internships at Navy operational and technical commands, Naval Research Labs and industries that are developing Autonomy
  - → Incorporate autonomy in wargames
- Give the autonomous technology to the operators (like we do with the SEALs)

R10: SECNAV must be the champion
- Carve out resources for autonomous systems
- Designate (e.g. COMPACFLT & CG MCCDC) Autonomy Domain Leads
- Choose charismatic leaders who can drive change and elicit support within the DoN, DoD, and Congress
- Tour lengths commensurate with the complexity of the activity

NRAC

28

R1: Create a comprehensive data plan to field superior autonomy

- Develop, institute and maintain a single, unified, Navy-wide plan for the use of corporate data
- Deem corporate data as a vital asset for strategic and operational decision-masking
- Define data governance roles (e.g., strategic leadership, data stewardship, education, data solutions architect, data quality lead, and data administrator)
- Establish policy to coordinate across the different Navy and Marine Corps units and address issues such as data redundancy
- Establish standards to support the fusion and interoperability of data from different sources
- Insure data quality for accurate training of learning systems
- Include provenance and metadata to provide trust
- Protect data
- Consider the ingestion of unstructured data (e.g., text, social media data, and images)
- Consider software as data and establish trusted software modules in the plan

R2: Create a world class research program spanning 6.1 to 6.3 for developing the tools and processes for VV&A for learning systems.

R3: Revise policies to enable rapid test and evaluation of unmanned and autonomous systems, commensurate with risk.

R4: Buy more using OTAs, as forcefully directed by Congress, whenever possible.

R5: Demand plans from type commanders to leverage existing autonomy to rapidly generate new capability.

R6: Establish an internal DoN "Shark Tank"

- Decision negotiated and made 'on the spot'
- $50M pilot
- Delegate spend authority 'sharks' (venture partners)

R7: Create Navy and Marine Corps incubators in Silicon Valley and Boston

- Access to Cooperative R&D agreements (CRADA)
- Access to appropriate Naval data sets, personnel, networks, etc.
- Navy right of first refusal to license tech
- Ownership retained by startup to incentivize risk taking

R8: Create DoN-owned venture firm with $30M

R9: Shape the workforce for autonomy

- Remove barriers to recruit/retain the best talent
- Waive educational requirements to hire autonomy talent
- Speed up the hiring process (e.g., hire talent in 30 days)
- Offer total compensation comparable to the private sector
- Provide incentives to attract/retain the best talent
    - Award Autonomy Fellowships funded by ONR
    - Offer paid internships (10 weeks) at Naval Laboratories
    - Require inclusion of internships at Naval Labs and/or autonomy-related startups in all research proposals
    - Sponsor faculty to work in Naval Laboratories
    - Provide strategic opportunities for non-national talent (n=100) to stay in the USA and join the Navy and Marine Corps workforce
    - Allow autonomy talent that leaves the Naval workforce to return at a later time without jeopardizing their career development
- Insert autonomy in education and training (like Cyber)
    - Insert across the Navy and Marine Corps, including the Navy enlisted "A" school, Marine Corps MOS schools, Naval Academy, ROTC, and Naval graduate schools
    - Integrate internships at Navy operational and technical commands, Naval Research Labs and industries that are developing Autonomy
    - Incorporate autonomy in wargames
- Give the autonomous technology to the operators (as we do with the Navy SEALs)

R10: SECNAV must be the champion

- Carve out resources for autonomous systems
- Designate (e.g. COMPACFLT & CG MCCDC) Autonomy Domain Leads
- Choose charismatic leaders who can drive change and elicit support within the DoN, DoD, and Congress
- Tour lengths commensurate with the complexity of the activity

Appendix A

Use and Acquisition of Unmanned Systems in the Department of the Navy
Naval Research Advisory Council (NRAC)

TERMS OF REFERENCE

These terms of reference establish the Secretary of the Navy (SECNAV) objectives for the Naval Research Advisory Committee (NRAC), a permanent subcommittee of the SECNAV Advisory Panel (SNAP), to conduct a study on the use and acquisition of unmanned systems throughout the Department of the Navy (DoN).

Mission Statement: Examine current DoN Unmanned Systems (UxS) policies and acquisition strategies. The goal of the study is to provide recommendations to the Secretary on the application of autonomous and robotic systems across the naval mission set to include levels of autonomy, learning machines, and human-machine teaming.

