Advancing Autonomous Systems

An Analysis of Current and Future Technology for Unmanned Maritime Vehicles

Bradley Martin, Danielle C. Tarraf, Thomas C. Whitmore,
Jacob DeWeese, Cedric Kenney, Jon Schmid, Paul DeLuca

For more information on this publication, visit www.rand.org/t/RR2751

Published by the RAND Corporation, Santa Monica, Calif.
© Copyright 2019 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This document and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited. Permission is given to duplicate this document for personal use only, as long as it is unaltered and complete. Permission is required from RAND to reproduce, or reuse in another form, any of its research documents for commercial use. For information on reprint and linking permissions, please visit www.rand.org/pubs/permissions.

The RAND Corporation is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Support RAND
Make a tax-deductible charitable contribution at
www.rand.org/giving/contribute

www.rand.org

Preface

The U.S. Navy is interested in developing autonomous capabilities to execute tasks that are increasingly hazardous for humans and to enhance warfighting capabilities. This report focuses on two Navy platform classes—unmanned undersea vehicles and unmanned surface vehicles—and explores the potential for increasing the numbers and capabilities of autonomous Navy systems. The report examines both the technological development of such systems and the warfighting requirements of the Navy. The authors analyze the following four areas: the current state of the art of autonomous technology, current kill chains and capabilities, future fleet architecture and its autonomous capabilities, and autonomy in alternative concepts of operation.

This research was sponsored by the Office of Naval Research and conducted within the Acquisition and Technology Policy Center of the RAND National Defense Research Institute, a federally funded research and development center sponsored by the Office of the Secretary of Defense, the Joint Staff, the Unified Combatant Commands, the Navy, the Marine Corps, the defense agencies, and the defense Intelligence Community.

For more information on the RAND Acquisition and Technology Policy Center, see www.rand.org/nsrd/ndri/centers/atp or contact the director (contact information is provided on the webpage).

Contents

CHAPTER FIVE

CHAPTER SIX

APPENDIXES

Figures and Tables

Figures

Tables

Summary

Increases in computing power have enabled machines to do more of the repetitive and demanding work that humans used to do. These work duties range from simple automated tasks to more-sophisticated applications of artificial intelligence, up to the point of self-driving vehicles and farm equipment. These advances in artificial intelligence and robotics are very likely to continue and result in even wider use of machines, especially for repetitive tasks in hazardous environments, a type of environment that is very familiar to people serving in the military. Indeed, the U.S. military is increasingly employing autonomous and semi-autonomous machines either to reduce the risk to humans or to enhance military capabilities, and the U.S. Navy is no exception.

Although autonomous technologies offer benefits in a wide range of military tasks, this report focuses on platform autonomy—specifically, unmanned undersea vehicles (UUVs) and unmanned surface vehicles (USVs) (that is, vehicles that operate on the surface of the water). Several reasons underpin this choice. First, the most useful applications of autonomy are likely to be in environments where humans do not have a reasonable chance of surviving. Thus, UUVs and USVs are excellent candidates for analysis of how and where such systems might be employed most effectively. Second, unmanned aerial vehicles have similar characteristics, and the findings from UUVs and USVs may be transferable. However, much of the development effort for unmanned aerial vehicles is taking place outside the purview of the Office of Naval Research, which sponsored this study. Per guidance from that office, we focused on UUVs and USVs.

In this report, we seek, first, to establish the state of the art in the development of autonomous systems and map how these might be effective in advancing warfighting as it exists and is currently projected. Second, and more broadly, we evaluate how advances in autonomy might change the nature of several key elements of naval warfare and to ascertain how such changes might affect investment and development decisions.

To accomplish these tasks, we proceed along two paths: one looking at technological development and one looking at warfighting requirements. We overlay our findings to develop a gap analysis with mission and projected technology, then we further examine concepts that go beyond existing warfighting concepts of operation (CONOPs) potentially to exploit features of autonomy that current CONOPs do not

take into account. We chose four areas of analysis: the current state of the art of autonomous technology, current kill chains (the end-to-end means for achieving a warfighting effect) and capabilities, future fleet architecture and its autonomous capabilities, and autonomy in alternative CONOPs.

What We Conclude

Our conclusions come in three broad categories: the state of the art of autonomous technology, applications of autonomy in current missions, and possible future uses and constraints on autonomy.

We conclude the following about autonomous technology's current state of the art:

- Advances in autonomy have been steady, but the transition to systems capable of reacting to unexpected changes in the environment has not occurred and might not occur for a decade or more.
- Military applications contemplated for unmanned vehicles are unlikely to materialize without substantial targeted investment and development. Reliance on commercial off-the-shelf technology is not likely to support complex military missions.
- Although recent advances in machine learning are promising, enthusiasm for the use of machine learning in military applications should be tempered by the need for embedded computation and retraining, as well as such systems' potential vulnerability to countermeasures.

We conclude the following about applications of autonomy in current missions:

- The major limitations associated with the most complex missions are not necessarily associated with autonomy, and it might not be useful to accelerate autonomy while such issues as power generation and storage are still being worked out.
- Under current kill chains and CONOPs, autonomy is generally employed to replicate items in the kill chain exactly as they are carried out by manned systems.

We conclude the following about potential applications and constraints on autonomy:

- Future fleet architecture and possible uses of autonomy do not align well, and our analysis suggests that some features desired for the future fleet are unlikely to be reached under the known program of record.
- Policy issues related to autonomous systems applying rules of engagement do exist, and it is unlikely that these systems will be engineered in a way that avoids these issues. Slowing the decisionmaking of autonomous systems by interpos-

ing a human in the loop will likely cede a critical time advantage in a high-intensity environment. This delay is a choice and cannot be mitigated by technical improvements.

- Although development efforts focus on multifunctional, highly complex systems, some of the more promising uses of autonomy might be in using simple systems with limited autonomy for most functions—but adding the capacity for the systems to coordinate with each other. This is particularly the case for missions (such as mine countermeasures) in which the current extended timelines are related to the need to have one platform go through the full detect, classify, identify, and engage sequence. Having a large number of cooperating single-sensor platforms could significantly speed up the timeline.

What We Recommend

Using the findings from our analyses in the four areas we studied, we recommend that the U.S. Navy do the following:

- Revisit assumptions concerning technological progress in autonomy. Our research indicates that the capability for autonomous systems to interpret context and make independent decisions, particularly in a dynamic environment, is not realistic in the short term.
- In systems requiring high degrees of autonomy, align the development of autonomy with the development of other capabilities, such as power generation and storage, that might be limiting factors.
- Use the unique features of autonomy to enable new CONOPs. It seems particularly promising to employ simple but numerous systems carrying different kinds of sensors. This capability will require development of a system capable of fusing multiple inputs to create a common operational depiction of the battlespace as a reference for the individual operating systems.
- Reevaluate future fleet design architecture requirements in light of the state of the art of autonomous technology. Some features of the proposed architectures are more aspirational than likely.
- Accept the reality that autonomous systems will need to make engagement decisions if those decisions are to be effective. Modern weapon system timelines simply preclude human intervention.
- Develop a mechanism that allows humans to periodically assess whether an autonomous system is misinterpreting its environment.
- Critically evaluate the viability of complex multimission platforms, and consider emphasizing simple but cooperating platforms.

Acknowledgments

We gratefully acknowledge the guidance of Eric Gulovsen and the Disruptive Technology team at the Office of Naval Research, which sponsored this study. We also would like to recognize several RAND Corporation staff for their contribution to this work. In particular, we greatly appreciate the thorough and insightful reviews provided by Osonde Osoba and Angela Putney. Furthermore, we appreciate the administrative support of Sunny Bhatt, Maria Falvo, and Leslie Thornton, as well as the leadership and oversight of Cynthia Cook, Joel Predd, and Christopher Mouton. Finally, we are grateful for the editorial assistance of Jerry Sollinger and Allison Kerns.

Abbreviations

3D	three-dimensional
A2AD	anti-access/area denial
ACTUV	Anti-Submarine Warfare Continuous Trail Unmanned Vessel
ASW	antisubmarine warfare
CONOP	concept of operation
CNS	celestial navigation system
CSG	carrier strike group
DARPA	Defense Advanced Research Projects Agency
DoD	U.S. Department of Defense
DSB	Defense Science Board
DVL	doppler velocity log
FFDA	future fleet design architecture
GPS	Global Positioning System
IEEE	Institute of Electrical and Electronics Engineers
INS	inertial navigation system
IPB	intelligence preparation of the battlefield
IPC	International Patent Classification
ISR	intelligence, surveillance, and reconnaissance
LDUUV	large displacement unmanned undersea vehicle
MCM	mine countermeasures
PV	photovoltaic
S&T	science and technology
SLAM	simultaneous localization and mapping
UAV	unmanned aerial vehicle
UGV	unmanned ground vehicle

USV	unmanned surface vehicle
UUV	unmanned undersea vehicle
V&V	verification and validation

CHAPTER ONE

Introduction

The technology that enables machines to carry out functions that have traditionally been done by humans has continued to evolve and develop. Not all of this technology is associated with military applications—indeed, much of it is not. But it is important to understand how these developments might affect the way the U.S. military—and, for purposes of this study, the U.S. Navy in particular—carries out operations.

The autonomous technology that is developed by commercial or private entities can have far-reaching effects on the economy, transportation, education, and other spheres. But the impact of autonomous technology for military use extends well beyond those applications. For example, the military might—and probably should—invest in the development of systems that better survive and counteract advances in weapon lethality and capability that have made operations in some denied environments untenable. Improved autonomy might be key to gaining a warfighting advantage. Developing such systems transcends a desire to minimize human casualties and extends into the need to operate in an environment that is so physically stressing and hostile that humans cannot operate in it, even assuming that the decisionmakers have a tolerance for human casualties. That is, the point of advanced autonomous technology for the military is not solely to avoid casualties; it is to enable operations in environments where manned systems simply could not operate.

As we discuss in this report, interpretations of *autonomy* differ, as do many applications. For purposes of this study, we are focusing mostly on autonomy associated with unmanned undersea vehicles (UUVs) and unmanned surface vehicles (USVs) (that is, vehicles that operate on the surface of the water). After consulting with the sponsor of this study (the Office of Naval Research), we elected to concentrate on these particular systems for two key reasons. First, although high levels of autonomy might be built into manned systems for such functions as processing information and helping make tactical decisions, this type of improvement in autonomy is largely a matter of enhancing human decisionmaking and thus requires a person somewhere on the platform. But autonomy might be most effective in environments where humans do not have a reasonable chance of surviving. Capitalizing on that potential will require the development of systems capable of doing many of the same interpretation and discrimination functions that humans perform but in an environment that is unsafe for humans. And because UUVs and USVs operate in and above water, respectively, they are excellent candidates for analyzing such capabilities. Second, while the findings of our analysis are likely transferable to unmanned aerial vehicles (UAVs), which have similar characteristics to the unmanned systems we examined, the development of UAVs is taking place outside the purview of the study sponsor. Thus, for this study, we focus on UUVs and USVs.

Purpose

Our overall purpose is, first, to establish the state of the art in autonomous vehicle technology and map how such capabilities might be effective in advancing warfighting as it exists and is currently projected. The U.S. Navy has some very ambitious plans for integrating autonomous systems into its future fleet, and it is not obvious that all these elements are on a path for delivery. A related objective is to describe the types of investment and development decisions that might need to be made to support these changes. It may be that industry progress will be so complete that little additional development will be required by the military other than to adapt the capabilities to military use. Or it may be that significant military investment will be required from inception to delivery. Finally, we consider how improved autonomy might change the practice of warfare in unexpected ways.

Approach and Report Organization

In this report, we explore two separate paths—one looking at technological development and one looking at requirements. We overlay our findings to develop a gap analysis with mission and projected technology, then we further examine concepts that go beyond existing warfighting concepts of operation (CONOPs) to potentially exploit features of autonomy that current CONOPs do not recognize. We examine the following four areas of analysis: the current state of the art of autonomous technology, current kill chains and capabilities, future fleet architecture and its autonomous capabilities, and autonomy in alternative CONOPs. Each of these is considered in a separate chapter.

The State of the Art

We begin with an exhaustive review to establish and assess the current state of the art of autonomous technology. Our review covers the algorithms that enable autonomy, as well as the UUV and USV platforms and payloads currently in service or development. Our review of the algorithms is based on an extensive survey of the published technical and academic literature. Our review of platforms and payloads includes civilian systems, civilian systems already adapted for military use, wholly military systems already in service, and military systems under an established program of record. The main findings of our review and analysis are summarized in Chapter Two. Some of the details of our deep dives into algorithms, platforms, and payloads are relegated to the appendixes. Although we strive to identify both solved problems and remaining challenges from a technical viewpoint with an eye toward the naval operating environment, we make no presumption about the effectiveness or suitability of the available technologies for various missions at this point. Chapter Two and its supporting appendixes thus primarily serve as an introductory exposition on vehicle autonomy and its enablers for our readers, laying a common ground for subsequently exploring vehicle autonomy for Navy missions.

Current Kill Chains and Capabilities

Kill chain refers to the end-to-end sequence of events required to achieve a warfighting effect. In Chapter Three, we review the military missions whose kill chains currently contain or envision a role for unmanned autonomous systems. For purposes of this project, we consider three missions that have been proposed, discussed, or included in war games:

- mine countermeasures (MCM) against mines laid by an adversary in contested waters
- penetration into denied undersea space for covert intelligence, surveillance, and reconnaissance (ISR) or offensive mining
- use of USVs as decoys, electronic warfare assets, and potentially independent surface action groups.

Initially, for these missions, we use CONOPs that the Navy already uses for its manned systems and has, to a degree, adapted for its unmanned systems. In this chapter, we simply demonstrate how closely capability development matches the needs of the CONOPs the Navy has established. Where investment gaps occur, we note them.

Future Fleet Architecture and Its Autonomous Capabilities

The National Defense Authorization Act for Fiscal Year 2017 directed the U.S. Department of Defense (DoD) to develop three separate fleet architectures that would help illuminate the requirements and interrelationships among capabilities for Navy fleet operations in the future. These architectures were laid out in three documents mandated by Congress: one by the Navy, one by a private research institute, and one by a federally funded research and development center.[1] In Chapter Four, we compare the

---

[1] Megan Eckstein and Sam LaGrone, "Trio of Studies Predict the U.S. Navy Fleet of 2030," *USNI News*, February 14, 2017.

capability needs of the proposed future fleet design architectures (FFDAs) with the development and likely state of autonomous technology in the future. Because the FFDAs are projected so far into the future (2045), the kill chains required are, in many cases, still undefined. However, some of the capability development needs that will be required are clear, and we focus our analysis there.

Autonomy in Alternative Concepts of Operation

Our fourth area of analysis, discussed in Chapter Five, evaluates how missions might be better accomplished using unmanned autonomous systems but with different CONOPs from those envisioned in either current kill chains or the proposed FFDAs. This is an effort to examine the ways in which likely technology might enable a different use or approach from the ones used today, which, among other things, seek to minimize risk to manned platforms.

Recommendations and Conclusions

Using the findings and conclusions from our analyses, we make recommendations in four areas:

- the viability of proposed uses of autonomy based on the current state of technology
- the short-term development and investment requirements to overcome gaps in existing CONOPs
- the longer-term development and investment requirements for the future fleet
- the alternative CONOPs for the use of autonomous technology.

CHAPTER TWO

The State of the Art of Autonomous Technology

This chapter and the appendixes to the report provide a comprehensive picture of the relevant state of the art of autonomous technology, focusing particularly on autonomous vehicle technologies for use underwater and on the surface of the water.[1] Although autonomous systems are increasingly present in people's lives, agreement on what constitutes *autonomy* remains elusive. We thus begin this chapter with a brief survey of the use of the term among various sources. We then introduce our analytic approach, focusing on algorithms as the central component of autonomy, with payloads and platforms as enablers. This approach sets the stage for the deeper dives and findings that follow for each of these three components. After we explore the individual building blocks of autonomy, we end the chapter with an integrative assessment that highlights promising technological advances, as well as remaining technological gaps, with an eye toward the unique requirements of the naval operating environment.

What Is Autonomy?

Autonomy and unmanned systems is one of the nine focus areas in the Office of Naval Research's Naval Science and Technology (S&T) Strategy, which has a vision of achieving an integrated manned and unmanned force with the "ability to sense, comprehend, predict, communicate, plan, make decisions and take collaborative action to achieve operational goals."[2] Although *autonomy* is not explicitly defined in this context, its goals clearly are. These goals will guide our investigation into the current state of the art of autonomous underwater and surface vehicles.

A Defense Science Board (DSB) task force defines *autonomy* as "a capability (or a set of capabilities) that enables a particular action of a system to be automatic or, within

---

[1] In this chapter, we present a comprehensive but high-level overview of the state of the art, geared toward a general audience. We relegate most of the technical details, of particular interest to technical audiences, to the appendixes. Appendix D includes an analysis of patent data to supplement our study.

[2] Office of Naval Research, *Naval S&T Strategy: Innovations for the Future Force*, Washington, D.C., 2015, p. 28.

programmed boundaries, 'self-governing.'"[3] A DSB report on technology and innovation enablers recommends research and development investments in long-endurance, autonomous, networked UUVs to achieve superiority through cost-imposing strategies, although the authors do not precisely define the "mission-level cooperative autonomy" that they advocate.[4] According to the 2016 DSB *Summer Study on Autonomy*, "Autonomy results from delegation of a decision to an authorized entity to take actions within specified boundaries."[5] The study focused on autonomous *capabilities* rather than autonomous *systems* because, as noted in the report, no machine is truly autonomous in the strict sense of the word. A DoD community of interest on autonomy defines the term as the "computational capability for intelligent behavior that can perform complex missions in challenging environments with greatly reduced need for human intervention, while promoting effective man–machine interaction."[6] This definition highlights the role of software and computation. The autonomy community of interest's Test and Evaluation, Verification, and Validation Working Group further differentiates between automation and autonomy.[7] This limited survey of how various groups define autonomy highlights the difficulties inherent in precisely defining the term and the ongoing debates about what that definition should be.

Similarly, multiple attempts have been made to define the *levels* of autonomy. Some of the definitions aim simply to categorize a given system into one of several levels on a scale covering the spectrum from fully human operation to fully autonomous operation, although the number and nature of these levels vary widely. For instance, the autonomy scale created for the Uninhabited Combat Air Vehicle program consisted of four levels, and that created for the Army's now-defunct Future Combat

---

[3] DSB, *Task Force Report: The Role of Autonomy in DoD Systems*, Washington, D.C.: U.S. Department of Defense, July 2012, p. 1.

[4] DSB, *Technology and Innovation Enablers for Superiority in 2030*, Washington, D.C.: U.S. Department of Defense, October 2013, p. 39.

[5] DSB, *Summer Study on Autonomy*, Washington, D.C.: U.S. Department of Defense, June 2016, p. 4.

[6] Kris Kearns, "DoD Autonomy Roadmap," briefing slides, National Defense Industrial Association 19th Annual Science and Engineering Technology Conference, March 21, 2018. Reliance 21, the overarching framework of the DoD S&T joint planning and coordination process, is led by the S&T Executive Committee and underpinned by an ecosystem of 17 technical groups known as communities of interest. See DoD, *Reliance 21: Operating Principles*, Washington, D.C., January 2014.

[7] The working group employs the following definitions: *Automation* corresponds to the system functioning with no or little human operator involvement, but the system performance is limited to the specific actions it has been designed to do (i.e., through the use of simple rule-based responses). *Autonomy* corresponds to the system having a set of intelligence-based capabilities that allows it to respond to situations that were not preprogrammed or anticipated prior to system deployment (i.e., decision-based responses). See Test and Evaluation, Verification, and Validation Working Group, *Technology Investment Strategy 2015–2018*, Washington, D.C.: U.S. Department of Defense, Autonomy Community of Interest, May 2015, p. 2.

Systems program consisted of ten levels.[8] A 2013 RAND study for the Navy employed a scale with seven levels.[9]

Although such scales are relatively straightforward to define, it is not difficult to foresee their potential limitations. In particular, a system may have multiple components carrying out different tasks with varying degrees of human intervention. Does it make sense then to talk about the level of autonomy of the entire system? Thomas Sheridan, while advocating an autonomy scale consisting of ten levels, additionally proposed evaluating the autonomy of a system on four dimensions (information acquisition, information analysis, decision selection, and action implementation).[10] The Committee on Autonomous Vehicles in Support of Naval Operations proposed looking at levels of *mission autonomy*, incorporating two degrees of freedom (mission complexity and degree of automation).[11] In addition, the National Institute of Standards and Technology developed its generic Autonomy Levels for Unmanned Systems framework encompassing three degrees of freedom (human independence, mission complexity, and environmental difficulty).[12] Finally, a DSB task force on the role of autonomy in DoD systems recommended that "DoD should abandon the debate over definitions of levels of autonomy and embrace a three-facet (cognitive echelon, mission timelines, human-machine system trade spaces) autonomous systems framework."[13]

Rather than enter these debates by proposing our own definitions of autonomy and the levels of autonomy or by adopting existing definitions, we opted to focus our efforts on understanding the vast array of technological advances that enable engineers to remove humans from the loop, with an emphasis on vehicles, particularly underwater and surface ones. Understanding these technological advances, including their current limitations, enables us to have a fruitful discussion about how they might be developed and applied in Navy missions.

---

[8] Committee on Autonomous Vehicles in Support of Naval Operations, *Autonomous Vehicles in Support of Naval Operations*, Washington, D.C.: National Academies Press, 2005.

[9] Scott Savitz, Irv Blickstein, Peter Buryk, Robert W. Button, Paul DeLuca, James Dryden, Jason Mastbaum, Jan Osburg, Phillip Padilla, Amy Potter, Carter C. Price, Lloyd Thrall, Susan K. Woodward, Roland J. Yardley, and John M. Yurchak, *U.S. Navy Employment Options for Unmanned Surface Vehicles (USVs)*, Santa Monica, Calif.: RAND Corporation, RR-384-NAVY, 2013.

[10] The four dimensions are inspired by a simple, four-stage, sequential model of human information processing, consisting of sensory processing, perception and working memory, decisionmaking, and response selection. The system may thus exhibit a different level of autonomy along each of its four dimensions. See Thomas B. Sheridan, *Humans and Automation: System Design and Research Issues*, Hoboken, N.J.: John Wiley, 2002.

[11] Committee on Autonomous Vehicles in Support of Naval Operations, 2005.

[12] Hui-Min Huang, ed., *Autonomy Levels for Unmanned Systems (ALFUS) Framework*, Vol. 1: *Terminology*, Gaithersburg, Md.: National Institute of Standards and Technology, Special Publication 1011-I-2.0, October 2008.

[13] DSB, 2012, p. 2.

Analytic Approach

Our goals for this report are to assess and convey the following:

1. what technologies are being developed to enable autonomy in naval vehicles
2. what capabilities currently are and are not possible
3. where there are challenges, why these challenges exist, and why they are difficult to overcome.

To achieve these goals, we carried out an extensive data-collection, analysis, and synthesis effort covering Navy and industrial investment areas in science, technology, research, and development. We also conducted extensive searches into the relevant technical and academic literature.[14] We structured our efforts around the three main building blocks of autonomy: algorithms, payloads, and platforms.

Algorithms can be viewed as the "brains" behind autonomous capabilities. They are thus the central and most critical building blocks in the quest to remove humans from the loop. They are also the bridge by which we can most naturally connect autonomous capabilities to vehicle tasks (and, eventually, missions, as we discuss in later chapters). Payloads, on the other hand, are the "peripherals" that enable vehicles to interact and interface with the world around them. They serve as important enablers of autonomy, although they are also used to augment humans when humans remain in the loop. Finally, platforms are the "chassis" of the vehicles, where the algorithms and payloads come together. Depending on their design and characteristics, platforms may either hinder or enable the quest toward autonomy, as we discuss later in this chapter.

In the next sections, we summarize our analysis and assessment of the state of the art of each of these three building blocks of autonomy. To maintain a streamlined exposition here, we relegate the details of our analysis of each building block to Appendixes A, B, and C.

Algorithms

A Historical Perspective

The quest for a systematic, principled study of autonomy—and, in particular, machine thinking or intelligence—dates back several decades.[15] The current toolbox of algorithms traces its roots to two primary thought communities that arose in the middle of the 20th century: the *cybernetics* community and the *artificial intelligence* community.

---

[14] We supplemented these with an analysis of recent patent data, as described in Appendix D.

[15] The quest for building autonomy goes back to antiquity: The first engineered feedback control system is widely thought to be the water clock, variously attributed to the ancient Greeks, Arabs, Indians, Romans, and Persians.

In his 1948 book, Norbert Wiener used the term *cybernetics* to refer to self-regulating mechanisms.[16] The definition posited by Andrey Kolmogorov, "science concerned with the study of systems of any nature which are capable of receiving, storing and processing information so as to use it for control,"[17] is a crisp description of the interdisciplinary field envisioned by Wiener, bringing together various branches of engineering and mathematics. Although the use of the term *cybernetics* has since evolved very far from its roots,[18] several rigorous academic disciplines grounded in mathematics and physics arose out of the original cybernetics thought community; these disciplines include information theory, communication theory, and control theory. The third field, of particular interest for our study, is significantly represented in our review and assessment of algorithms that enable autonomy.

John McCarthy coined the term *artificial intelligence* in 1955 when he organized the Dartmouth Summer Research Project on Artificial Intelligence that effectively launched the field. The summer project's goal was to bring together a group of eminent researchers to brainstorm about thinking machines. As noted in the project's proposal,

> The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves.[19]

Since the first book on artificial intelligence was published in 1963,[20] the concept has courted its fair share of controversy, owing partly to its quick rise to popularity and its ambitious claims, including the imminent arrival of artificial general intelligence possessing the flexibility and ingenuity of a human brain—claims that remain largely unfulfilled. Nonetheless, machine learning has emerged as a serious technical field grounded in computer science and statistics. It is of particular interest for our study and is significantly represented in our review and assessment of algorithms.

---

[16] Norbert Wiener, *Cybernetics or Control and Communication in the Animal and the Machine*, Cambridge, Mass.: MIT Press, 1961. Wiener's book was met with significant public interest and effectively introduced the term *cybernetics* into public discourse. In 1954, Marie Neurath produced a children's book that introduced one of the objects of study in cybernetics, analog feedback control, in an accessible format (Marie Neurath, *Machines Which Seem to Think*, London: Parrish, 1954).

[17] Quoted in Stuart Umpleby, "Definitions of Cybernetics," *Larry Richards Reader 1997–2007*, 1998.

[18] See American Society for Cybernetics, homepage, 2016.

[19] John McCarthy, Marvin Minsky, Claude Elwood Shannon, and Nathaniel Rochester, *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, Hanover, N.H.: Dartmouth College, 1955.

[20] Edward A. Feigenbaum and Julian Feldman, eds., *Computers and Thought*, London: McGraw-Hill, 1963. This book consists of 20 chapters, of which six had been previously published as RAND reports. See Philip Klahr and Donald A. Waterman, *Artificial Intelligence: A RAND Perspective*, Santa Monica, Calif.: RAND Corporation, P-7172, January 1986.

Control theory and machine learning are thus complementary fields underpinning autonomy and are quite different in spirit, perhaps owing to their cultural roots in different thought communities.[21] Indeed, viewed from a high level, control theory relies heavily on physically justifiable models of systems and the environment in which they operate, while machine learning bypasses the explicit use of physical models by focusing on large-scale data for pattern matching and inference. And, in an ironic twist,[22] public perception has come to lump both fields under the broad banner of artificial intelligence.

Analytic Approach

While the academic exercise of organizing the algorithms underpinning autonomy on the basis of their inherent characteristics (e.g., model-based, data-driven) and cultural roots (e.g., control theory, machine learning) is intriguing from an intellectual point of view, the reality is that developing and fielding autonomous naval vehicles requires solving a wide range of engineering problems enabling these vehicles to carry out various tasks without human input or intervention.

For our study, we thus opted to pursue a task-centric analytic approach in which we developed a taxonomy to organize tasks into families and subfamilies, covering both single vehicles and teams of unmanned vehicles. This task-centric taxonomy contrasts with capability-centric taxonomies, such as the OODA (observe, orient, decide, and act) loop of decisionmaking, to organize technologies critical to the development of autonomous systems. A useful feature of our taxonomy is that the tasks provide a natural way of relating the algorithms to operations and missions. In addition, our study was scoped to focus exclusively on vehicle autonomy and thus stops short of looking into human-machine teaming.

The proposed taxonomy is shown in Table 2.1 for a single unmanned vehicle (also referred to in this report as an *agent*). The *navigation* task family encompasses algorithms enabling an unmanned vehicle to accurately pinpoint its position; move from a point of origin to a destination; do so safely by avoiding obstacles and collisions; and maintain its absolute position or its station relative to a point of reference, possibly another vehicle, in the face of ocean currents. The *exploration* task family encompasses algorithms that enable an autonomous vehicle to monitor its environment to detect, identify, and track objects of interest; detect changes; explore its surroundings, gen-

---

[21] In reality, it would be more accurate to talk about classes of approaches than about complementary fields because both control theory and machine learning are vast fields that encompass a large range of algorithmic and mathematical developments. An in-depth historical perspective can be found in Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge, UK: Cambridge University Press, 2009.

[22] It is generally thought that McCarthy coined the term *artificial intelligence* to differentiate his vision for machine intelligence from that of the more established Wiener, thereby avoiding direct competition. It is thus ironic that Wiener's intellectual agenda has come to figure significantly into, if not dominate, McCarthy's marketing terminology.

Table 2.1
Taxonomy for a Single Unmanned Vehicle

Task Family	Subtasks
Navigation	• Localization • Path-planning • Collision avoidance • Reference-tracking and path-following • Station-keeping
Exploration	• Object recognition • Detection and cuing • Tracking • Simultaneous localization and mapping (SLAM) • Oceanography
Effect delivery	• Kinetic strike • Nonkinetic strike • Manipulation
Countermeasures	• Evasion • Fault tolerance • Cyber resilience • Deception
Resource management	• Power management

erate a map of them, and place itself on that map; and explore its surroundings for research and study purposes (oceanography). The demarcation between the navigation and exploration task families is not sharp; for example, processing of sensor signals and filtering and estimation using sensor data are needed for both navigation and exploration. The *effect delivery* task family includes algorithms allowing the unmanned vehicle to deliver effects, both kinetically and nonkinetically, and manipulate or otherwise affect its environment and the objects or systems in it. The *countermeasures* task family incorporates algorithms to evade threats, both kinetic and nonkinetic; mitigate against and recover from degradation of physical or software parts, as a result of either system failures or targeted attacks; and deceive adversaries. Finally, the *resource management* task family encompasses algorithms for managing onboard resources, primarily power.

The corresponding taxonomy for teams of unmanned vehicles is shown in Table 2.2. In contrast to the single-vehicle taxonomy, our proposed taxonomy for teams of vehicles consists of only two families: *swarming*, which is used here to refer to coordinated navigation and movement of the vehicles, and *cooperation*, which is used here to refer to other endeavors that a team of vehicles might jointly undertake.

Table 2.2
Taxonomy for Teams of Unmanned Vehicles

Task Family	Subtasks
Swarming	• Flocking and synchronization • Rendezvous • Formation control
Cooperation	• Area surveillance and distributed sensing • Cooperative search and tracking • Multi-agent SLAM • Coordinated deception

There is one remaining class of algorithmic developments that are critical for effective development and use of autonomous vehicles, whether individually or in teams, that is not explicitly captured in our task-centric taxonomy but that is nonetheless covered in our algorithm survey and assessment. This class of developments relates to *verification and validation* (V&V) of autonomous systems.

Having developed our taxonomy of tasks, we carried out an extensive search of the academic literature to identify and assess classes of algorithms available for each task family and subfamily, as well as for the V&V problem. Our search included control theory and machine-learning literature, among others. Our main findings, including general trends and remaining gaps, are presented next, while the details of the literature survey can be found in Appendix A.

General Trends and Remaining Gaps

Although navigation, in all its facets, has seen much recent progress, the algorithms remain generally computationally intensive. Moreover, they often require tailoring for real-time performance. Finally, they often rely on the Global Positioning System (GPS), which is not always available for underwater navigation.

Machine-learning algorithms are natural candidates for exploration tasks—for instance, object classification based on image or other sensor data. However, the recent successes are mostly due to supervised-learning algorithms that require training on large, labeled data sets; that require significant retraining to maintain performance; and that run on powerful machines in data centers (known as cloud computing). The availability of labeled data sets, sufficient opportunities to retrain in theater, and the extent to which these algorithms can be implemented with embedded computation on board the vehicle may all pose challenges for naval applications. Finally, classifier algorithms have been shown to be vulnerable to *adversarial examples* in which small tweaks to the data have been shown to cause gross errors by the classifiers.

Mission sets and specifications for both single vehicles and multi-agent teams remain limited. Additionally, there remains a marked lack of algorithmic development for operating in adversarial environments, including development of evasion and deception strategies. Work on swarm adaptation to losses and swarm reconfiguration to complete missions following failure or attack is also limited at present.

Using unmanned vehicles (or agents) with limited power, sensing, and communication capabilities in cooperative multi-agent teams may hold the promise for significant advantages over using a single, more sophisticated agent, particularly to achieve robustness and large-scale spatial coverage. However, these promises have yet to be fulfilled. Indeed, decisionmaking in multi-agent teams generally follows one of three categories of schemes: centralized, decentralized, and distributed.[23] Centralized schemes allow for global optimization but require large communication bandwidth and potentially provide a single point of failure. Decentralized and distributed schemes, on the other hand, require less communication bandwidth and reach and may, in theory, allow for better scalability. However, they still require minimal communication connectivity, and their design has proved to be challenging for large-scale systems and systems with complex specifications or performance objectives.

Finally, V&V remains a challenge. Control-based design approaches are generally well grounded in theory, but their implementation in software code is not an exact process, and the underlying models may misrepresent important features of reality. Both factors open the door to errors with potentially serious consequences. While some V&V approaches exist, their scalability remains limited. On the other hand, where machine-learning approaches are used, progress in the way of explanatory principles has lagged the rapid advances in algorithms and their demonstrated use to solve practical problems—a fair characterization even while noting some progress on ensuring robustness to distributional shifts. This may be a particular challenge in the undersea environment, where the Navy contemplates using many of its unmanned systems.

Closing Thoughts and Recommendations

We conclude our discussion of algorithms with some closing thoughts about the current state of algorithmic developments and where they may need to go to enable true autonomy.

Notwithstanding their cultural differences, both control theory and machine learning effectively encode assumptions about the system and environment, whether explicitly (through the models) or implicitly (through the data). As a result, while they have varying degrees of robustness to uncertainty as it arises, it is fair to say that neither enables true autonomy at present or in the foreseeable future. Indeed, should the world or circumstances (that is, system, environment, and objectives) drastically change, it is highly unlikely that the current suite of best-in-class algorithms would enable an unmanned vehicle to recover, adapt, and complete its mission.

---

[23] See Appendix A for details.

Although it is unclear what the path to true autonomy might look like, we can nonetheless make two general high-level recommendations that may help open up new windows of opportunity:

1. For a variety of reasons, including cultural ones, control theory and machine learning have developed in parallel for the most part and continue to do so with few exceptions.[24] In all likelihood, a combination of model-based and data-driven approaches will be needed to solve problems that neither approach on its own can solve. We thus recommend support for research aiming to bridge the divide between the two approaches.
2. While big data and machine learning go hand in hand, it may be fruitful to focus research efforts on the questions of what can be learned from data derived from local operation of systems and sensors and how. This is especially important given that the dynamic nature of the battlefield environment may preclude the collection of large data sets in the likely fast timescales on which the world is changing.

Payloads

The second autonomy building block is payloads. Payloads are devices or equipment carried by autonomous systems that make it possible for them to interact and interface with their environment and that are needed for autonomous mission completion. We first outline the approach we took to survey and categorize existing payload technologies for UUVs and USVs. We then discuss notable types of payloads, the aspects of existing missions they can contribute to, and unmanned maritime vehicle design trade-offs associated with using different types of payloads. Finally, we discuss growth areas for specific technologies, including where commercial research is headed and where future DoD investment may be beneficial to drive the technology to be more militarily useful.