Additionally, conduct a detailed examination of current acquisition strategies for Navy and Marine Corps UxS including how to ensure common open systems architectures. These common architectures should be addressed on three levels: autonomy, the physical vehicles, and the integration and data sharing across and throughout the multiple physical domains in which the Navy and Marine Corps operate.

The acquisition strategy study should include an examination of the verification, validation, and accreditation (VV&A) policy for autonomous and robotic systems including an assessment of the role of modeling and simulation. The study should further include a discussion of the approach for risk acceptance for non-deterministic autonomous and robotic systems.

Issue Statement: SECNAV has established proliferation of UxS as a DoN strategic goal. However, UxS does not fit the traditional acquisition model. Challenges include: the speed of technological advance; reduced involvement of human operators; non-deterministic behavior of the autonomous systems; different VV&A requirements; and the cost and interoperability advantages of modular software and hardware components. These challenges are amplified by the desire for commonality in architectures throughout the multiple operational domains. Additionally, systems to date have largely relied on discrete interfaces that impede shared functionality and interoperability across subsequent programs of record and domains.

Traditional VV&A procedures can impose testing schedule and monetary costs that account for a large share of UxS program resources. Moreover, traditional testing regimes can result in extended delays. Modeling and simulation is increasingly used in programs of record such as the heavyweight torpedo to offset the number of in-water experiments required for certification. The current fiscal environment requires design choices that reduce cost and schedule across the DoN UxS portfolio rather than just in a single program of record.

Most UxS use proprietary software that cannot be readily adapted by DoN without incurring high costs. The prevalence of proprietary software also limits the interoperability of UxS with many Navy platforms. Improving interoperability and commonality at the system level will increase opportunities for modular mission sets. Furthermore, commonality promotes opportunities for small businesses to develop software and hardware payload modules and empowers program managers to reduce cost and schedule burdens by accessing a library of software and hardware that function across DoN UxS. Increased commonality will expedite enhanced autonomy, expand the capabilities of learning machines, and facilitate human-machine teaming throughout the entire naval mission set.

Objectives and Scope: NRAC will address the following specific objectives:

a. Assess how autonomous systems operating across the different domains will inform the common architecture discussion on three levels: architecture of the autonomy, architecture of the unmanned vehicles, and architecture of the integration and data sharing across multiple domain platform hosts.

b. Examine the issue of establishing effective and efficient Verification, Validation, and Accreditation (VV&A) policy for autonomous systems. Assess the role that modeling and simulation will play and how to approach risk acceptance for VV&A policy.

c. Assess how DoN should apply different levels of autonomy, learning machines, and human-machine teaming across the naval mission set.

Methodology: NRAC assessments will be conducted in compliance with all pertinent regulations of the Federal Advisory Committee Act (FACA).

Deliverables: Because of the scope of the work and relevance of the problem to future UxS policy and acquisitions, NRAC shall provide a report by June 2017. NRAC will deliver progress updates in October 2016 and February 2017.

Membership: The members will address the task as delineated below:

a. NRAC: Dr. Bellingham; VADM Bowes (Ret.); Dr. Bruno; Dr. Gates; Dr. Padilla; Dr. Walsh; RADM Young (Ret.).

b. Per the SNAP Charter, non-voting subject matter experts (SMEs) may be appointed to assist SNAP or its subcommittees on an ad hoc basis to address specific issues under consideration. These SMEs are not members of SNAP or its subcommittees and will not engage or participate in any deliberations.

<u>Support</u>: The Department of Defense through the Office of the Secretary of the Navy, shall provide support, as necessary, for the performance of the committee's functions, and shall ensure compliance with requirements of the FACA and established DoD policies and procedures.

Appendix B: Study Members

NRAC Members

Dr. James Bellingham (Chairman)

Dr. Bellingham is currently the Director for the Center for Marine Robotics at the Woods Hole Oceanographic Institute (WHOI), Woods Hole, MA where his combination of a strong academic research background, entrepreneurial business savvy, and experience working with other academic, nonprofit, private sector, and government sectors will help advance the field of marine robotics and speed its application for science and a broad range of other uses. Prior to his appointment to WHOI, he was the Chief Technologist at the Monterey Bay Aquarium Research Institute (MBARI), Dr. Bellingham identifies technical and engineering opportunities for the ocean sciences aligned with MBARI's strategic plan. He also acts as a liaison for MBARI to the ocean engineering community. Prior to serving as Chief Technologist, Jim was Director of Engineering at MBARI from 1999 to 2006. During his tenure he was responsible for elevating MBARI's Engineering Department to international stature, and expanding its focus from ROV design to advanced ocean observing system development. As Director of Engineering, Jim served as a member of MBARI's Management Team, and continues to provide advice on priorities for technical and engineering projects, particularly during MBARI's internal proposal process. Dr. Bellingham returned to full time research in 2006, and his current research focus is the development of observing systems that leverage mobile robotic capabilities to provide an unprecedented view of the physical, chemical and biological ocean. He led the multi-institutional Autonomous Ocean Sampling Network effort through two phases of development, beginning in 1995. AOSN has generated a range of widely adopted AUV technology. Today Jim is developing a new generation observation systems tailored to the needs of global climate and ocean ecosystem studies. Cyber infrastructure for data management, decision support, and data exploration is a growing element of Jim's technology research program.

VADM William "Bill" Bowes, USN (Ret.) (Vice Chairman)

VADM Bowes is a retired Navy Vice Admiral, having served 33 years in the Navy in numerous operational and acquisition assignments. As a Vice Admiral he served as the Commander of the Naval Air Systems Command, the Principal Deputy Assistant Secretary of the Navy for Research, Development and Acquisition (RDA), and the Acting Assistant Secretary of the Navy for RDA. As a Naval Aviator he commanded an A-7E squadron, flew 350 combat missions during the Viet Nam conflict and has flown over 5000 hours in more than 50 different US and foreign military aircraft types. He is an accomplished test pilot and program manager. He served as the program manager for the F-14 and Phoenix missile program, the Joint Cruise Missiles Project, which developed and deployed the Tomahawk cruise missile, and was the first director of DoD's Joint Unmanned Aerial Vehicles Project. After retiring from the Navy, Bowes joined Hughes Aircraft as a Senior Vice President and Deputy General Manager of the newly forming

Sensors and Communications Sector. After Hughes was acquired by Raytheon, Bowes joined Litton Industries as the Vice President, Corporate Strategic Planning, and subsequently led the creation of the Military Aircraft Electronics Systems business unit after Litton was acquired by Northrop Grumman. Since retiring from Northrop Grumman he has served as a director on a number of public company, private company and non-profit boards as well as serving on the Carnegie Mellon University Software Engineering Institute Board of Visitors.

Dr. Ann Gates

Dr. Gates is Professor and Chair of the Computer Science Department at the University of Texas at El Paso. Her areas of research are in software engineering and cyberinfrastructure with an emphasis on workflows, ontologies, and formal software specification. Gates directs the NSF-funded Cyber-ShARE Center that focuses on developing and sharing resources through cyber-infrastructure to advance research and education in science. She was a founding member of the NSF Advisory Committee for Cyber-infrastructure. Dr. Gates served on the IEEE-Computer Society (IEEE-CS) Board of Governors 2004-2009, chaired the IEEE-CS Educational Activity Board's Committee of Diversity and External Activities and was given the IEEE-CS Golden Core Award for her service to the IEEE-Computer Society. In addition, she served on the Computer Science Accreditation Board (2011-2013). Dr. Gates leads the Computing Alliance for Hispanic-Serving Institutions (CAHSI) and is a founding member of the National Center for Women in Information Technology (NCWIT). She received the 2015 Great Minds in STEM's Education award, the 2015 A. Nico Habermann Award, the 2010 Anita Borg Institute Social Impact Award, the 2009 Richard A. Tapia Achievement Award for Scientific Scholarship, Civic Science, and Diversifying Computing and was named to Hispanic Business magazine's 100 Influential Hispanics in 2006 for her work on the Affinity Research Group model.