Analytic Approach

We organize our survey of existing UUV and USV payload technologies around four categories: *navigation*, *sensors*, *communications*, and *weapons*.[25]

---

[24] Although efforts that bridge the two fields do exist (e.g., in deep reinforcement learning), they remain a small fraction of the research undertaken by either community.

[25] Notably, some types of payloads (e.g., specialized data-processing equipment or other special mission equipment, such as manipulator arms) do not fit cleanly into one of these four categories. We identified several such payloads in our survey and have grouped them in a separate category, *miscellaneous payloads*, which is described in further detail in Appendix B.

Payloads used for navigational purposes provide a host platform with data about its current position, velocity, and acceleration. Approaches to navigation for UUVs and USVs differ in that UUVs must contend with the effects of signal attenuation presented by the ocean, making such techniques as GPS-aided navigation impractical. The subcategories of payloads we discuss include inertial, GPS, radar, and acoustic positioning systems.

The primary purpose of sensor payloads is collecting data relevant to a platform's mission beyond platform navigation. Types of data collected by sensors include acoustic profiles of nearby contacts, radar and infrared signatures, electro-optical images, local ocean depth, and local ocean turbidity, among others.[26] Major subtypes of sensor payloads that we consider are sonar, radar, environmental, and light or optic sensors. These are discussed in further detail in Appendix B.

Communication payloads facilitate transmission of data between the autonomous platform and a central controller, a collection node, or another autonomous platform. Communication methods vary greatly between UUVs and USVs based on the surrounding environment's effects on signal transmission paths; these methods include radio, laser, fiber-optic, satellite, and acoustic means.

We survey weapon systems that can support MCM and surface decoy missions, with particular focus on minesweeping gear, electronic warfare systems, and USV weapon suites.

In Appendix B, we briefly review payloads that do not cleanly fit into any of the aforementioned categories (and that we thus grouped into the miscellaneous category).

Having identified key elements of each of these four payload types, we now assess the strengths and weaknesses of each. As part of this assessment, we consider the current application of these payloads to Navy missions, as well as additional applications of these capabilities to existing maritime kill chains. The next sections summarize our findings, and a more in-depth discussion of each area is included in Appendix B.

General Trends and Remaining Gaps

Navigation

To navigate maritime vessels, combinations of sensor payloads (and algorithms) are employed to provide reference points and mitigate errors that accumulate over time. The selection of the navigation payload suite is dependent on the operational environment and the precision required. USVs typically use GPS for navigation but can use radars or the inertial navigation system (INS) if satellite signals are unavailable. UUVs do not have the capability to use GPS signals (due to attenuating effects of water) without resurfacing, instead relying on acoustics, sonar, cameras (light and optical capabili-

---

[26] Some of these data may be useful to an autonomous vehicle's navigation systems in addition to its mission systems. But the sensor payloads that we identify support primarily an autonomous vehicle's mission systems as opposed to its navigation systems.

ties), INSs, or combinations thereof to navigate. Fusing data from multiple sensors can improve navigation but requires more onboard processing power, a challenge for vehicles that have power constraints but still require high levels of navigational precision. SLAM, surveyed in the section on algorithms, is one such development that leverages multiple data feeds.[27] Improving the navigation capabilities of maritime vessels is also a subject of interest for the commercial world, and we expect that further commercial development of navigational payloads and algorithms will address challenges faced in such industries as oil and gas. Table 2.3 summarizes current trends and findings pertaining to navigation payloads.[28]

Table 2.3
Trends and Findings: Navigation Payloads

Major Subtype	Capability	Design Trade-Off	Implementation Challenge or Limitation
Inertial	Collect data from accelerometers and gyroscopes to determine position, orientation, and velocity	Data-processing capability, power demand, sensor calibration	INS requires data processing and fusion of data from multiple sensors to correct for drift errors.
GPS	Provide three-dimensional (3D) positioning continuously while in the satellite coverage area	Data rate of communication link, signal frequency	GPS is not completely covert, is susceptible to interception and jamming, and is unavailable under water.
Acoustic	Use acoustic transponders to determine positioning relative to receivers or features (e.g., seafloor)	Sensor geometry	Some sensors require fixed infrastructure, and some require bottom-lock; water presents environmental constraints; and some systems have speed restrictions.
Radar	Combine radar imagery with sea charts to determine positioning	Sensor geometry, power demand, data-processing capability	Using radar as a navigational tool requires feature-rich environments and is limited to use above water.
Depth	Measure the ambient pressure in the water column to calculate depth	Sensor configuration	Limitations for this subtype are minimal. Measurement sensors will function at depths much greater than projected platforms are intended to go.
Orientation	Calculate platform heading from one or several sensors	Power demand	Performance is degraded during acceleration.
Light and optical	Determine positioning using environmental features (e.g., stars, pipeline) as a guide	Data-processing capability	Environmental constraints, such as water and fog, limit accuracy.

---

[27] Hugh Durrant-Whyte and Tim Bailey, "Simultaneous Localization and Mapping: Part I," *IEEE Robotics and Automation Magazine*, Vol. 13, No. 2, 2006.

[28] In Tables 2.3–2.6, *design trade-off* describes features or capabilities that are degraded as a result of using that payload type.

Sensors

As summarized in Table 2.4, a variety of sensor types are available to autonomous maritime system designers as they construct systems to carry out Navy missions. Each sensor subclass is covered in detail in Appendix B. Following our review of existing sensor technologies, we assess that, with a few exceptions, Navy investment will likely be the major contributor to future game-changing developments in maritime sensing technology, particularly undersea. Civilian parties with an interest in maritime sensing technology include academia, environmental protection, and the oil and gas industry. While these parties have an interest in improved sensing and data collection, the military's needs in this area are more specialized and likely warrant additional specific investment.

One apparent growth area is the fusion of diverse sensor data, as previously discussed in the algorithm review. Sensor fusion in an operational environment is challenged by the aforementioned limited available processing capability aboard unmanned maritime vehicles. However, improving the fusion of data collected by multiple sensor types may be a more effective investment than improving sensor technology because rapid shifts in sensor paradigms appear to be uncommon (with the possible exception of light detection and ranging).

Another potential investment area in the realm of unmanned maritime sensors is the collection of large amounts of environmental and other data for future process-

Table 2.4
Trends and Findings: Sensor Payloads

Major Subtype	Capability	Design Trade-Off	Implementation Challenge or Limitation
Sonar: single beam, multibeam, sidescan, synthetic aperture	Passive monitoring, target detection and identification, buried-object detection, imaging	Sensor geometry, power demand, data-processing capability	More-advanced systems need significant processing power; in addition, sonar systems limit platform speed and are available only under water.
Radar: S band, X band, dual band	Target detection and identification, targeting	Sensor geometry, power demand, data-processing capability	Radar systems have a high power demand, are large, and are limited to use above water.
Environmental	Environmental parameter measurement (temperature, salinity, turbidity, etc.)	Sensor geometry, data-processing capability	Environmental sensors require data processing and fusion of data from multiple sensors.
Light and optical: light detection and ranging, laser line scanning	Laser mapping and imaging, video feed generation	Sensor geometry, data-processing capability	The minimal optical wavelength propagation through water drives the requirement to be physically close to targets of interest when used under water.

ing. While various types of maritime data benefit Navy operations planning, the fact that environmental and other maritime conditions may change frequently presents a significant data-collection challenge. Extensive early use of inexpensive maritime environmental sensors for worldwide data collection would provide the Navy with large data sets that could be used to, among other things, develop improved sensor technologies and improved data-collection and analytic methods. This is especially important because the effectiveness of most payloads described in this section depends a lot on ambient ocean conditions that can change rapidly and are highly geolocation-dependent. Increased collection of a large amount of worldwide ocean data may be a relatively inexpensive venture that could contribute to a better understanding of ocean conditions and provide for predictive corrective modeling for sensors involved in future Navy missions.

Communications

In addition to enabling mission planning updates, communication payloads allow for (1) coordinating between vehicles and other manned or unmanned platforms and (2) providing remotely collected ISR data to a processing or decisionmaking authority, although the extent of communication availability and bandwidth required is largely dependent on the approaches used. Next, we discuss major communication technologies used by autonomous vehicles today.

Communication methods for maritime vehicles have remained relatively static for many years.[29] Though not covert, satellite and radio communications remain the most reliable high-data-rate methods of communication for UUVs and USVs. Communications tethered to fiber optics are also reliable but are limited by tether length. Acoustic communication technology provides the opportunity for more-covert communications, although the data rate for such communications is very limited, and in today's complex, high-volume data environment, such channels do not fully meet operational requirements. Other high-data-rate communication technologies, such as laser communications, are limited in the underwater environment because of the physical properties of seawater.

At the moment, Navy program-of-record investments do not address issues of data rate, requirements for covert communication, or the physical security of encryption equipment. While there may be advances in data rate from investment by the commercial sector, covertness and advanced encryption are military-specific needs. Table 2.5 summarizes current trends and findings related to communication payloads.

---

[29] Our review of declassified RAND research on Project AYMARA, a 1970s Defense Advanced Research Projects Agency (DARPA) program for an autonomous undersea vehicle, identified the same major communication technologies as we identified in this study.

Weapons

Maritime vehicle weapon payloads vary greatly depending on the mission and platform design. Additionally, the military is the only source of investment for improvements in these technologies. Fortunately, the high degree of commonality between manned and unmanned platform weapon capability requirements likely enables a high degree of reuse of systems that are already designed. A significant challenge facing weapon system use on unmanned platforms is control: Weapons are remotely controlled, and engagement decisions must be made by an operator. Table 2.6 reviews the trends and findings from our survey of weapon payloads.

Table 2.5
Trends and Findings: Communication Payloads

Major Subtype	Capability	Design Trade-Off	Implementation Challenge or Limitation
Radio and satellite	Long distance, relatively high-data-rate point-to-point communication or global communication via connection to a space-based asset	Antenna geometry, data rate of communication link	Radio and satellite communications are not completely covert and are vulnerable to interception and jamming.
Acoustic	Low-data-rate point-to-point method of undersea communication[a]	Data rate of communication link	Acoustic communications are vulnerable to jamming and ambient noise and have very low data rates.
Tethered	High-data-rate means of communication with platform	Length of tether	A tethered vehicle is limited in range to the length between the tether and the control station.

[a] Acoustic communications are the only reliable means of communication when an asset is submerged.

Table 2.6
Trends and Findings: Weapon Payloads

Major Subtype	Capability	Design Trade-Off	Implementation Challenge or Limitation
Minesweeping gear	Cutting or detonation of mines in the water column	Sweep gear type (contact, acoustic, magnetic)	Influence minesweepers must predict the corresponding signal, and the gear is susceptible to anti-sweeping mechanisms.
Electronic warfare systems	Electronic support, detection, and countermeasure protection	Sensor geometry, power demand, data-processing capability	Such systems have a high power demand and are large, and more-advanced systems need significant processing power.
Weapon suites	Force protection	Weapon type (machine gun, rocket launcher, missile launcher)	The systems are remotely operated and require a human decision to engage targets.

Platforms

Analytic Approach

We investigated platforms being developed or used by academia, industry, or DoD (particularly DARPA and the Navy) and developed a catalog to track our findings, listing up to 25 key parameters for each platform.[30] We categorized all platforms as either UUV or USV and report our findings in those terms. Our search resulted in the identification of 178 UUV platforms and 89 USV platforms, for a total of 267 platforms. Using our catalog, we identified the trends related to the platform design parameters that were readily available for most platforms—namely, platform length, weight, speed, endurance, and power source. We also identified trends in industry and nation of origin, as we report in this section.

Summary of Findings

As detailed in Appendix C, industry-developed platforms accounted for a significant portion of our findings. While this demonstrates significant industry investment in unmanned maritime platforms, there appears to be overlap in the capabilities being developed by the plethora of industry entities. With a few specialized exceptions,[31] our survey found commonality among platforms designed by all three development sources in terms of both platform design parameters and design basis missions for the platforms. From a platform design perspective, the majority of platforms were small to medium-size with a relatively low maximum speed and endurance, as discussed in further detail later. From a platform mission perspective, surveillance, mapping, and monitoring or inspecting were among the most common listed missions for the platforms we surveyed. The most common military-specific mission identified in a platform's open-source design documentation was MCM.[32]

Our survey identified open-source literature for platforms designed by academic and industry entities from 24 different nations. Of industry-developed platforms identified in our survey, about half were developed by U.S. organizations, and European Union nations developed a large percentage of the remaining platforms. This highlights

---

[30] In an ideal setting, we would investigate stand-alone platforms devoid of the payloads and detached from specific uses. In practice, much of the data we were able to obtain pertain to specific uses of platforms equipped with particular payloads suited for the use. This is what we report on in this section.

[31] DARPA's Anti-Submarine Warfare Continuous Trail Unmanned Vessel, or ACTUV, program is one such exception. See Sandra Jontz, "DARPA Christens New Sea Robot Vessel Sea Hunter," *SIGNAL Magazine*, April 7, 2016.

[32] Antisubmarine warfare (ASW) and security were also military-specific missions identified for platforms we surveyed. However, as we discuss in further detail later in this report, this should not be taken to mean that these platforms had end-to-end capability in any given kill chain. Rather, most of the ASW and security platforms we identified were capable of the early stages of these kill chains—detecting and identifying targets—which made them very similar to other unmanned maritime vehicles that performed more-generic survey or monitoring missions.

a challenge with using open-source literature for a survey of this kind: As discussed in further detail later in Appendix D, significant investment in unmanned maritime platforms has occurred in China and Russia, among other nations, but these platforms were not included in our search because information about them that is openly available in literature searches and translated into English is limited.

Platform Length and Weight

Data on platform length were available for 242 platforms. As shown in Figure 2.1, most UUVs and USVs identified were relatively small: More than half of the UUVs were 10 ft long or less, and the majority of USVs were 20 ft long or less. This is significant because small platforms have reduced onboard space, weight capacity, and power capacity to carry out complex sensing missions or missions with intensive data-processing needs. The mission sets of the few large platforms (50 ft long or more) that we identified included (1) hydro-acoustic research and mapping and (2) oceanographic survey. These platforms had longer relative endurance and needed to carry large sensing and data-processing payloads that may have driven the size of the design.

Data on platform weight were available for 192 platforms. The weight distribution of platforms presented in Figure 2.2 corresponds with length: Shorter platforms were typically lighter. It is not inherent in platform design that lighter weight implies any particular limitation. Our observation is that, among the platforms currently in use, lighter weight corresponds to platforms with less payload or power-generation capability, while heavier platforms generally include applications that require greater

Figure 2.1
Platform Length Among Surveyed Vehicles

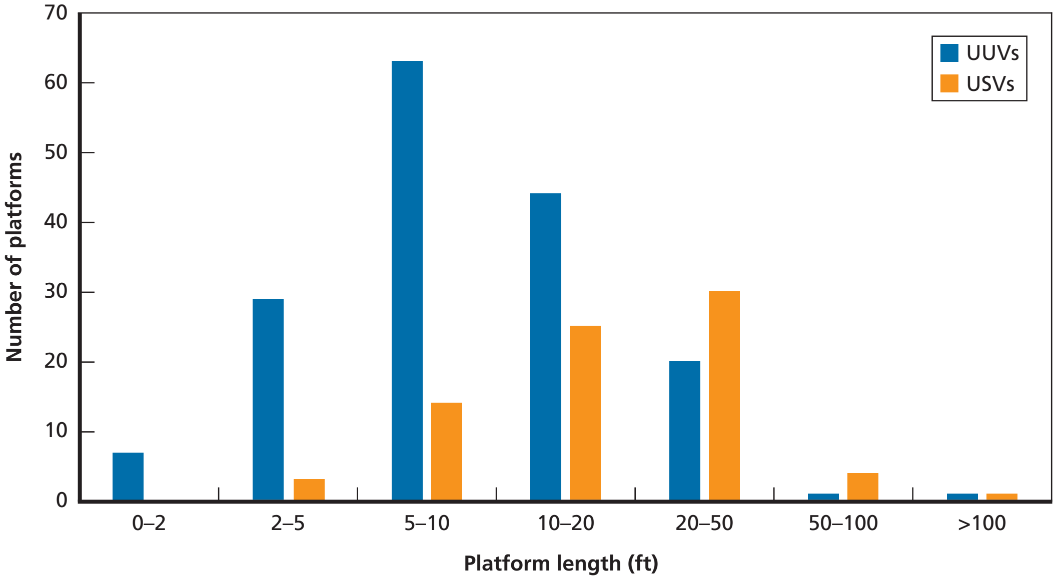

Figure 2.2
Platform Weight Among Surveyed Vehicles

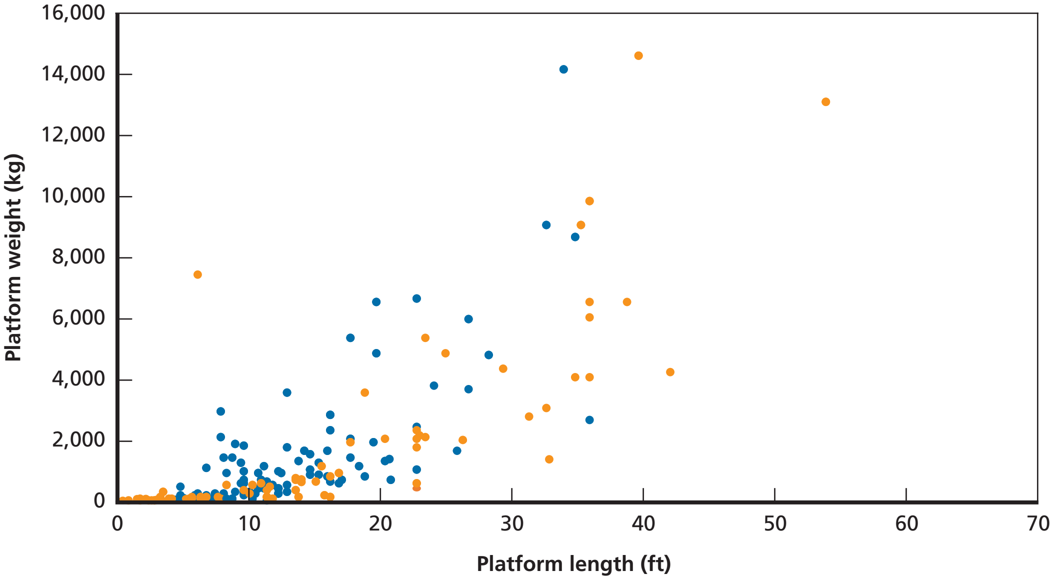

sensor capability and endurance. Notably, the design basis missions of the heavier outlier platforms include undersea mapping and tethered salvage operations.

Platform Speed

The requirements for platform speed vary by mission. For some missions, platform speed is a vital parameter for multiple kill chains. For instance, unmanned vehicles attempting to serve as decoys for manned vessels must be able to travel at speeds consistent with the platforms whose signatures they simulate. Platforms tracking adversary ships or submarines must also be able to travel at high enough speeds to keep pace with the vessels they are tracking. However, for missions in which the primary sensor is a sonar, such as minehunting, speed is less of a consideration because the speed will be limited by acoustic conditions and the need to look at the sonar contact with multiple sensors over a period long enough to enable classification and identification. To a degree, the speed observed for platforms is conditioned by the mission the platform and sensor are intended to support.

Data on maximum speed were available for 170 of the platforms in our survey. Across all platforms, UUV maximum speed averaged 5.7 knots and USV maximum speed averaged 22 knots. To understand potential relationships between platform size and maximum speed, we took the Figure 2.1 platform length distribution histogram and overlaid it with an indication of platform maximum speed for each previously identified length range. The result is shown in Figure 2.3.

Figure 2.3
Platform Length and Maximum Speed Among Surveyed Vehicles

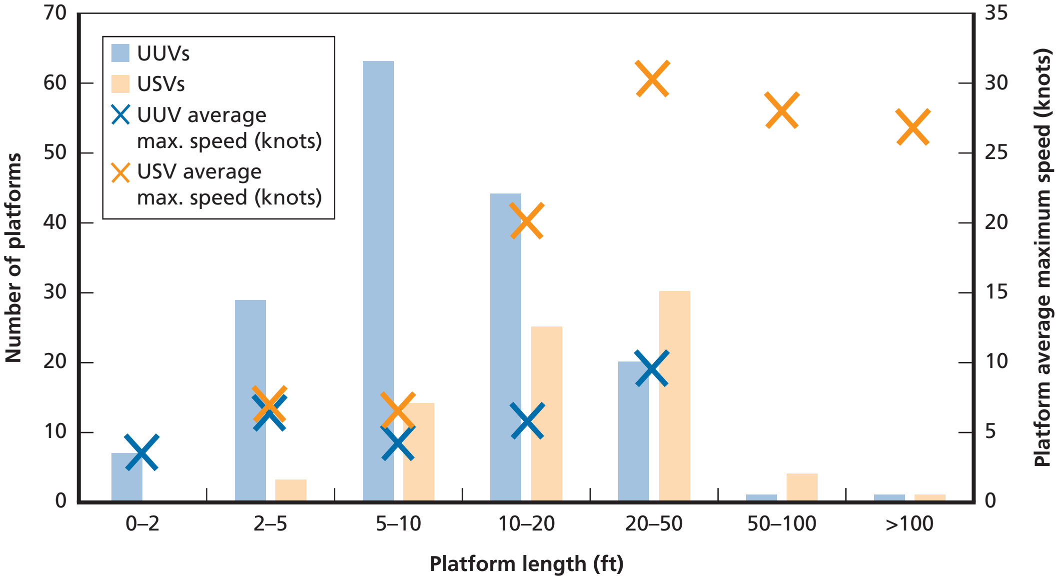

Comparing platform length and maximum speed demonstrates that platforms with higher maximum speeds were larger than those with lower maximum speeds. While multiple factors contribute to this, a higher platform maximum speed requires exponentially more propulsion power as desired speed increases for a given weight.[33] Notably, traveling at maximum speed is a high power draw that challenges the limited power resources aboard unmanned vehicles. Among vehicles we surveyed, the maximum speeds of UUVs were markedly lower than that of USVs, particularly at lengths greater than 10 ft.

In addition to collecting platform maximum speed data, we collected platform nominal speed data. However, only 70 data points were available for the latter, and the distribution of these points limited our ability to draw meaningful conclusions regarding this parameter. That said, when maximum and nominal speeds were both listed for a given platform, available documentation indicated that the endurance for the platform design was based on traveling at nominal speed rather than maximum speed. Traveling at maximum speed required significantly greater power usage and reduced

---

[33] The relationship between maximum speed and required propulsion power for ships varies based on hull type, but, generally, increased speed requires exponentially more propulsion power. Platforms with a higher maximum speed must be relatively large to carry sufficient propulsion power to achieve the speed; as platform weight increases, still more power is required. These relationships combine to drive platform size quite high when high maximum speed is required. See Thomas C. Gillmer and Bruce Johnson, *Introduction to Naval Architecture*, Annapolis, Md.: Naval Institute Press, 1982.

the endurance of these vehicles. This finding is important because, in many cases, platform nominal speed was less than 70 percent of platform maximum speed; even for the highest-speed USVs, it was less than half.

Platform Endurance

Figures 2.4 and 2.5 show platform endurance data. As with speed, the requirements for platform endurance may vary, and there is no assumption that endurance is good or necessary in every case. These figures simply report the state of the art of platform endurance as we observed and cataloged it. Figure 2.4 shows a histogram of platform average maximum endurance, and Figure 2.5 shows the Figure 2.1 platform length histogram with an overlay of platform average maximum endurance. As demonstrated by Figure 2.4, of 140 collected endurance data points, 96 had an average maximum endurance of one day or less, and only 24 platforms boasted endurance greater than five days. With one exception, the platforms we identified in our survey with a maximum endurance greater than five days were small, passive sensing devices with very limited onboard data-processing or propulsion capability.[34] These platforms, generally referred to as *gliders*, are designed as long-duration environmental sensing platforms with limited maintenance or intervention required.

Figure 2.4
Platform Average Maximum Endurance Among Surveyed Vehicles

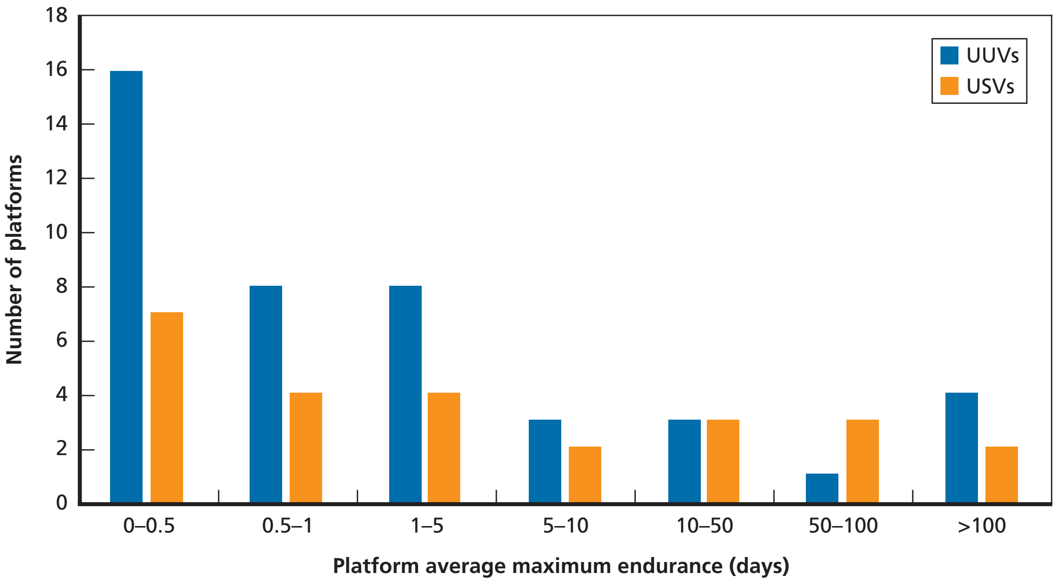

---

[34] The exception was DARPA's ACTUV, which is a USV.

Figure 2.5
Platform Length and Average Maximum Endurance Among Surveyed Vehicles

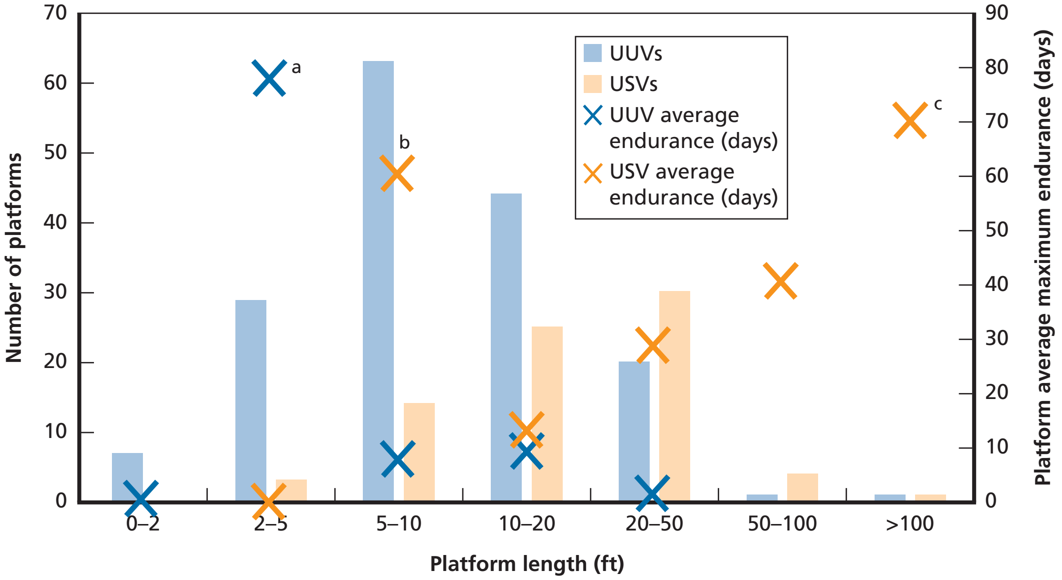

[a] There was a high average maximum endurance of UUVs in this bracket resulting from two passive, long-duration sensing vehicles—one with a 500-day endurance and one with a 1,460-day endurance. Without these two data points, the bracket average reduces from 77 days to 0.965 days.

[b] There was a high average maximum endurance of USVs in this bracket resulting from two passive sensing platforms with endurances listed at 365 days. Without these vessels powered solely by waves and ocean current, the bracket average reduces from 60 days to 4.54 days.

[c] There was a high average maximum endurance of USVs in this bracket resulting from the listed endurance of DARPA's ACTUV (the only platform in this bracket).

Platform Power Source

Given the importance of platform endurance to multiple mission areas and the limitations in existing platform endurance described in the previous section, an understanding of unmanned maritime vehicle power sources is important to help the Navy identify technology investment areas that could improve platform endurance. As is true of manned platforms, preferred power sources differ greatly between undersea and surface platforms. In terms of the current state of the art, battery technology is extremely common, especially for undersea vehicles.[35] Among both military and industry platforms, lithium-ion batteries are the most common because their energy density is higher than other batteries. Lithium-ion batteries are often associated with safety concerns related to fire hazards, but this is more of a concern for manned platforms than for unmanned ones. Figure 2.6 shows the distribution of platform power sources among the UUVs we surveyed.

---

[35] Undersea vehicles are less likely to use diesel engines because such engines generate exhaust and noise, and the vehicles cannot take advantage of wind or solar power while under water.

Figure 2.6
Platform Power Sources Among Surveyed Unmanned Undersea Vehicles

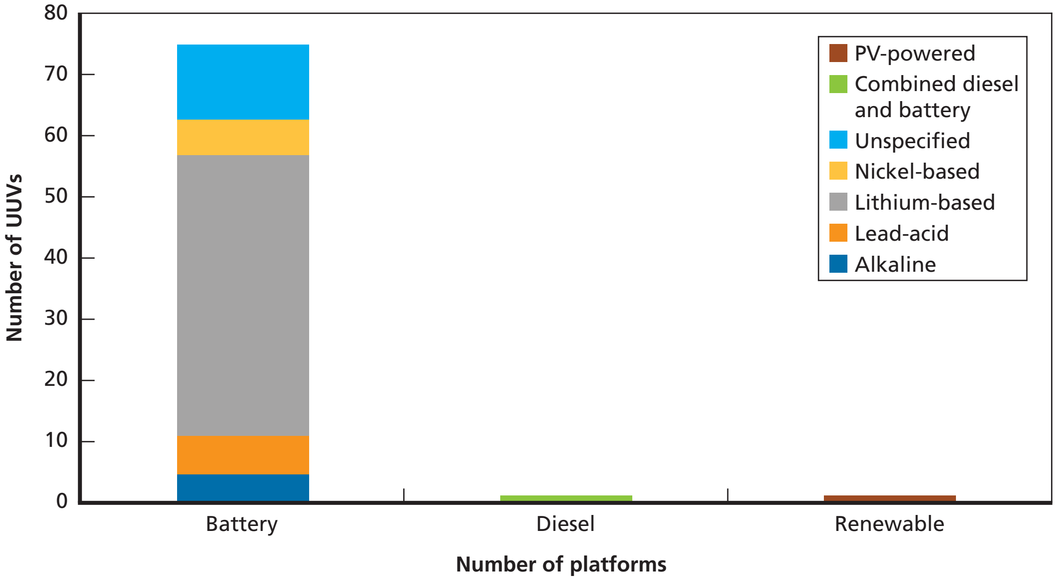

NOTE: Platforms using diesel and renewable power sources are designed to surface intermittently to exhaust their engines and charge their onboard batteries using solar power, respectively.
PV = photovoltaic.

For the USV platforms we surveyed, diesel or gasoline engines were preferred to batteries, as shown in Figure 2.7, because of their higher instantaneous power output and resulting ability to enable higher maximum speed. However, using such engines brings an additional complication. With no manned operators on board, USVs using diesel or gasoline engines require additional machinery-control autonomy features to start, stop, and change engine operational speeds on the fly. This requirement adds complexity to system control features. While all autonomous platforms must have an understanding of their maximum available power for mission management, battery energy management requires less machinery-control complexity because it is more a function of battery discharge rate and remaining energy capacity.[36]

Platform Survey Conclusion

Although significant resources are being invested in unmanned maritime vehicles, much progress remains to be made in multiple areas to enable the performance of more-complex missions. Our review of currently available unmanned maritime plat-

---

[36] Power source information was available for only 108 of the platforms we identified. The literature identified two additional power sources—fuel cells and nuclear power—as promising possibilities for unmanned maritime vehicles, but we did not identify a specific platform that used these technologies. And although the use of nuclear technology to power unmanned maritime vehicles could enable increased platform endurance, it would present numerous other technology and security concerns that we did not evaluate for this report.

Figure 2.7
Platform Power Sources Among Surveyed Unmanned Surface Vehicles

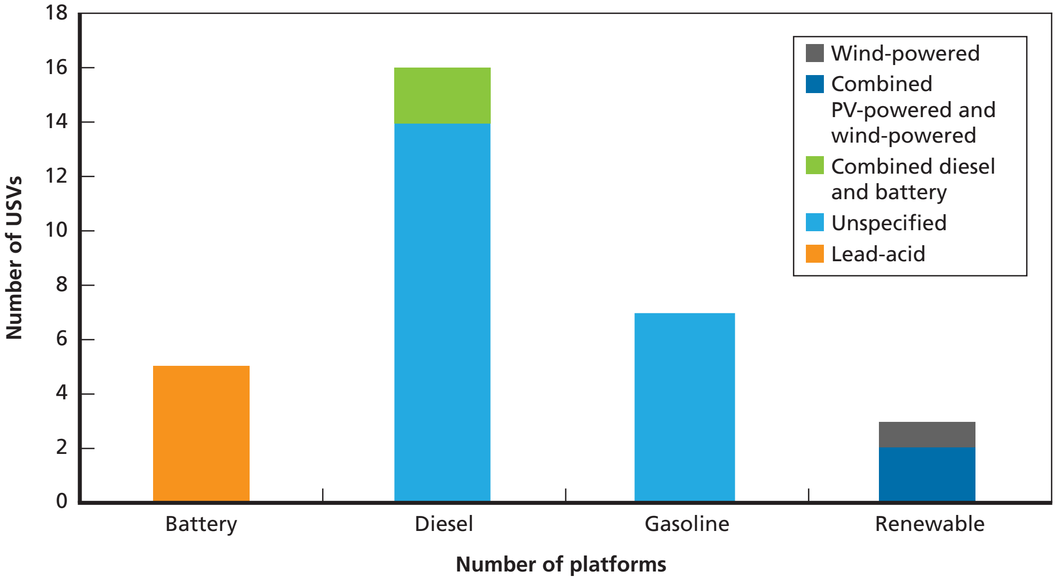

form technologies leads us to conclude that platform speed and endurance are two areas in which significant growth stands to be made. To some degree, both measures are driven by power source—specifically, the energy density available in battery technology. Fortunately, improved battery technology is an area of interest for various commercial industries and will likely draw heavy investment from many nonmilitary sources. However, because endurance is a key mission enabler, Navy investment in this area could help drive capability in a useful direction.

Conclusions

Following our in-depth analyses of the three building blocks of autonomy (algorithms, payloads, and platforms), we have two concluding thoughts.

First, significant codependence and interdependence exist among the building blocks. Advances in one area may help offset deficiencies in another. All three areas would benefit from further development, both individually and in tandem and taking into account design trade-offs, to push the state of the art in autonomous vehicles. Additionally, the use of several lower-performance systems (whether several lower-quality sensors on a given platform or several lower-cost unmanned vehicles) may sometimes be advantageous to the development of a single, more sophisticated system, provided that the requisite algorithms for fusion and cooperation are also developed.

Second, we are not generally optimistic that the reliance on developments in the commercial world and the use of commercial off-the-shelf technologies will be sufficient to provide the Navy with the capabilities it will need. Indeed, certain capabilities are Navy-specific and thus unlikely to receive much attention in the commercial or academic world. Moreover, even when the tasks or capabilities sought are not Navy-specific, the operational environment's scale, dynamic tempo, and unique resource constraints are such that solutions developed for the commercial world are unlikely to prove adequate in the battlespace environment. Accordingly, we make specific recommendations for targeted Navy investments as we proceed with our analysis of Navy missions in the subsequent chapters of this report.