Dr. Ingrid Padilla

Dr. Padilla is currently a professor in Environmental and Water Resources Engineering in the Department of Civil Engineering and Surveying at the University of Puerto Rico, Mayagüez. She has been at the University of Puerto Rico since 2001. Prior to that, she worked with Greg Morris and Associates and the U.S. Geological Survey. Dr. Padilla holds a Ph.D. from the University of Arizona; a M.S. from the University of Michigan; and a B.S. from the University of Maryland. She has directed laboratory, field-scale, and modeling environmental engineering and ground-water investigations in the academic, government, and private sectors, and has been awarded with numerous scientific grants (NSF, NIH, DoD, DoE, EPA). Her training and experience include chemical detection in subsurface environments, characterization and quantification of fate and transport processes in soils, environmental remedial technologies, collection and analysis of hydrologic data; hydrogeological characterization of alluvial, karst, and volcanic-rock aquifers; evaluation and characterization of aquifer in coastal zones; aquifer testing; water quality sampling and monitoring; well evaluation and design; characterization and modeling of surface water and groundwater interactions; and sorption of organic contaminant onto particles. Dr. Padilla has received numerous awards through her career, including: Distinguished Professor in Civil Engineering; Innovative Woman in Engineering Education, Ford Foundation Fellowship, and several outstanding minority awards. She has presented her work in many local, regional, and national conferences, and published in distinguished journals and proceedings.

Dr. Joseph (Jay) T. Walsh Jr.

Dr. Walsh is the Vice President for Research at Northwestern University where he helps develop and implement the strategic plan for the university's research operations and where he oversees the research infrastructure on the campuses in both Evanston and Chicago, Illinois. Dr. Walsh is also a Professor of Biomedical Engineering. He has served on the Board of Governors for Argonne National Laboratory since 2008 and the Board of Directors for the Chicago Council on Science and Technology since 2007. Previously, Dr. Walsh served on the Board of Directors for Fermi National Laboratory and as a member of the Illinois Governor's Innovation Council. His research area is the study of light-tissue interactions. Previously, he investigated the photophysics and photobiology of laser-based ablation. More recently he has investigated tissue birefringence feedback systems, the propagation of polarized light in tissue, optically induced stimulation of the auditory system, and nanostructured surfaces for biosensing applications. He has been the principle investigator on several NSF and NIH grants as well as industry sponsored translational research. Dr. Walsh has been a program chairman for 5 major conferences in his field. He is a past-president of the American Society for Laser Medicine and Surgery, the world's premier medical laser society. Dr. Walsh conducted his doctoral research on the medical applications of laser and other optical sources in the Wellman Laboratories at the Massachusetts General Hospital, received his Ph.D. in Medical Engineering from the Harvard-MIT Division of Health Science and Technology, and BS and MS degrees in Electrical Engineering from MIT.

RADM Charles Young, USN (Ret.)

RADM Young served on the *USS ULYSSES* S. *GRANT* (SSBN 631B); *USS PLUNGER* (SSN 595); *USS SAND LANCE* (SSN 660); USS *SAN JUAN* (SSN 751) and *USS HOLLAND* (AS 32). Shore duty assignments included instructor duty at Nuclear Power School, Bainbridge, Maryland; Squadron Material Officer on the staff of Commander Submarine Squadron Sixteen in Kings Bay, Ga.; Director of Tactical Training at the Navy Fleet Ballistic Missile Submarine Training Center in Charleston, S.C.; Deputy Commander for Readiness and Training for Submarine Squadron TWO and Undersea Warfare Assistant Office Director for Advanced Submarine Technology in the Defense Advanced Research Projects Agency. Returning to Washington, DC in August 1994, Admiral Young assumed duties as Director, Resources and Evaluation on the staff of the Assistant Secretary of the Navy for Research, Development and Acquisition. He was the Program Manager for the Navy's Unmanned Undersea Vehicles Program Office from June 1995 to October 1997. From October 1997 to July 2001 he served as Deputy Commander, Naval Sea Systems Command, Undersea Technology. Rear Admiral Young was the Commander, Naval Undersea Warfare Center from October 1998 to July 2001.

He served additional duty as the Vice Commander, Naval Sea Systems Command from August 1999 to January 2000 and was the Program Executive Officer for Undersea Warfare from February to April 2000. Admiral Young is a graduate of both the Program Management Course and the Executive Program Management Course at the Defense Systems Management College. He served as Vice Commander, Naval Sea Systems Command from April 2001 to July 2002. Rear Admiral Young became the 11th Director of Strategic Systems Programs in July 2002 where he was responsible for all aspects of the research, development, production, logistics, storage, repair, and operational support of the Navy's Fleet Ballistic Missile Weapon Systems, that include the TRIDENT I and II missiles and their associated shipboard subsystems. He was also the U.S. Project Officer responsible for managing U.S. Government support of the British POLARIS/TRIDENT Force. Since retirement from the Navy, Admiral Young has served on the Defense Science Board Task Force on the National Security Industrial Base for the 21st Century. He has also served as an advisor to the Threat Reduction Advisory Committee Nuclear Deterrent Transformation (NOT) Panel.