CHAPTER THREE

Autonomy in Existing Kill Chains

Kill chain refers to the sequence of events required to achieve a warfighting effect. The use of *kill* is not meant to imply that the result of the process is physical death for someone or something; rather, it is intended to imply that multiple related process steps have resulted in a definite and desired outcome. Although *kill chain* appears in numerous documents and is understood by most users, it does not appear to have a doctrinal definition. However, across warfare areas, the steps of a kill chain generally contain some combination of find, fix, and finish—or, more broadly, intelligence preparation, detection, localization, targeting, and engagement. The kill chain approach has been used to describe processes ranging from targeting high-value individuals to enhancing force readiness.[1]

Our literature search led us to three areas in which the Navy believes that it can best exploit autonomy: MCM, denied-area ISR, and operational deception. These areas share the characteristic of being hazardous for manned systems, potentially requiring operation in hostile or uncertain environments. Depending on the CONOP, the missions might also be tedious and time-consuming, which can tax the endurance of human operators. There is a general propensity to assign unmanned autonomous systems to missions that are "dull, dirty, and dangerous."[2] For each of the three areas identified as potentially benefiting from autonomy, we examine the existing kill chain and then look at how existing or projected unmanned autonomous systems could compress or otherwise affect the kill chain's timeline.[3]

---

[1] Andrew Cockburn, "Kill Chain: The Rise of the High-Tech Assassins," *Huffington Post*, Thought Matters blog, December 6, 2017; and W. E. Gortney and C. D. Haney, *Introduction to the Readiness Kill Chain*, Norfolk, Va.: U.S. Fleet Forces Command, April 2013.

[2] Bernard Marr, "The 4 Ds of Robotization: Dull, Dirty, Dangerous and Dear," *Forbes*, October 16, 2017.

[3] Our research also shows that there is considerable interest in using autonomous systems to perform missions to defend homeland critical infrastructure. Because this is a mission more for local authorities or the U.S. Department of Homeland Security than for the Navy, we did not examine it in detail. However, many of the same issues and approaches that we developed here would also apply in a variety of other applications.

Mine Countermeasures

MCM is the process of neutralizing sea mines to a degree that the mines do not impair the ability of ships to transit sea lanes safely. Mines vary widely in explosive yield, actuation mechanism, placement (in the water column or at the bottom of the sea), and ocean conditions in which they operate (e.g., turbidity, bottom-scouring, and sediment movement). Some mines are tethered in the water column and may have a contact actuation mechanism—that is, they detonate upon contact—or may be actuated by acoustic, magnetic, or pressure influences or a combination thereof. Other mines are on the ocean bottom and actuate only through an influence or combination of influences.

The Current Kill Chain

Most kill chains begin with intelligence preparation of the battlefield (IPB), which is "an analytic process used to organize and analyze information on terrain, weather, and the threat within a unit's area of operations and associated area of interest."[4] This process allows more-efficient prosecution of subsequent steps. There are two basic kinds of MCM:

- minehunting, which is the process of positively detecting, identifying, and engaging mines to allow a high degree of confidence that all the mines in the field have been detected and neutralized
- minesweeping, which is the process of replicating the mine target's signature to cause influence mines to actuate or using a mechanical cable and cutter system to cut the mooring cables of a moored mine.

Minehunting allows for greater certainty but is more time-consuming; minesweeping is faster, but there is no guarantee that a sweep has been effective. Where acoustic conditions allow, minehunting, which involves the detection, classification, and identification of mines with sonar and visual sensors before neutralization, is the preferred tactic. This process provides a relatively high degree of certainty that mines are detected, separated from surrounding clutter, and positively verified to be neutralized.

Current Minehunting Kill Chain

The kill chain for current MCM vehicles is best depicted as sequential, with little opportunity for overlap (see Figure 3.1). The same general kill chain applies for both moored and bottom mines, although the neutralization mechanism may differ. These processes are manpower-intensive, and the kill chain can require up to an hour to complete for a single mine.[5] Given that there may be thousands of objects resembling mines

---

[4] Jamison Jo Medby and Russell W. Glenn, *Street Smart: Intelligence Preparation of the Battlefield for Urban Operations*, Santa Monica, Calif.: RAND Corporation, MR-1287-A, 2002.

[5] Scott Savitz, "Rethink Mine Countermeasures," *Proceedings Magazine*, Vol. 143, July 2017.

Figure 3.1
Minehunting Sequence

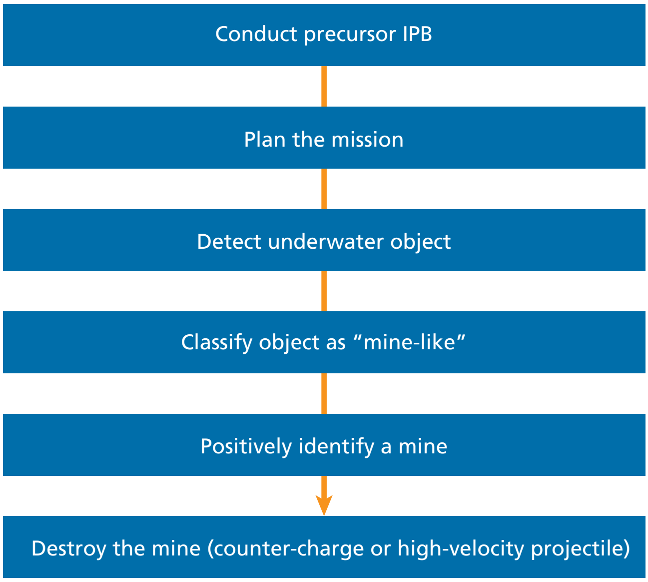

on a given ocean bottom, the timelines for mine clearance in some environments can be long and extensive. Moreover, because these are sequential processes, MCM vehicles must reacquire each mine multiple times, creating a greater chance of losing sensor contact with the mine if environmental conditions change or if the reacquiring sensor is not optimized for the environment.

Current Minesweeping Kill Chain

Unfavorable environmental conditions, such as mine burial, dense clutter, or poor acoustics, can make minehunting ineffective. Acoustically based minehunting may be completely impossible in very shallow water. When minehunting is not practical, MCM vehicles may be required to perform minesweeping. There are minesweeping systems to target acoustic and magnetically actuated mines. There is not, at the moment, any means of sweeping a pressure-actuated mine, other than by navigating a vessel that can generate a pressure signature sufficient to actuate the mine. The kill chain used in minesweeping (Figure 3.2) has features in common with that for minehunting, but there is no step requiring positive identification of the mine and, unfortunately, no real way of knowing whether the sweep has been effective if no influence mines actuate. The lack of actuation might be because there actually were no mines in the field or because the sweep did not correctly replicate the needed influence. Because minesweeping is being done because of poor conditions for sensors, there is no way to verify whether there is a mine on the bottom of the ocean.

Figure 3.2
Minesweeping Sequence

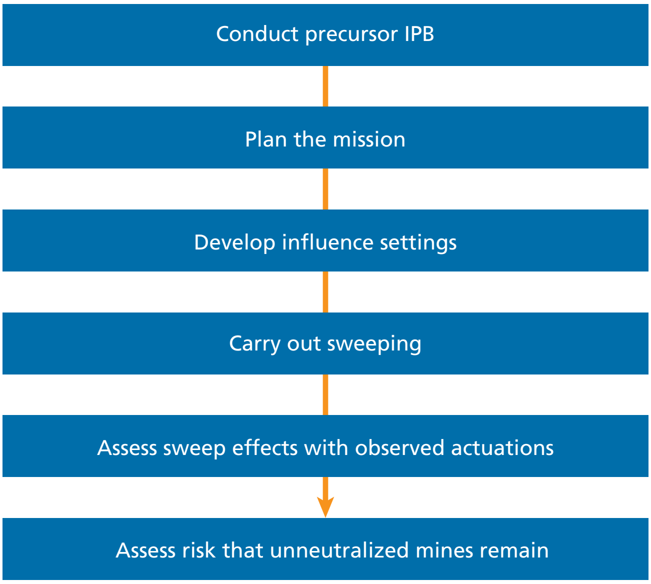

Although minesweeping is generally a more rapid process than minehunting, it carries with it some significant operational risk. First, it places a manned air or surface platform into a minefield without a clear idea of whether mines are there and what their character may be. The platforms are expected to use sweep systems whose success can be exactly measured only by mine actuations; if no actuation occurs, that does not necessarily mean that no mine is present; it might only mean that the sweep was ineffective. When actuation does occur, there is a standoff distance between the platform and the sweep (to protect the platform); however, calculation of the standoff range is not perfect, so explosions can occur closer to the platform than is intended. Such explosions too close to a ship can cause serious damage, and the explosive plume can be a hazard to airborne MCM systems. However, the more significant risk to the MCM mission is the lack of certainty about whether the sweep has been effective. When a sweep does not actuate any mines, a mine's presence might be announced only by a non-MCM ship triggering the explosion.

Autonomous Systems' Potential Contribution to Minehunting

As discussed in Chapter Two, several autonomous vehicles are intended for, or could have applications for, MCM. We begin with the applications that might improve the chances of positively identifying mines prior to their neutralization as part of the minehunting kill chain. As we have discussed, the current kill chain requires multiple passes at the same contact with different sensors and, as a result, imposes a long and difficult timeline. This timeline is, to a degree, the result of limitations of different sensors, but, primarily, it is the result of sensors not being connected together in a coherent way that enables them to share target information. Mine-like objects are detected and classified, reacquired and identified, then reacquired again and finally neutralized.

Autonomous systems currently being developed by and for the Navy, in some respects, just replicate this process with unmanned systems. Sensor detection is done with a towed sonar deployed from a manned helicopter or a USV, with both hosted on a littoral combat ship equipped with an MCM mission package. The mission package has several systems and sensors, some hosted on the organic MH-60 helicopter and some hosted on a USV, which replaces the canceled Remote Multi-Mission Vehicle.[6] The package will also include an autonomous UUV called Knifefish, which hosts forward-looking sonar sensors intended to detect, classify, and identify mines in a single pass. While, to date, the Knifefish has not successfully demonstrated this single-pass capability, developers and sponsors are confident that this capability is achievable.[7] None of these systems overcomes the requirement for the platform, even if it has identification and neutralization systems organically hosted, to mark and then reacquire the mine to deliver a neutralization charge. If the mine were neutralized in stride, it would pose a hazard to the platform and disrupt the environment for the minehunting sensor to the point that the search could not continue until after conditions returned to normal after the explosion of a neutralized mine. Neutralization generally occurs after the sensor search is complete, and this requires reacquisition.

Autonomous Systems' Potential Contribution to Minesweeping

The major requirements for minesweeping are the abilities to navigate along pre-set paths and correctly set influences for the expected target mines. The first task is relatively easy for autonomous systems; the second is not something that would be adjusted in stride, so the need for autonomous ability to correct settings might not be critical. Overcoming the inherent limitations of assessing sweep effectiveness will require the ability to identify and assess the state of mines, and this belongs back in the realm of minehunting. Having many autonomous sweeps with a variety of settings might raise the confidence that correct influences have been applied to the area of interest simply by dint of numbers and variety. But this does not depend in any significant way on any advanced features of autonomy.

---

[6] The Remote Multi-Mission Vehicle was a semi-autonomous platform intended to host the AQS-20 sonar and operate from the littoral combat ship. It was canceled in 2016 because of repeated failures in operational testing and repeated cost overruns. See Sydney J. Freedberg, Jr., "Navy Scraps RMMV Mine Drone, Accelerates CBARS," *Breaking Defense*, February 26, 2016.

[7] "U.S. DoD Finds More LCS Mission Package Problems," *Maritime Executive*, November 11, 2016.

Denied-Area Intelligence, Surveillance, and Reconnaissance

Historically, the U.S. Navy has used the stealth and high endurance of manned submarines to operate in areas where surface ships and aircraft cannot go (under water). The purpose of submarine missions has typically been to collect intelligence and perform surveillance and reconnaissance. Some of the missions have achieved notoriety for the risks involved.[8]

Typically, the material collected in this kind of operation is not used in an immediate tactical action—for example, in an attack that the submarine itself initiates. The submarine is operating in an area known to be well defended, and carrying out some kind of attack would potentially reveal the submarine's presence and its location. If the submarine attacked, it would alert the force defending the denied area and compromise the submarine's further ability to carry out the ISR mission.

Accordingly, the kill chain for the denied-area ISR mission does not generally involve an engagement step. There will be extensive mission planning; navigation to the area of interest; collection of intelligence and other relevant environmental information; some amount of onboard processing, exploitation, and dissemination of information; and then offboard transmission of the information in a way that preserves the platform and the mission. Figure 3.3 depicts this process.

Figure 3.3
Denied-Area Intelligence, Surveillance, and Reconnaissance Sequence

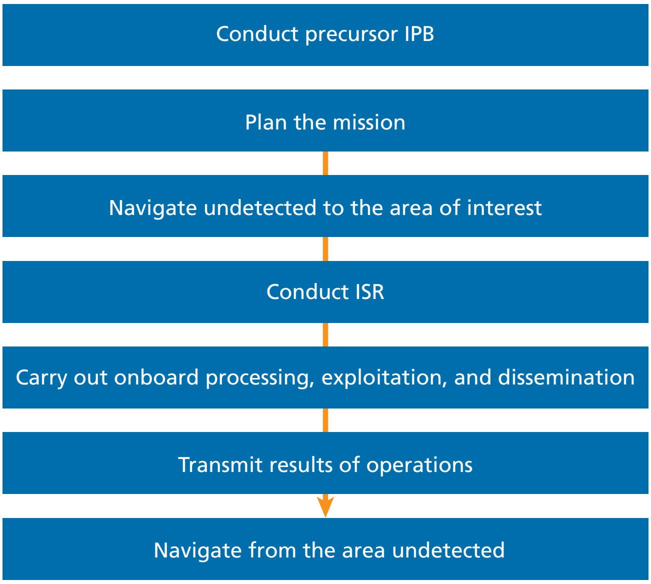

---

[8] Sherry Sontag, Christopher Drew, and Annette Lawrence Drew, *Blind Man's Bluff: The Untold Story of American Submarine Espionage*, New York: PublicAffairs Press, 1998.

Autonomous Systems' Potential Contribution

The denied-area ISR mission is hazardous for the crew of the submarine and requires the dedication of an exceptionally expensive platform that has other potential uses. While the submarine is performing this mission, it is not available for anything else. For this to be a cost-effective use of the submarine, the ISR must be shown to have exceptional value, and it may be difficult to assess this in advance of the mission. The United States has had access to information collected in this way, has undoubtedly found the results useful, and has good reasons for continuing to want it. However, the Navy has a strong interest in a platform that can carry out the mission at less cost and less risk to crews than those imposed by using a submarine.

When using unmanned autonomous systems for denied-area ISR, the kill chain would be similar in some ways to that when using manned systems. At the most rudimentary level, the platform would navigate into an area, move along a preplanned route, and passively collect data that will be analyzed once the platform finishes its mission and returns to its support platform (or port). However, this would result in a lengthy mission in which it is unclear until after the mission is complete whether valuable information is being obtained. If the mission is connected to a tactical mission requiring rapid dissemination of intelligence, a UUV with only simple autonomous ability to navigate and collect could not make the kinds of judgments that a submarine commander would make that some matters are of sufficient importance to break the stealth of the mission. The UUV would be very safe but of significantly less value than its manned counterpart.

Thus, if the UUV is going to do an effective job at denied-area ISR, it will need to have additional autonomous capabilities. These include the ability to do the following:

- Navigate from over the horizon.
- Internally manage energy consumption and production.
- Optimize sensors and collection systems.
- Judge the value of collected data to minimize transmissions.
- Transmit relevant data securely.
- Sense and appropriately react to threats.

All these actions involve degrees of cognition, learning, and the ability to react appropriately to changing circumstances. Our technological assessment does not indicate that many of these capabilities are imminent, at least not at a level that allows the UUV to perform the tasks that make the activity valuable.

Navy Program of Record

The requirement to navigate into a denied area from a long distance and remain in the area for a long period creates power-generation and storage requirements that are generally possible only in a large platform. Batteries and power-generation systems require space, and this requirement precludes use of smaller UUVs. The Navy's program of record for performing denied-area ISR is the large displacement UUV (LDUUV), which received Milestone A acquisition approval in 2015.[9] The LDUUV will be 23 ft long and able to conduct missions of 30 days at a time in open and littoral ocean waters.[10] The LDUUV is expected to reach squadron initial operational capability by 2020 and full-rate production by 2025.[11] It is being designed using a modular open-system approach, which is intended to provide flexibility in mission capability. At initial operational capability, the LDUUV ought to be able to (1) operate in complex littoral environments and identify and characterize a variety of undersea objects, (2) conduct low-profile continuous underwater operations, (3) have minimum shore- or ship-based communications during operations, and (4) have subsequent designs able to complete progressively complex missions.[12]

The LDUUV is being delivered with the ability to navigate autonomously and to avoid vessels in its area of operations.[13] It will also be able to monitor its own power requirements. It will be able to collect information in an area of interest, but, at initial operational capability, it will not be able to replicate what a manned submarine can accomplish—and indeed will not provide information within tactical timelines. Although the system is designed to host upgrades enabling more-complex missions as they become available, the Navy should carefully assess its investments to support the denied-area ISR mission for the LDUUV. In particular, the Navy needs to gauge the actual value of the mission, given the likely expense and difficulty of developing the autonomy necessary for the platform to carry out more-complicated missions. Even if such autonomous capability is developed, it is unclear that the LDUUV would provide timely and actionable information that other units could exploit. Unless the LDUUV has the ability to carry out engagements, the nature of the denied-area ISR mission may make it mostly a feature of long-term peacetime intelligence collection. If this becomes the major employment of the LDUUV, the simple ability to navigate and passively collect images, signals, and other components of ISR may be sufficient. However, this means that the LDUUV would not be performing the more expansive mission now carried out by submarines.

---

[9] Program Executive Office Littoral Combat Ship Public Affairs, "Large Displacement Unmanned Underwater Vehicle Program Achieves Acquisition Milestone," press release, Washington, D.C.: U.S. Navy, September 3, 2015.

[10] U.S. Navy, *Capability Development Document for Large Diameter Unmanned Undersea Vehicle System*, Washington, D.C., February 2014.

[11] "Large Displacement Unmanned Underwater Vehicle Innovative Naval Prototype (LDUUV INP)," *NavalDrones*, April 13, 2015.

[12] U.S. Navy, 2014.

[13] John Keller, "Navy Asks Metron for Autonomy and Control Software for Future Large-Displacement UUV," *Military and Aerospace Electronics*, May 22, 2013.

Operational Deception

Anti-access/area denial (A2AD) capabilities preclude surface ships—including aircraft carriers, amphibious ships, and cruiser–destroyer classes—and manned aircraft from operating in certain areas. The layers of sensors and overlapping weapon rings create multiple opportunities for adversaries to attack detectable platforms; in the most challenging A2/AD environments, targeted platforms are unlikely to survive, even with advanced kinetic interception capabilities.[14]

However, some capabilities might be effective in degrading the detection and targeting capabilities of A2AD networks, thus improving the ability of formations of surface ships—such as carrier strike groups (CSGs)—to penetrate and overcome them. There are three main ways to degrade detection and targeting. The first is by having the CSG operate in a restricted-emission environment and thus deny the enemy the ability to target the formation passively. Adversaries can still search using active sensors, but not receiving emissions complicates the task. This is most likely to be effective when accompanied by some other system that replicates the signature of a CSG and thus draws effort toward the false target. This tactic is intended to cause the adversary to search for and possibly even launch large raids on false targets.

To carry out long-range deception missions, the Navy has historically used manned platforms to replicate false targets.[15] During the Cold War, small combatants were equipped with “blip enhancers” to make them appear larger than their normal radar cross-sections.[16] Similarly, auxiliaries could load a suite of equipment and transmitters that could convincingly replicate air traffic control transmissions and other emissions generally unique to aircraft carrier operations.

If the attacking platform succeeds in generally locating the CSG, the next chance to disrupt the attack is to create false signals—that is, “spoof”—the targeting radars that allow attackers to more precisely place weapons. This tactic can be done by jamming, which requires generating sufficient power in the operating bands of the targeting radars to limit their power output or generating decoys sufficiently like the targets that the radars lock on them instead. The power requirements for jamming are sufficiently large that only ships or aircraft whose singular mission is electronic attack are capable of carrying it out. For ships, they have to be uniquely positioned to jam a targeting radar successfully, and this position might make them vulnerable in other ways. Note that this mission assumes that the attacker generally knows the position of the CSG but does not know with certainty the identity of valid targets.

---

[14] Ronald O’Rourke, *China Naval Modernization: Implications for U.S. Navy Capabilities—Background and Issues for Congress*, Washington, D.C.: Congressional Research Service, RL33153, August 21, 2018.

[15] Jonathan F. Solomon, *Defending the Fleet from China’s Anti-Ship Ballistic Missile: Naval Deception’s Roles in Sea-Based Missile Defense*, thesis, Washington, D.C.: Georgetown University, April 15, 2011.

[16] Solomon, 2011, p. 44.

Finally, if units of the CSG are localized to the point at which the attacking platforms have launched missiles, the last opportunity to carry out countermeasures is against the seeker heads (i.e., radars that actively search for and attack the target platform) of the launched missiles. Most cruise missiles have an active seeker head, but many of these also have other modes that home on jamming, heat signature, or even optical contrast. These countermeasures are organic to the ships being attacked and represent a last-ditch effort at self-defense, in conjunction with kinetic point defense systems.

Figure 3.4 depicts the operational deception kill chain, which is intended primarily to disrupt the attack kill chain of the adversary. Manned platforms carry these out now, and, in the case of deploying decoys against missile seekers, the systems employed are organic to the ships and used as a matter of self-defense.

Autonomous Systems' Potential Contribution

Autonomous vehicles could be very useful for many parts of the operational deception mission, particularly at the point at which deception is still a possibility. Unmanned autonomous systems can be sent off-axis in sufficient numbers and with emitters similar to those of CSGs. However, to be truly effective, these systems need to credibly replicate the behavior and signatures of CSGs. This deception force should not be a one-engagement type of capability. It should have the ability to shut down all its emitters,

Figure 3.4
Operational Deception

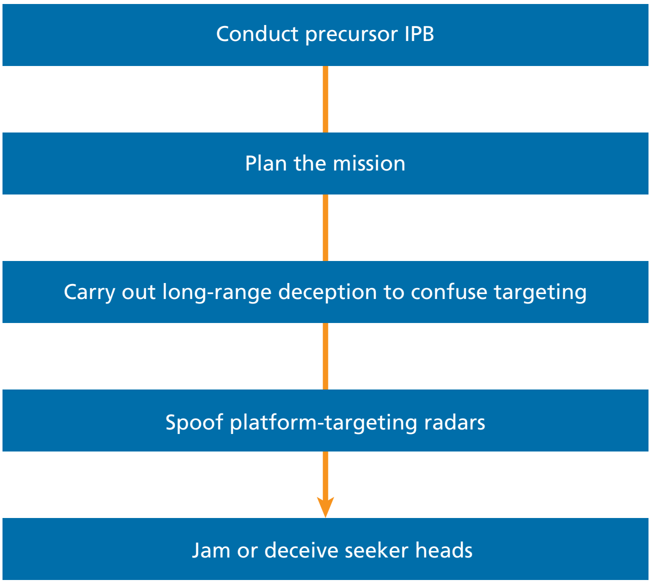

relocate, and then reenergize in tactically plausible places. This requirement implies the following capabilities:

- Independently navigate to areas out of the line-of-sight control capability of CSG assets.
- Coordinate positioning with other unmanned assets to allow a credible presentation of CSG-like force on radar.
- Operate systems that credibly replicate the emissions of a CSG.
- Perform active or passive self-defense.

The last capability could also be potentially useful in spoofing targeting radars. In this situation, the attacker has a general idea of where the CSG is located, so there is no opportunity to completely deceive the opponent to devote reconnaissance or expend munitions on empty ocean. The decoys are trying to generate a variety of plausible targets for the targeting radar, which may help make selection among several targets difficult. Rather than having the decoys widely separated, it is probably beneficial for the decoys to be relatively close to the platforms the radar is attempting to precisely locate. Consequently, there is no particular requirement for the vehicle to have the ability to navigate independently. It will, however, need the ability to sense the presence of a targeting radar signal and, without operator intervention, generate a signal that makes the radar identify it (rather than manned elements of the CSG) as the actual target. This has to be done based on processing and assessment that is on board the vehicle, not transmitted to it from some other source. Doing this successfully requires a high degree of autonomy, greater than anything currently in the Navy program of record.

These same decoys might be useful in the final stage of an enemy missile attack—after the missile's active seeker has locked onto a target and is commencing a terminal attack. At this point, the targeted unit will be attempting various techniques to either jam the seeker head or cause the missile to break lock on the ship and attack a chaff cloud. An autonomous system could generate additional chaff or a jamming signal, but it would need to do this either in response to the signals it receives from the missile or from commands given from the attacking platform. However, it cannot rely on close control by an operator. The time frames are far too compressed to require operator intervention. The system has to be assessing and reacting on its own. This application of autonomy is already feasible in shipboard systems, and it is realistic to think that this could be readily adopted to offboard systems. However, it would be a new requirement for the offboard systems to coordinate with one another and shipboard systems. If done incorrectly, the systems could counteract rather than reinforce one another. However, although the ability to react autonomously from a single platform is well within technological reach, this type of coordination is not currently contained in the Navy program of record.

Common Mission Area Themes

All three missions discussed in this chapter (MCM, denied-area ISR, and operational deception) benefit from the possible uses of autonomous systems. However, one common theme is that seeking autonomy for the purpose of replicating the kill chain used in manned systems might not be the best use for autonomous technologies. In fact, the thing that might prove most productive is the ability to share and collaborate, and this may be one of the more difficult technical challenges for autonomous systems. We consider these alternative CONOPs in Chapter Five, but the clearest sign to date is that this type of collaboration may be the most valuable but least developed of capabilities.

Another theme is that factors other than autonomy may be the most limiting for platform development. In the MCM mission, the problem is principally in the effectiveness of sensors or influence sweeps, not in whether the platform operating them is more or less autonomous. In the denied-area ISR mission, the greatest demands are in endurance and range, and these are not issues generally related to autonomy. Only in operational deception and surface decoys is autonomy the key platform limitation, and, even here, the limitation is not primarily in the ability to process information and react but rather to process information and react in a coordinated fashion. The fact that autonomy is not generally the limiting factor is important in considering effort and investment. Developing a highly autonomous system that lacks important elements of mission capability adds little warfighting value.

CHAPTER FOUR

Future Fleet Architecture and Its Autonomous Capabilities

We have considered applications of autonomy primarily in terms of the kill chains that currently exist and are likely to persist over the next several years—generally the five years of the Future Years Defense Plan. However, the Navy has a 30-year shipbuilding plan and three proposed FFDAs that are intended to project capability requirements further into the future.[1] Autonomous systems play a major role in these FFDAs and extend past the kill chains currently contemplated. In this chapter, we examine the potential role of autonomous systems in the future fleet, focusing on the mission areas we identified, with some modifications based on projected changes in the operating environment.

Key Elements Among the Future Fleet Design Architectures

The National Defense Authorization Act for Fiscal Year 2017 directed DoD, and the Navy in particular, to develop three separate fleet architectures that would help identify key capabilities and force structure for the Navy in 2030 and beyond.[2] Three organizations developed architectures: the Center for Strategic and Budgetary Assessments, the MITRE Corporation, and the Navy itself.

The three fleet architectures vary in their identification of promising technologies and reach different conclusions concerning required force structure. Although there were some differences among the studies, the following common themes emerged:

- The need to function in limited-connectivity environments. This implies a need for organic over-the-horizon sensing that can pass targeting information back to the host platform via a direct path rather than through a complicated communication architecture.

---

[1] Eckstein and LaGrone, 2017.

[2] Public Law 114-328, National Defense Authorization Act for Fiscal Year 2017, December 23, 2016; Eckstein and LaGrone, 2017.

- The need for "distributed lethality," in which dispersed units are capable of collaborating and providing coordinated fires across a wide area. This implies a need for an organic capability to overcome local anti-access threats, such as mines.
- The continued need to collect intelligence in denied areas, including where space-based surveillance systems may have limited availability.
- The preferability of preventing adversary subsurface assets from leaving bastion areas, thus minimizing the need to devote assets to open-ocean ASW.
- The need to deny the enemy easy targeting, with the presumption that any targeted platform is likely to be overwhelmed with the variety of missiles that adversaries are capable of delivering.
- The idea of autonomous surface action groups, with weapon delivery capability, operating in contested areas (e.g., the Black Sea or near denied areas of the South China Sea). This is not a radical extension of the decoy kill chain we discussed in previous chapters, but the weapons delivered would be kinetic instead of the jamming or electromagnetic deception that the systems discussed in Chapter Three would deliver.

Requirements for Autonomous Operations in the Future Fleet Design Architectures

When we look at these common themes, the essentiality of unmanned and autonomous systems is clear. The assumption is that many environments will be untenable for manned platforms and that operations may be feasible only with unmanned platforms. The need for manned platforms to have organic surveillance and offboard communication systems is also clear. These are likely to be hosted on unmanned platforms that do not require significant direction from the manned platforms, given an assumption of intermittent and asynchronous connectivity. The most general theme that emerges from the three architectures is that there will be a continuing need for concealment, dispersion, and coordinated action—but conducted in an environment with limited communications. Autonomy is likely to be essential in these conditions. However, the particular applications of autonomy are not necessarily included in the Navy's program of record, and it is not clear that they will be readily developed. In the following sections, we consider each of these applications by mission area. This approach differs slightly from the one we took in Chapter Three for current kill chains. Although elements of each kill chain will likely remain the same, the proposed FFDAs change the mission approach in some cases and expand the mission in others.

Overcoming Local A2/AD Challenges

All three proposed fleet architectures suppose that ships will be required to operate in areas subject to mining, even in areas that might be outside enemy land-lab missile envelopes. This supposition requires that surface action groups and larger formations need some kind of organic MCM capability, the components of which we have previously discussed. However, the MCM kill chains as currently conceived by the Navy still require a host platform that is effectively dedicated to the mission, even if that mission is to oversee the operation of offboard systems. The Navy's FFDA requires wide dispersion of ships, all with some capability to deliver missiles or other kinetic payloads. This requirement may mean that ships cannot be spared to perform single missions and that MCM systems will need to demand less space and operate nearly without intervention. This requirement might logically grow from kill chains as they exist but will likely require development of systems capable of autonomously evaluating underwater objects. Such a capability would require advancement in sensor processing and onboard decisionmaking.

Local A2/AD challenges could also come in the form of missile-firing surface platforms or submarines. Against these, the possibility of unmanned surface platforms operating as decoys might help confuse targeting and perhaps draw adversary platforms into the weapon envelopes of manned platforms. An unmanned platform capable of hosting an ASW helicopter, at least for takeoff, landing, and refueling, could add range and endurance for this mission and thus enhance defense. All these capabilities, however, are extensions beyond the program of record and depend on the development of other capabilities.

Collecting Intelligence in Denied Areas

The challenge in this mission is that denied-area bastions (complete with mines, ASW sensors, and countermeasures) are likely to get more rather than less formidable and thus stress autonomous UUV systems attempting to penetrate them to an even greater extent than in current situations. The vehicles will need to not only navigate past mines and obstacles but also, while in the denied area, overcome underwater sensors and weapons nearly continuously. Their ability to transmit collected information will be extremely limited because any emission or surfacing is likely to be detected and countered.

No doubt, ways exist of making already very expensive autonomous systems more capable of operating in these environments, but if the mission is questionable in the short term, it appears effectively impossible in the long term. It will likely be possible for vehicles to enter denied areas and perhaps even remain undetected. It will not be possible for them to move around to any degree while in these bastions or to pass collected information. The ability to penetrate and wait may prove very useful for other missions, such as offensive mining, and we explore those later. But the value of using autonomous systems for denied-area ISR might be highly questionable.

Performing Offensive Mining

Many of the capabilities intended for use in denied-area ISR may also be highly useful in allowing offensive mining and thus impairing the ability of adversary submarines to successfully leave bastions. Delivery vehicles' requirements for offensive mining are in some ways very similar to the requirements for denied-area ISR. Those offensive mining requirements are as follows:

- Navigate undetected from distances outside the sensor range of A2/AD systems. In other words, navigate from a pier, submarine, or ship outside the maximum range of coastal sensors.
- Avoid obstacles and countermeasures.
- Carry a large enough number of mines to make an effective minefield.
- Potentially, control activation of deployed influence mines.

A beginning premise is that the mining should be done covertly; the adversary's awareness that a minelayer has been deployed ensures that the adversary will either attempt to interdict the minelayer or apply countermeasures once the mines have been laid. To avoid these countermeasures, the minelayer needs to deploy in a manner that prohibits adversaries from knowing that it is inbound, which implies launch either from a submarine or from far over the horizon (to include from a friendly port). An equally important premise is that, for a minefield to be effective, it needs to be dense with adjacent actuation circles. Although a lightly seeded minefield might have some psychological effect on a risk-averse opponent, these types of fields can be readily cleared and, even when they are not successfully cleared, impose a low probability of damaging transiting vessels. An effective minefield requires a large number of mines, and this imposes a significant payload requirement on the minelaying UUV. Thus, the ability to monitor and regulate energy production and expenditure to ensure endurance with a heavy payload may be particularly important for this mission.

The last steps in laying a minefield are setting and then enabling the actuation characteristics of the mines. While there are different CONOPs for how the minelayer might seed the minefield, the most likely is for the minelayer to transit to the field and lay the mines, either in a preplanned pattern or in a pattern determined after the vehicle performs an onboard environmental analysis. This process might occur before hostilities, so it is clear that the mines should not be capable of actuation until after authorization, which will require transmitting a signal that tells the mines to turn themselves on. This signal could come from a manned platform or from the UUV that remains in the vicinity of the minefield. The UUV would have to receive an authorizing signal, but its ability to reliably transmit an arming signal to the mines that it laid is likely greater than one sent to several mines from a more distant transmitter.

Denying Targeting Information and Operating Autonomous Surface Action Groups

Whatever version of the FFDA comes to fruition, it assumes an environment in which A2/AD threats have reached a level of sophistication and pervasiveness that a located surface ship within an envelope will have great difficulty surviving. Consequently, the FFDA depends heavily on the ability to deny targeting, which can be accomplished by dispersion, spoofing, emission control, and use of offboard sensors that communicate asynchronously with the host platform. We have discussed many of these in the context of single kill chains. For the entire FFDA, this type of capability, which relies heavily on autonomy, goes beyond desirable; it is essential.

One proposed version of the FFDA goes beyond stating a requirement to deny targeting and defines a need for unmanned surface action groups with the ability not only to navigate and provide deception but also to target and deliver ordnance, including surface-to-air missiles and land-attack missiles. Parts of this mission would involve a capability that does not currently exist—in particular, the ability of platforms to coordinate among themselves for positioning and firing doctrines. Currently, human planners and operators station ships for the best sensor coverage and firing position. This positioning depends on the ability to perceive not only what is best for a particular platform but also what is best for the group as a whole. Machines have a fairly easy time optimizing their own environments. They do not yet have the ability to share an operational and tactical depiction (or *picture*) of the battlespace and coordinate with one another.

The degree to which these platforms ought to be able to fire without human intervention is not so much a technology problem as it is a policy problem. The ability of weapon systems to complete the last link in a kill chain and fire autonomously has been present for decades in the Aegis weapon system and, for that matter, the close-in weapon systems installed on ships throughout the Navy. Giving firing criteria to a weapon system is a relatively simple programming problem. The policy decision revolves around whether doing so could lead to engagements with undesirable consequences that a person exercising contextual judgment would have been able to avoid. In evaluating this issue, it is important to recognize that human judgment is fallible also, and mistaken engagements can occur from failure to follow disciplined threat assessment and reacting from fear rather than following doctrine. Another consideration is that decisions about threat may need to take place in seconds or fractions of seconds. In the absence of context, a machine is more likely than a human to process threats and responses correctly. Ceding speed of command might just be a way of ceding the engagement to the side most willing to rely on speed and reliability.