Subject Matter Experts

VADM Michael J. Connor, USN (Ret)

VADM Connor founded ThayerMahan Inc. in order to accelerate the United States' ability to effectively and efficiently monitor ocean activity using autonomous systems. Recognized globally as one of the foremost authorities in undersea operations, Mike brings a wealth of experience. In a 35 year career in which he rose to the rank of Vice Admiral in the United States Navy, Mike commanded at the ship, squadron, and task force levels. His assignments include command of USS SEAWOLF, a nuclear-powered attack submarine, Submarine Squadron EIGHT, Undersea Forces in the Western Pacific, and the Arabian Gulf, the United States Submarine Force, and NATO's Allied Submarine Command. He led the US Navy Submarine Force move into robotic undersea systems, achieving key milestones including the first operational deployment and recovery of an unmanned vehicle from a submarine. He led an innovation effort that began the shift away from undersea search operations based on expensive platforms and moved toward operations based on large numbers of inexpensive vehicles. Mike has written extensively on the future of undersea warfare and is a sought-after speaker on undersea warfare topics. His education includes a B.A. in Physics from Bowdoin College and an M.A. in National Security Studies from the United States Naval War College.

Dr. Kevin R. Fall

Dr. Fall is the former Chief Technology Officer and Deputy Director of the Software Engineering Institute (SEI) at Carnegie Mellon University. As CTO and the director of research at SEI, his responsibilities included oversight and selection of a technical portfolio in software engineering and cybersecurity. Prior to this position, Dr. Fall held positions at Qualcomm as a Principal Engineer and Consultant and the Sandia National

Laboratories as Consultant. As a member of the US Air Force Scientific Advisory Board (AFSAB), he provided technical advice on matters of concern to the Chief of Staff and Secretary of the Air Force. His scholarly work includes teaching at the University of California, Berkeley, University of California, Santa Cruz, University of California, San Diego and postdoctoral work at the Massachusetts Institute of Technology. Dr. Fall holds a B.A. in Computer Science from the University of California, Berkeley and a Ph.D. in Computer Science from the University of California, San Diego.

Mr. Scott M. O'Neil

Mr. O'Neil retired from civil service in January 2016 after working for the Navy for over 43 years. For the decade prior to his retirement, Mr. O'Neil served as Executive Director and Director for Research and Engineering at the Navy's premiere weapons laboratory, the Naval Air Warfare Center Weapons Division (NAWCWD) at China Lake, CA. There he directed a highly technical civilian workforce of over 5000 scientists and engineers. Mr. O'Neil was appointed to the Senior Executive Service in November 1998 as the Head of NAWCWD's Weapons and Targets Department. In April 2002, he was assigned as the head of the Naval Air Warfare Center Aircraft Division, Atlantic Ranges and Facilities Department at Patuxent River, MD. From May through December 2010 he served as the Acting Deputy Commander, Naval Air Systems Command. Prior to moving into general management, Mr. O'Neil worked for 15 years in the fields of solid rocket propulsion and thrust vector control systems. He was a program manager for Vertical Launch ASROC propulsion and control system, an in-house advanced development program; that had initial operational capability in 1993. He was also Deputy Program Manager for Tomahawk Cruise Missile rocket motors. He led the development MK III Improved Rocket Motor and the WDU 36 lightweight warhead system for Tomahawk. Mr. O'Neil received his Bachelor's degree in Mechanical Engineering from Seattle University in 1972 and his Master's in Mechanical Engineering from the University of Southern California in 1977. The Massachusetts Institute of Technology (MIT) selected Mr. O'Neil as an Alfred P. Sloan Fellow and he received his Master's in Management from MIT in 1992.

Mr. George Nolfi

Mr. Nolfi is a member of the Secretary of the Navy's Advisory Panel. He received his BA from Princeton University's Woodrow Wilson School of Public and International Affairs with highest honors. He was awarded a Marshall Scholarship to study Political Philosophy at Oxford University and holds an M.A. from UCLA in quantitative Political Science. He is a writer, director, and producer of film and television. His work includes the films *The Bourne Ultimatum, The Adjustment Bureau, Ocean's Twelve, Birth of the Dragon* and the NBC television show, *Allegiance*. He has advised non-profit groups, companies, and other government agencies, on strategy and messaging.