However, part of the ability of humans to assess a situation more correctly is the capability to share a depiction of the battlespace and coordinate actions. As we noted, autonomous technology is not well advanced in this area. For the future periods that we consider in this report, it may be that force-level firing decisions are best left with human operators, with the understanding that, at some point, the autonomous features of individual systems are likely to be employed as a system shifts to a self-defense mode. The need to keep force-level firing decisions with humans is partly related to the desire to avoid unintended engagements, but it is also meant to ensure the efficient employment of force assets. A system optimizing its own functions might fire at targets that are better targeted by a differently positioned system. Until autonomy reaches the point at which machines understand that and know how to cooperate, humans are likely to remain an essential part of the architecture.

Conclusions

The FFDA depends, in several important respects, on continued improvements in autonomy. Autonomy will be essential to allow continued operations in communication-limited environments, to give ships the ability to coordinate fires without continuously available networks, and to use varieties of stealth to operate in denied areas or complicate the enemy's ability to function in those denied areas.

Following our assessment of the state of autonomous technology, we conclude that the requirements do not appear to be greatly out of line with the arc of technological progress. However, the applications of autonomy considered in the three proposed FFDAs appear to be heavily oriented toward optimization of single-system capability. In some ways, this is a rational approach to technology that we know is better at optimizing the performance of individual platforms than promoting cooperation and a common picture of the battlespace. However, when the military relies on individual platform performance, it makes itself particularly vulnerable to countermeasures. If, for example, an MCM system relies very heavily on identification that is based on pattern recognition, disruption of those patterns by laying mines in ways that the machine has never seen is a straightforward way to defeat it.

This emphasis on single-system optimization opens up numerous vulnerabilities in all of the mission areas we have discussed. In the next chapter, we look at possible applications of autonomy that are oriented more toward simple application of some tasks but with greater capability for multiple systems to coordinate and specialize. The difference in emphasis might lead to different investment priorities, and we want to understand clearly what those differences might be.

CHAPTER FIVE

A Different Direction for Autonomous Systems

Both the systems intended to support current kill chains and the future fleet are effectively intended as direct replacements for manned vehicles. The detect, classify, identify, and engage systems are hosted on a single platform, and the systems are engineered to enable autonomous operations from that single platform. We have seen that there are limitations and challenges associated with developing tightly integrated multifunction vehicles. These limitations and challenges include the following:

- extended timelines imposed by sequential CONOPs
- limited ability to cover a wide area
- complex interactions between subsystems
- potential vulnerability to countermeasures.

These obstacles are, to a degree, inherent in relying on individually complicated systems. The missions are also complicated, and there may be no way to completely eliminate the need for complex interaction. However, that notion suggests that CONOPs used for manned platforms are necessarily the best for autonomous systems, and this seems to be an assumption worth questioning.

The Value of Single-Task Autonomous Systems

In developing our catalog on the state of the art of autonomous technology, we encountered several platforms that are intended to swarm, which implies a level of coordination between small and possibly individually disposable platforms that still organize themselves into a coherent and purposive body. To a degree, the ability of unmanned systems to cooperate with each other has been overstated. Our research indicates that the organization for these systems is generally based on individually programming simple platforms to move in ways that do not interfere with neighboring systems. There is no actual coordination taking place.

This kind of real-time coordination may be more aspirational than actual for the next decade or more. However, the principle of using multiple small platforms carrying out single tasks in a kill chain, and cooperating to the point that they are all addressing a common target, appears to be both feasible with projected technology and a possible way of compressing timelines and putting more systems into the environment. This would be a change in CONOPs, but it would also exploit already accessible autonomous technology. In the next sections, we illustrate this idea for our three mission areas: MCM, denied-area ISR, and operational deception.

Mine Countermeasures

The MCM kill chain as carried out currently and envisioned for the future requires that a single platform carry out a full detect-to-engage sequence. An alternative concept would be to use several different platforms, each performing a different set of tasks in the kill chain sequence. For example, one group of platforms with wide-area sensors could carry out a search and identify the mine-like objects in the field. These platforms would then pass the locating information to platforms with sensors optimized for identification, and these, in turn, would pass the information to expendable neutralization platforms. There is no requirement for the system to approach slowly and then reacquire or for a particular system to continue any contact with a mine-like object. Objects are discovered with the search sensor and passed to other sensors as tracks.

The individual requirements for autonomy within each platform are relatively limited. Numerous platforms could be applied to the minefield problem, which, in turn, could reduce timelines. However, the requirements for sharing information and keeping a common reference picture of the battlespace are considerable. The platforms would require reliable short-range communications across a wide variety of platforms. The platforms would need precise and common navigation. Although each individual platform would not necessarily need the entire tactical picture, a fused picture would have to be available to some unit that is monitoring the overall situation. However, no one platform is unduly complex, and the required collaboration is effectively limited to sharing information. The platforms are not required to exercise anything like judgment.

Denied-Area ISR

Single large platforms have endurance and may have the ability to process some of the information collected on board. However, the platform's sensor range will be limited, and any movement for improved sensor coverage will be at the expense of endurance. Moreover, devising countermeasures against a single large platform may be as simple as nets or other underwater obstacles. Multiple small platforms distributed across a wide area could help provide broader sensor coverage and would not require more than limited relocation. In effect, these platforms could function as a buoy field or deploy-able underwater sonar array. Moreover, they would be less susceptible to individual countermeasures. However, this concept requires at least one unit capable of collating sensor reports, fusing them, and periodically reporting the results to the organization desiring the intelligence. The platform would not need to be as large or complicated as an LDUUV, but it would be larger than the platforms in the distributed field and could indeed be a weak link.

A more significant limitation for single-task systems in a denied-area ISR mission could be that, although the capability of autonomous systems to process and fuse information is already well developed, reliable underwater communication is not. Unless the distributed platforms are capable of passing the information to the processing unit with a high probability that the information will be reliably received, the collected information will be of little or no value. This type of communication is also highly vulnerable to disruption by natural factors and human-made countermeasures. However, despite these limitations, the concept uses feasible applications of autonomy, expands sensor coverage over the program of record, and does not require the development of platforms that understand context or tactical judgment.

Operational Deception: Surface Decoys and Air Defense

An obvious use of smaller single-task platforms is simply as a smart and persistent airborne decoy. These could be spread across a wide area, transmitted on pre-set frequencies to simulate aircraft, and generally create enough noise across the spectrum to complicate enemy location and targeting. These decoys would not be frequency-agile, and they would likely not be able to replicate ship signatures. But they would serve some of the same functions as chaff or short-duration decoys—and for a longer period.

If a more robust anti-air capability were desired, the same concept used for distributed MCM might be effective. Instead of having one platform attempt to carry out a full detect-to-engage sequence, the detecting units, tracking units, and firing units would be dispersed and placed in the locations best suited for their respective tasks. Detection units, for example, could be down a threat axis to enable early detection, while the firing units might be stationed farther away to allow multiple shots as a target passes through the network. This setup would imply the presence of tactical data links and the ability of each unit to recognize common operational and tactical pictures. Unlike underwater networks, however, this kind of network has been in use for decades on manned platforms and would not be difficult to adapt to unmanned units.

The policy issues relating to autonomous engagements might be a more difficult challenge for operational deception mission area than for the others. Engagements will unfold rapidly, and decisions to fire may take place in seconds. The ability to consider context probably will not be a feature of these systems. The risk of engaging a friendly or neutral target will have to be considered and resolved before the systems are put at higher readiness postures.

Future Application Conclusions

We suggest alternatives to the currently envisioned uses of autonomous systems, with particular emphasis on widely distributing relatively simple systems that can perform collaborative tasks. Without question, the applications discussed in this chapter face challenges, but they may be more in line with the promise of autonomy than applications that try to make autonomous systems direct replacements for manned systems.

CHAPTER SIX

Conclusions and Recommendations

As part of this study, we explored two paths to assess the state of technological development in autonomy and warfighting requirements—particularly, how the Navy has applied or intends to apply the autonomous technology that it has. We found that there appears to be neither a short-term commercial explosion of autonomy nor a completely coherent strategy for applying autonomy. It may be that some of the more promising uses of autonomy are not being exploited, and our recommendations pertain directly to how to reorient the Navy's approach.

Conclusions

Using the findings from our analyses in the four areas we studied (the current state of the art of autonomous technology, current kill chains and capabilities, future fleet architecture and its autonomous capabilities, and autonomy in alternative CONOPs), we draw the following conclusions:

- Advances in autonomy have been steady, but the transition to systems capable of reacting to unexpected changes in the environment has not occurred and might not occur for several years.
- The military applications contemplated for unmanned vehicles are unlikely to be developed without substantial investment and development. Reliance on commercial off-the-shelf technology is not likely to support complex military missions.
- The limiting features associated with the most-complex missions are not necessarily associated with autonomy, and it might not be useful to accelerate autonomy while such issues as power generation and storage are still being worked out.
- Under current kill chains and CONOPs, autonomy is generally employed to replicate items in the kill chain exactly as they are carried out by manned systems.
- The alignment between the expected future fleet and possible future uses of autonomy is not close, and our analyses of the building blocks of autonomy (algo-

rithms, platforms, and payloads) suggest that some features desired for the future fleet are unlikely to be reached under the known program of record.
- Policy issues related to autonomous systems applying rules of engagement do exist, and it is unlikely that these systems will be engineered in a way that avoids these issues. Slowing the decisionmaking of autonomous systems by interposing a human in the loop will likely cede a critical time advantage in a high-intensity environment. This delay is a choice and cannot be mitigated by technical improvements.
- Because machine learning depends heavily on pattern recognition, it may be particularly vulnerable to spoofing or misidentification.
- Although development efforts are focused on multifunction, highly complex systems, some of the more promising uses of autonomy might be in using simple systems with limited autonomy for most functions—but adding the capacity for the systems to coordinate with each other. This is particularly the case for missions (such as MCM) in which the current extended timelines are related to the need to have one platform go through the full detect-to-engage sequence. Having a large number of cooperating single-sensor platforms may significantly speed up the timeline.

Recommendations

We recommend, based on the findings described in this report, that the U.S. Navy do the following:

- Revisit assumptions concerning technological progress in autonomy. Our research indicates that the capability for autonomous systems to interpret context and make independent decisions, particularly in a dynamic environment, is not realistic to expect in the short term.
- In systems that require high degrees of autonomy, align the development of autonomy with the development of other capabilities that might be the limiting factors. For example, long-range, high-capacity UUVs are more limited by power generation and storage than by the autonomous features required for the long-range mission.
- Support research that bridges the gap between control theory and machine-learning approaches to autonomy.
- Rather than using autonomy simply to replicate the kill chains of manned systems, use the unique features of autonomy to enable new CONOPs. It seems particularly promising to employ simple but numerous systems carrying different kinds of sensors. This capability will require the development of a system capable

of fusing multiple inputs to create a common operational picture to serve as a reference for the individual operating systems.

- Reevaluate FFDA requirements in light of the state of the art of autonomous technology. Some features of the proposed architectures are more aspirational than likely.
- Accept the reality that autonomous systems will need to make engagement decisions. While the policy ramifications of autonomous systems are interesting, they may be taking up more of the debate than is essential; the central issue should be about making the best use of autonomy. If a system is being employed in an environment where there are a large number of potential targets and a very compressed timeline for making engagement decisions, a desire to retain human intervention except for general ability to override may not be compatible with the tactical picture. There is, moreover, no obvious way to engineer a way out of this dilemma. Rather than the Navy attempting to inject a human into this loop who makes the decision to engage a target, the Navy should accept that modern weapon system timelines simply preclude such intervention. If engagement decisions are to be effective, development efforts should not focus on finding ways around the fact that machines will likely have to make such decisions.
- Develop a mechanism that allows humans to periodically assess whether an autonomous system is misinterpreting its environment. Because autonomous systems have not, to date, learned adaptive behavior or an ability to interpret context, they appear to be particularly vulnerable to countermeasures that alter some feature of their expected environment. A human operator's ability to recognize that a system is misinterpreting some part of the environment may be a particularly important oversight mechanism. Such oversight also implies the ability to access the machine's learning capability to provide images or other means of recognizing something unexpected. So, human operator intervention will occur less when an engagement decision is being made and more when it is apparent that the system is behaving in ways that indicate misunderstanding of the events and conditions around it. An autonomous system must possess an interface that allows periodic assessment of what it senses and how it is reacting to what it senses. The degree to which this is done will depend, in part, on the system operating undetected by adversaries.
- Critically evaluate the viability of complex multimission platforms, and consider emphasizing simple but cooperating platforms.

APPENDIX A

In-Depth Analysis: Algorithms

In this appendix, we review the academic and technical literature on the algorithms that underpin vehicle autonomy, organized by task family as defined in Chapter Two. We chose not to restrict our search to naval vehicles; instead, we also reviewed the literature on unmanned vehicles operating on land and in the air. The reason for our choice was twofold: (1) research focusing on these application domains is broadly available and (2) the ideas behind many of the algorithms are domain-agnostic. The survey presented here provided input for our assessment of the general trends and remaining gaps in autonomous algorithms, as well as for our recommendations for future research directions in algorithms, as presented in Chapter Two.

Task Families for a Single Unmanned Vehicle

Navigation

The navigation task family encompasses algorithms that allow an unmanned vehicle to accurately pinpoint its position (*localization*); move from a point of origin to a destination (*path-planning*, *reference-tracking*, and *path-following*); do so safely by avoiding obstacles and collisions (*collision avoidance*); and maintain its absolute position relative to a point of reference, possibly another vehicle, in the face of ocean currents (*station-keeping*). Navigation in an underwater environment brings a unique set of challenges owing to the limited communication environment and susceptibility to currents. Approaches for localization may be anchor-based, anchor-free, or a combination, such as in terrain-relative navigation.[1] Path-planning approaches include genetic algorithms, particle swarm optimization, and a broad array of sampling-based approaches and tree planners (e.g., expansive-space tree planning, rapidly exploring random tree planning, and prerequisite tree planning).[2] Various techniques are used for the low-level control of the vehicle to ensure path-following and reference-tracking. Likewise, the research

---

[1] S. Dektor and S. Rock, "Improving Robustness of Terrain-Relative Navigation for AUVs in Regions with Flat Terrain," *2012 IEEE/OES Autonomous Underwater Vehicles (AUV)*, 2012.

[2] M. Elbanhawi and M. Simic, "Sampling-Based Robot Motion Planning: A Review," *IEEE Access*, Vol. 2, 2014.

on obstacle and collision avoidance is rich, and the current trend is toward reactive algorithms for real-time obstacle avoidance.

Note that sensor fusion is an important technical challenge underlying several of these subtasks; such fusion allows one sensor to combine input from multiple, often heterogenous sensors to improve estimation and, thus, performance.[3]

Exploration

The exploration task family encompasses algorithms that enable an autonomous vehicle to carry out area surveillance by monitoring its environment to detect, identify, and track objects of interest (*object recognition*, *detection and cuing*, and *tracking*); generate a map and localize itself on that map *(SLAM*); and explore its surroundings for research and study purposes (*oceanography*). The demarcation between the navigation and exploration task families is not sharp; for example, processing of sensor signals and filtering and estimation using sensor data and sensor fusion are needed for both navigation and exploration.

Various approaches, including feature detection and convolutional neural network–based classification algorithms, have seen successes in object recognition.[4] Approaches for combining classifiers have also been pursued.[5] Approaches for 3D object detection and passive–active handoff allow detection and cuing of objects.[6] Correlation filters and various other approaches have been used for multi-object track-

---

[3] Sajjad Safari, Faridoon Shabani, and Dan Simon, "Multirate Multisensor Data Fusion for Linear Systems Using Kalman Filters and a Neural Network," *Aerospace Science and Technology*, Vol. 39, December 2014.

[4] On feature detection, see Juan Luo and Gwun Oubong, "A Comparison of SIFT, PCA-SIFT and SURF," *International Journal of Image Processing*, Vol. 3, No. 4, 2009; and Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua, "LIFT: Learned Invariant Feature Transform," ArXiv, July 29, 2016. On convolutional neural network–based classification, see Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *NIPS '12: Proceedings of the 25th International Conference on Neural Information Processing Systems*, Vol. 1, New York: Curran Associates Inc., 2012.

[5] Saso Džeroski and Bernard Ženko, "Is Combining Classifiers with Stacking Better than Selecting the Best One?" *Machine Learning*, Vol. 54, No. 3, March 1, 2004.

[6] On 3D object detection, see Y. Guo, M. Bennamoun, F. Sohel, M. Lu, and J. Wan, "3D Object Recognition in Cluttered Scenes with Local Surface Features: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 36, No. 11, 2014. On passive-active handoff, see C. Wakayama, D. Grimmett, and R. Ricks, "Active Multistatic Track Initiation Cued by Passive Acoustic Detection," *2012 15th International Conference on Information Fusion*, 2012.

ing, with promising developments in multitracker setups.[7] Also widely addressed are SLAM problems, as well as various uses and approaches for oceanography.[8]

Note that, by choice, much of the literature we cite focuses on the tasks rather than the specific role of autonomous vehicles in carrying out the task. Also worth noting is the vast array of competitions and benchmarks that focus on object recognition, detection, and tracking, as well as sensor fusion, which challenge results often published subsequent to the competitions.[9] Finally, because of the significant use of machine-learning techniques in this task family, it is apt to point out the vulnerability of supervised-learning algorithms to adversarial examples, as well as current research on developing potential mitigations.[10]

---

[7] On correlation filters, see David S. Bolme, J. Ross Beveridge, Bruce A. Draper, and Yui Man Lui, "Visual Object Tracking Using Adaptive Correlation Filters," *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010; and J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-Speed Tracking with Kernelized Correlation Filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 37, No. 3, 2015. On multi-object tracking, see Wenhan Luo, Junliang Xing, Anton Milan, Xiaoqing Zhang, Wei Liu, Xiaowei Zhao, and Tae-Kyun Kim, "Multiple Object Tracking: A Literature Review," ArXiv, May 22, 2017; and Samuel Scheidegger, Joachim Benjaminsson, Emil Rosenberg, Amrit Krishnan, and Karl Granstrom, "Mono-Camera 3D Multi-Object Tracking Using Deep Learning Detections and PMBM Filtering," ArXiv, February 27, 2018. On multitracker setups, see Isabelle Leang, Stéphane Herbin, Benoît Girard, and Jacques Droulez, "On-Line Fusion of Trackers for Single-Object Tracking," *Pattern Recognition*, Vol. 74, 2017.