Appendix C: Fact Finding Contributors

Contributor	Organization
VADM David Johnson, USN	Principal Military Deputy for the ASN for Research, Development and Acquisition
Dr. Lawrence Schuette	ONR Director of Research
RADM Girrier, USN	Director of Unmanned Warfare Systems (N99)
Dr. William Soper	Director of Strategic Capabilities Office, OSD
Dr. Brad Tousley	DARPA Director of the Tactical Technology Office (TTO)
BG Frank Kelley, USMC (ret)	Deputy Assistant Secretary of the Navy for Unmanned Systems
Mr. James Shields	Defense Science Board
Dr. Paul Nielsen	Defense Science Board
Dr. Marc Steinberg	Office of Naval Research, Science of Autonomy
Dr. Paul Schneider	Naval Studies Board
Dr. Jason Stack	Office of Naval Research, Ocean Engineering
RADMs David Antanitus and Nevin Carr, USN (ret)	Leidos Inc.
Dr. John Red-Horse	Sandia National Lab
VADM Grosklags, USN	Commander, Naval Air Systems Command
Mr. Steve Kracinovich	Director, Autonomous Systems Initiative, Naval Air Systems Command
Mr. Robert Floyd	Threat Systems, Naval Air Systems Command
Mr. Kerry Neace	The Johns Hopkins University Applied Physics Laboratory
Mr. Robert Finlayson	The Johns Hopkins University Applied Physics Laboratory
Mr. Patrick Buckley	Autonomous Systems, Naval Air Systems Command
Mr. John Kriz	Battlespace Simulation & Test, Naval Air Systems Command

Ms. Traci McCormick	Long Range Anti-Ship Missile, Naval Air Systems Command
Mr. Sam Lauber	The Johns Hopkins University Applied Physics Laboratory
Mr. Steve Olson	Autonomy Roadmap， Naval Air Systems Command
Mr. Matt Williams and Mr. Ricardo Ferro	Office of Naval Intelligence
Ms. Amy J. Markovich	Battle Simulation Office, Naval Air Systems Command
Colonel James T. Jenkins, USMC	Futures Directorate, Marine Corps Combat Development Command
Mr. Daniel P. Corbin	C4, Head Quarters Marine Corps
Mr. Charles Nemfakos	DoD Acquisition Process
RADM David J. Hahn, USN	Chief of Naval Research
Dr. Matt Daniels	Net Assessment, Office of the Secretary of Defense
Dr. John D. Burrow	Deputy Assistant Secretary of the Navy RDT&E
RADM Peter Fanta, USN	Director Warfare Integration (N91)
Dr. Ashish Kapoor	Microsoft Research
Dr.Andrey Kolobov	Microsoft Research
Dr. Debadeepta Dey	Microsoft Research
Dr. Sudipta Sinha	Microsoft Research
Mr. Jim Piavis	Microsoft Research
Mr. Dave Cuthbert	Amazon Web Services
Mr. Matt Wood	Amazon Web Services
Mr. Matt Lyman	Amazon Web Services
Mr. Tim Walden	Space-X
Dr. Gerald Mather, Dr. Jeff Dexter, Dr. John McNeal,	Space-X Engineers

and Mr. Andrew Howe	
Mr. Eddie Cabrera	Jet Propulsion Laboratory of the California Institute of Technology
Dr. Abbi Jain	Jet Propulsion Laboratory of the California Institute of Technology
Dr. Issa Nesnas	Jet Propulsion Laboratory of the California Institute of Technology
Mr. Rich Doyle	Jet Propulsion Laboratory of the California Institute of Technology
Dr. Terry Huntsberger	Jet Propulsion Laboratory of the California Institute of Technology
Dr. Mark Maimone	Jet Propulsion Laboratory of the California Institute of Technology
Dr. Mike Wolf	Jet Propulsion Laboratory of the California Institute of Technology
RADM Paul Sohl, USN	Commander, Operational Test and Evaluation Force
CAPT Paul Mitchell, USN	NAVAIR PMA 281
Mr. Matthew Clark	Verification and Validation Office, Air Force Research Laboratory
Dr. Marv Langston	Langston Associates
RADM Michael Manazir, USN	DCNO Warfare Systems (N9)
Mr. Don Christison	Long Range Anti-Ship Missile (LRASM), Naval Air Warfare Center Weapons Division, China Lake (NAWCWD)
Ms. Traci McCormick	Weapons Modeling, Simulation and Analysis Naval Air Warfare Center Weapons Division
Ms. Ashley Lemons	Weapons and Energetics, Naval Air Warfare Center Weapons Division
Mr. H. Wayne Willhite	Integration and Interoperability, Naval Air Warfare Center Weapons Division
Ms. Michelle Kilikauskas	Modeling and Simulation, Naval Air Warfare Center Weapons Division
Dr. Edward Berndt	Modeling and Simulation, Naval Air Warfare Center Weapons Division
Dr. Alan Van Nevel	Research Directorate, Naval Air Warfare Center Weapons Division
Dr. Gary Hewer	Image and Signal Processing, Naval Air Warfare Center Weapons Division
Dr. Katia Estabridis	Image and Signal Processing, Naval Air Warfare Center Weapons Division

Mr. Felipe Jauregui	Range Data Systems, Naval Air Warfare Center Weapons Division
Mr. James Walters	Evolved Sea Sparrow Missile (*ESSM*), Naval Air Warfare Center Weapons Division
Mr. Ken Hayes	Applied Manufacturing, Naval Air Warfare Center Weapons Division
Dr. Stephen Fallis	Chemistry Group, Naval Air Warfare Center Weapons Division
Dr. Gary Hewer, Dr. Katia Estabridis	Autonomy Demonstration, Naval Air Warfare Center Weapons Division
Dr. Jacob Dennis, Mr. Nick Quigley	Solid Fuel Ramjet, Naval Air Warfare Center Weapons Division
Mr. Scott Paulsen	F/A-18 and EA-18G, Naval Air Warfare Center Weapons Division
Mr. Mike Hefley	Unmanned Systems Integrated Programs, Naval Air Warfare Center Weapons Division
Ms. Rebecca Avitia	Threat Signal Processing In-The-Loop (T-SPIL), Naval Air Warfare Center Weapons Division
Dr. Lauren Merriman and Mr. Walter Merriman	Energetic Materials, Naval Air Warfare Center Weapons Division
Dr. Andrew Ihnen	Energetic Materials, Naval Air Warfare Center Weapons Division
Prof. John Leonard, PhD	Massachusetts Institute of Technology
Prof. Sangbae Kim, PhD	Biomimetic Robotics, Massachusetts Institute of Technology
Prof. James Kinsey, PhD	Humatics, Massachusetts Institute of Technology
Prof. Sertac Karaman, PhD	Aero/Astro, Massachusetts Institute of Technology
Various Individuals	Ocean Server Technology, Open Water Power, Autonomous Marine Systems in Greentown Labs, Somerville, MA
Ms. Margaret Palmieri	Digital Warfare Office, Navy Department
Ms. Jessica Shaffer	Annual Naval Technology Exercise, Naval Undersea Warfare Center, Newport
Mr. Jeremy Russell	UUV Family of Systems, Naval Undersea Warfare Center, Newport
Dr. Thomas Wettergren	Mathematical Methods for Autonomy Assessment, Naval Undersea Warfare Center, Newport

Mr. Michael Escobar	Weapons Analysis Facility Tour, Naval Undersea Warfare Center, Newport
Mr. Philip Campo and Mr. William Jones	Torpedo Software/VV&A, Naval Undersea Warfare Center, Newport
Mr. Stephen O'Grady	USW UxS Overview, Naval Undersea Warfare Center, Newport
Mr. Michael Pelczarski	Cross Domain Maritime Surveillance and Targeting, Naval Undersea Warfare Center, Newport
Mr. Steven Aguiar	Virtual Worlds, Naval Undersea Warfare Center, Newport
Mr. Peter Harrigan	Submarine UAS, Naval Undersea Warfare Center, Newport
Mr. Michael Incze	IVER UUV Operations, Naval Undersea Warfare Center, Newport
Mr. Thomas Merchant	UUV Architectures and Autonomy, Naval Undersea Warfare Center, Newport
Mr. Scott Houde	C3, Naval Undersea Warfare Center
Mr. Bryan Johnson	MK18 Explosive Ordnance Disposal Projects, Naval Undersea Warfare Center, Newport
Mr. Andrew Bouchard	Autonomy Verification and Validation, Naval Undersea Warfare Center, Newport
Mr. Jeffrey Witt	Architectures, Naval Undersea Warfare Center, Newport, Newport
Mr. Christopher Egan	Large Displacement UUV Architecture, Naval Undersea Warfare Center, Newport
Dr. Farnam Jahanian	Provost, Carnegie Mellon University
Dr. Andrew Moore	Computer Science, CMU
Dr. Bill Scherlis	Institute for Software Research, CMU
Dr. H. Herman	National Robotics and Engineering Center, CMU
Dr. Paul D. Nielsen	Software Engineering Institute (SEI), CMU
Dr. Phil Koopman	Department of Electrical and Computer Engineering, CMU
Mr. Nathan Flinn	Common Ground Control Station (GCS), Ship Combat System (SCS) Interface, NAVSEA Weapons Center, Dahlgren

Mr. Bob Gripshover	ASUW Mission Package, NAVSEA Weapons Center, Dahlgren
Mr. Ian Shafer	UAV Weaponization, NAVSEA Weapons Center, Dahlgren
Mr. Harry Dreany	USMC Expeditionary UxS, NAVSEA Weapons Center, Dahlgren
Mr. Travor Sutton	Battle Management System, NAVSEA Weapons Center, Dahlgren

Appendix D: Acronyms

AI	Artificial Intelligence
AUV	Autonomous Undersea Vehicle
C3-I	Command, Control, Communications and Intelligence
CFIT	Controlled Flight Into Terrain
CG MCCDC	Commanding General, Marine Corps Combat Development Command
COI	Community of Interest
CNO	Chief of Naval Operations
CRADA	Cooperative Research & Development Agreements
DARPA	Defense Advanced Research Projects Agency
DIUX	Defense Innovation Unit Experimental
DoD	Department of Defense
DoN	Department of the Navy
EDW	Electronic Data Warehouse
EHR	Electronic Health Record
FACA	Federal Advisory Committee Act
G-CAS	Ground Collision Avoidance System
GPU	Graphics Processing Unit
IED	Improvised Explosive Device
M&S	Modeling and Simulation
MOS	Military Occupational Specialty
NRAC	Naval Research Advisory Committee
NRL	Naval Research Laboratory
OODA (loop)	Observe, Orient, Decide, and Act
OTA	Other Transaction Authority
PACFLT	Pacific Fleet
R&D	Research and Development
ROTC	Reserve Officer Training Corps
SAM	Surface to Air Missile
SEAL	Sea Air and Land
SECNAV	Secretary of the Navy
SNAP	SECNAV Advisory Panel
T&E	Test and Evaluation
TTPs	Tactics, Techniques and Procedures
UUV	Unmanned Undersea Vehicle
UxS	Unmanned Systems
VV&A	Verification, Validation, and Accreditation

XLUUV	Extra Large Unmanned Undersea Vehicle

Appendix E: Glossary of Terms

Artificial Intelligence	Artificial intelligence is behavior by a machine when it mimics "cognitive" functions that humans associate with other human minds, such as "learning" and "problem solving".
Automation	Automation can be defined as the technology by which a process or procedure is performed without human assistance.
Autonomy	Autonomy or Autonomous behavior is – at best – imprecise terminology to express the concept of a non-human device acting on its own. It can range from simple pre-programmed activities to open-ended human out-of-the loop actions of AI learning systems.
Data Curation	Data Curation is a broad term used to indicate processes and activities related to the organization and integration of data collected from various sources, annotation of the data, and publication and presentation of the data such that the value of the data is maintained over time, and the data remains available for reuse and preservation.
Deep Learning	Deep learning is a class of machine learning algorithms that use multiple layers of nonlinear processing units for feature extraction and transformation. Each successive layer uses the output from the previous layer as input. A deep learning algorithm learns in supervised (e.g., classification) and/or unsupervised (e.g., pattern analysis) manners.
Machine Learning	Machine Learning is a field of computer science that enables software the ability to learn without being explicitly programmed.
Metadata	Metadata is a set of data that describes and gives information about other data.
Neural Nets (Networks)	Biological neural nets – like those found in humans enable perception through sensing one's environment (e.g.; sight, hearing, etc.). "Artificial" neural nets are computing systems, inspired by biological neural networks. Artificial neural nets perceive and learn (i.e.; progressively improve performance) to do tasks by considering examples, without task-specific programming.