[8] On SLAM, see Ling Chen, Sen Wang, Klaus McDonald-Maier, and Huosheng Hu, "Towards Autonomous Localization and Mapping of AUVs: A Survey," *International Journal of Intelligent Unmanned Systems*, Vol. 1, No. 2, 2013. On oceanography, see Russell B. Wynn, Veerle A. I. Huvenne, Timothy P. Le Bas, Bramley J. Murton, Douglas P. Connelly, Brian J. Bett, Henry A. Ruhl, Kirsty J. Morris, Jeffrey Peakall, Daniel R. Parsons, Esther J. Sumner, Stephen E. Darby, Robert M. Dorrell, and James E. Hunt, "Autonomous Underwater Vehicles (AUVs): Their Past, Present and Future Contributions to the Advancement of Marine Geoscience," *Marine Geology*, Vol. 352, June 2014; and Francis D. Lagor, Kayo Ide, and Derek A. Paley, "Incorporating Prior Knowledge in Observability-Based Path Planning for Ocean Sampling," *Systems and Control Letters*, Vol. 97, November 2016.

[9] Matej Kristan, Ales Leonardis, Jiri Matas, Michael Felsberg, et al., "The Visual Object Tracking VOT2017 Challenge Results," *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*, 2017. A sampling of such challenges and competitions includes ImageNet, "ImageNet Large Scale Visual Recognition Challenge (ILSVRC)," webpage, 2015; Common Objects in Context, homepage, undated; Visual Object Tracking Challenge, homepage, undated; Multiple Object Tracking Benchmark, homepage, undated; and Institute of Electrical and Electronics Engineers (IEEE), "Data Fusion Contest," webpage, undated.

[10] On vulnerabilities, see Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus, "Intriguing Properties of Neural Networks," ArXiv, February 19, 2014; Ian Goodfellow, Jonathon Shlens, and Christian Szegedy, "Explaining and Harnessing Adversarial Examples," ArXiv, March 20, 2015; and Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami, "The Limitations of Deep Learning in Adversarial Settings," ArXiv, November 24, 2015. On mitigations, see Nicholas Carlini and David Wagner, "Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods," in Association for Computer Machinery, *AISec '17: Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, New York, 2017.

Effect Delivery

The effect delivery task family encompasses algorithms allowing the unmanned vehicle to deliver effects, including kinetic and nonkinetic strike options, and manipulate or otherwise affect its environment and the objects or systems in it. This is the task family for which we found the least academic work, as is perhaps to be expected. The literature that does exist is focused primarily on higher-level command-and-control tasks, such as dynamic allocation of targets and shooters.

Countermeasures

The countermeasures task family encompasses algorithms to evade threats, both kinetic and nonkinetic; mitigate against and recover from degradation of physical or software parts, resulting from either system failures or targeted attacks; and deceive adversaries.

The research on evasive strategies remains somewhat limited. It includes work on pursuit-evasion games and so-called angel–devil problems, although the trend toward taking into account practical limitations, such as motion constraints, is encouraging.[11] Perhaps unsurprisingly, we were not able to find deception-related work in the public sphere. In contrast, there is an extensive history of work on fault detection and fault tolerance, as well as nascent research looking at resilience in cyber–physical systems, in terms of both detection and the design of safety controllers.[12]

Resource Management

The resource management task family encompasses algorithms for managing onboard resources, primarily power.[13] Approaches include algorithms for optimizing the use of multiple onboard power sources, as well as approaches for optimizing tasks (e.g., path-planning) to minimize energy use.[14]

---

[11] On pursuit-evasion games, see R. Vidal, O. Shakernia, H. J. Kim, D. H. Shim, and S. Sastry, "Probabilistic Pursuit-Evasion Games: Theory, Implementation, and Experimental Evaluation," *IEEE Transactions on Robotics and Automation*, Vol. 18, No. 5, 2002. On angel-devil problems, John H. Conway, "The Angel Problem," in Richard Nowakowski, ed., *Games of No Chance*, Vol. 29, Cambridge, UK: Cambridge University Press, 1996. On practical limitations, see William Lewis Scott and Naomi Ehrich Leonard, "Optimal Evasive Strategies for Multiple Interacting Agents with Motion Constraints," *Automatica*, Vol. 94, August 2018.

[12] M. K. Yoon, B. Liu, N. Hovakimyan, and L. Sha, "VirtualDrone: Virtual Sensing, Actuation, and Communication for Attack-Resilient Unmanned Aerial Systems," *2017 ACM/IEEE 8th International Conference on Cyber-Physical Systems (ICCPS)*, 2017.

[13] Eleftherios Amoiralis, Marina Tsili, Vassilios Spathopoulos, and A. Hatziefremidis, "Energy Efficiency Optimization in UAVs: A Review," *Materials Science Forum*, Vol. 792, August 2014.

[14] On optimizing multiple power sources, see Byeong Gyu Gang and Sejin Kwon, "Design of an Energy Management Technique for High Endurance Unmanned Aerial Vehicles Powered by Fuel and Solar Cell Systems," *International Journal of Hydrogen Energy*, Vol. 43, No. 20, May 2018. On optimizing tasks, see Adam Kaplan, *Path Planning and Energy Management of Solar-Powered Unmanned Ground Vehicles*, thesis, Ames, Iowa: Iowa State University, 2015.

Task Families for Teams of Multiple Vehicles

Several recent books and survey papers, and the references therein, provide a good overview of the existing literature on unmanned multi-agent teaming.[15] Although a good proportion of the published literature focuses on teams of UAVs and unmanned ground vehicles (UGVs), the fundamental ideas underpinning the algorithms, if not the algorithms themselves, are largely portable to unmanned underwater and surface vehicles.

Swarming

The swarming task family encompasses algorithms that focus on developing vehicle motion control laws, thereby enabling a team of unmanned vehicles to jointly navigate and move in a coordinated manner. A variety of swarming scenarios have been considered, including flocking, rendezvous, synchronization, and formation control (including leader–follower formations).[16]

Cooperation

The cooperation task family encompasses algorithms enabling a team of vehicles to jointly undertake endeavors beyond swarming. These can include, for example, distributed sensing or cooperative goals (e.g., cooperative learning), wide-area surveillance, target search and tracking, path coverage, multi-agent SLAM, and coordinated

---

[15] Venkatesh Saligrama, ed., *Networked Sensing Information and Control*, 1st ed., Boston, Mass.: Springer Science and Business Media, 2008; Y. Wang, E. Garcia, F. Zhang, and D. Casbeer, *Cooperative Control of Multi-Agent Systems: Theory and Applications*, New York: John Wiley & Sons, 2017; Danielle C. Tarraf, ed., *Control of Cyber-Physical Systems*, Cham, Switzerland: Springer International Publishing, March 2013; Jeff S. Shamma, ed., *Cooperative Control of Distributed Multi-Agent Systems*, New York: Wiley-Interscience, 2008; and Jorge Cortes and Magnus Egerstedt, "Coordinated Control of Multi-Robot Systems: A Survey," *SICE Journal of Control, Measurement, and System Integration*, Vol. 10, No. 6, 2017.

[16] On flocking, see Silvia Mastellone, Dušan M. Stipanović, Christopher R. Graunke, Koji A. Intlekofer, and Mark W. Spong, "Formation Control and Collision Avoidance for Multi-Agent Non-Holonomic Systems: Theory and Experiments," *International Journal of Robotics Research*, Vol. 27, No. 1, 2008; and R. Olfati-Saber, "Flocking for Multi-Agent Dynamic Systems: Algorithms and Theory," *IEEE Transactions on Automatic Control*, Vol. 51, No. 3, 2006. On rendezvous, see Feng Xiao, Long Wang, and Tongwen Chen, "Connectivity Preservation for Multi-Agent Rendezvous with Link Failure," *Automatica*, Vol. 48, No. 1, January 2012; Housheng Su, Xiaofan Wang, and Guanrong Chen, "Rendezvous of Multiple Mobile Agents with Preserved Network Connectivity," *Systems and Control Letters*, Vol. 59, No. 5, May 2010; and Michael Ouimet and Jorge Cortés, "Robust Coordinated Rendezvous of Depth-Actuated Drifters in Ocean Internal Waves," *Automatica*, Vol. 69, July 2016. On synchronization, see D. Sun, C. Wang, W. Shang, and G. Feng, "A Synchronization Approach to Trajectory Tracking of Multiple Mobile Robots While Maintaining Time-Varying Formations," *IEEE Transactions on Robotics*, Vol. 25, No. 5, 2009; and H. Zhang, F. L. Lewis, and A. Das, "Optimal Design for Synchronization of Cooperative Systems: State Feedback, Observer and Output Feedback," *IEEE Transactions on Automatic Control*, Vol. 56, No. 8, 2011. On formation control, see J. A. Fax and R. M. Murray, "Information Flow and Cooperative Control of Vehicle Formations," *IEEE Transactions on Automatic Control*, Vol. 49, No. 9, 2004.

deception.[17] Additionally, basic research on, for instance, multiple-viewpoint recognition and localization may be extended to the multi-agent setting, even though it is not specific to multi-agent cooperation.[18]

The general premise of multi-agent cooperation is that, although individual agents (unmanned vehicles) may have limited capabilities in terms of power, sensing, and communication, their coordinated collective use provides potential advantages, such as large-scale spatial distribution and inherent robustness owing to greater numbers and the lack of single points of failure. The manner in which these teams are designed has important implications in practice, as we describe next.

Centralized, Decentralized, and Distributed Schemes

At a high level, the algorithmic schemes used for decisionmaking in multi-agent teams (particularly for swarming and cooperation) fall into the following three categories:

1. Centralized schemes, in which a single decisionmaker has a global view of the state of every agent and of the team's global objective. This central decisionmaker thus designs the control actions of each agent and relays them for execution.
2. Decentralized schemes, in which every agent makes its own decisions based on its own local knowledge without knowledge of the other agents' states or actions. The global behavior of the team then emerges as a result of the individual decisions.
3. Distributed schemes, in which some amount of computation is carried out by each of the agents and some amount of communication occurs between an agent and its neighbors, however they are defined. Distributed schemes run the gamut

---

[17] On cooperative learning, see Naomi Ehrich Leonard and Alex Olshevsky, "Cooperative Learning in Multiagent Systems from Intermittent Measurements," *SIAM Journal on Control and Optimization*, Vol. 53, No. 1, 2015. On wide-area surveillance, see Zhijun Tang and U. Ozguner, "Motion Planning for Multitarget Surveillance with Mobile Sensor Agents," *IEEE Transactions on Robotics*, Vol. 21, No. 5, 2005. On target search and tracking, see Jin Yan, Liao Yan, A. A. Minai, and M. M. Polycarpou, "Balancing Search and Target Response in Cooperative Unmanned Aerial Vehicle (UAV) Teams," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, Vol. 36, No. 3, 2005. On path coverage, see M. Schwager, J. J. Slotine, and D. Rus, "Decentralized, Adaptive Control for Coverage with Networked Robots," *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007. On multi-agent SLAM, see N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Decentralized Active Information Acquisition: Theory and Application to Multi-Robot SLAM," *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015. On coordinated deception, see Z. C. Zhao, X. S. Wang, and S. P. Xiao, "Cooperative Deception Jamming Against Radar Network Using a Team of UAVs," *2009 IET International Radar Conference*, 2009.

[18] Scott Helmer, David Meger, Marius Muja, James J. Little, and David G. Lowe, "Multiple Viewpoint Recognition and Localization," in Ron Kimmel, Reinhard Klette, and Akihiro Suimoto, eds., *Computer Vision—ACCV 2010: 10th Asian Conference on Computer Vision*, Part I, Berlin: Springer-Verlag, 2011.

> from (1) bottom-up design setups in which each agent makes its own decisions based on local (its own and nearest neighbor's) information to optimize a local objective to (2) top-down design setups in which a global decisionmaker keeps track of a global objective that it decomposes and tasks to the agents, resulting in distributed computation toward a global objective.[19]

The advantages of centralized schemes are that, in principle, the central decisionmaker can globally optimize the actions of the team because it has global knowledge of the team and objectives. The disadvantages of such schemes are the communication bandwidth required, the complexity of the decisionmaking (which scales with the number of agents), and the potential for a single point of failure.

The advantages of centralized and distributed schemes are their potential scalability and the reduced need for communication bandwidth.[20] Their disadvantages are the difficulties of designing these schemes to ensure global objectives and the need to maintain connectivity of the underlying communication graph; the exact connectivity requirements are dependent on the specifics of the control scheme and its associated graph topology.

Validation and Verification

The problem of V&V remains largely an open one. Where control approaches are used, the design approaches are generally well grounded in theory. However, their implementation in software code is not an exact process, and the models used generally diverge from reality.[21] Both factors open the door to error, with potentially serious consequences, and require extensive V&V. While some V&V approaches exist, their scalability remains limited.[22] On the other hand, where machine-learning approaches are used—and despite rapid advances in algorithms and their demonstrated use to solve practical problems—there has been little progress made in explanatory principles. Thus, a much-needed fundamental understanding of why these algorithms succeed

---

[19] These categories are provided as general guidelines, with the caveat that the demarcation between one category and the next is not entirely clean.

[20] However, the current state of the art is far from realizing the scalability promises of these schemes.

[21] Approaches to turn this implementation into an exact process are currently being pursued. See, for example, Timothy Wang, Romain Jobredeaux, Marc Pantel, Pierre-Loic Garoche, Eric Feron, and Didier Henrion, "Credible Autocoding of Convex Optimization Algorithms," *Optimization and Engineering*, Vol. 17, No. 4, December 2016.

[22] See E. M. Clarke, T. A. Henzinger, and H. Veith, *Handbook of Model Checking*, Cham, Switzerland: Springer International Publishing, 2016; and T. Wongpiromsarn, U. Topcu, and A. Lamperski, "Automata Theory Meets Barrier Certificates: Temporal Logic Verification of Nonlinear Systems," *IEEE Transactions on Automatic Control*, Vol. 61, No. 11, 2016.

and fail when they do remains largely lacking. In addition, there has been much recent interest in explainable artificial intelligence and artificial intelligence safety, to include robustness to distributional shifts.[23]

---

[23] Distributional shifts occur when the testing distributions differ from the training distributions, potentially leading the machine-learning algorithms to perform poorly but for that performance not to be recognized as poor. On artificial intelligence, see the symposium debate at the 2017 Conference on Neural Information Processing Systems (Interpretable Machine Learning, "2V2 Debate: Caruana, Simard vs. Weinberger, LeCun, Interpretable ML Symposium, NIPS 2017," video, December 2017). On artificial intelligence safety, see Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané, "Concrete Problems in AI Safety," ArXiv, June 2016.

APPENDIX B

In-Depth Analysis: Payloads

This appendix contains a more detailed discussion of the unmanned maritime payload technologies identified during our literature review, as well as a more in-depth discussion of design trade-offs associated with these technologies.

Navigation

Inertial

Inertial navigation is a technique used by UUVs and USVs that leverages data provided by sensors to determine positioning and heading of a platform. Various systems can be used to provide inertial navigation and often are combined to reduce navigational errors, which can be exacerbated over time. INSs are one payload commonly used aboard UUVs and USVs. INSs contain sensors, such as gyroscopes and accelerometers,[1] that collect and supply data used to aid navigation.[2] High levels of precision with an INS can be achieved only for short periods. To circumvent this issue, information from external payloads can be fused with INS data to provide more-accurate measurements, although techniques have been developed to reduce INS errors without relying on additional sensors.[3] External aids that can be paired with an INS include GPS, celestial navigation systems (CNSs), and doppler velocity logs (DVLs). Fusing data from multiple external aids—for instance, an INS, CNS, and DVL configuration—is a proposed technique in which both the CNS and DVL provide reference points that reduce accu-

---

[1] These types of sensors are collectively referred to as the *inertial measurement unit*.

[2] Oliver J. Woodman, *An Introduction to Inertial Navigation*, Cambridge, UK: University of Cambridge Computer Laboratory, Technical Report No. 696, August 2007.

[3] E. Akeila, Z. Salcic, and A. Swain, "Reducing Low-Cost INS Error Accumulation in Distance Estimation Using Self-Resetting," *IEEE Transactions on Instrumentation and Measurement*, Vol. 63, No. 1, 2014.

mulating errors in INSs.[4] Similarly, as discussed in further detail later, acoustic navigation systems can also be configured to reduce INS drift.[5]

GPS

Because maintaining a GPS signal requires satellite coverage, GPS can provide 3D positioning to undersea and surface vehicles operating in shallow domains.[6] As described in further detail in the section on radar sensor systems, the radio frequency signals used by GPS are attenuated by water and can be jammed by adversarial countermeasures. Although underwater vehicles cannot directly use GPS to obtain positioning without surfacing, surface vessels (such as buoys or USVs) can use acoustic frameworks to track UUVs and fuse those data with GPS information to provide highly precise localization to the UUV.[7]

Acoustic

Long baseline, short baseline, and ultrashort baseline acoustic sensors each use single or multiple transponders operating at different frequencies to help UUVs navigate. The accuracy of these systems increases as the frequency of the signals increases (ultrashort baseline has the highest frequency, long baseline the lowest); however, range is reduced at higher frequencies. Each system is often mounted on ship hulls; long baseline sensors are often implemented on surface buoys, the sea floor, or ice formations.[8] The historical reliance on these systems for navigation is being shifted toward systems that can be deployed rapidly with little required infrastructure.[9]

A DVL is a system of transducers that measure velocity and can be outfitted on either UUVs or USVs.[10] While DVLs can support underwater operations, they have

---

[4] Qiuying Wang, Xufei Cui, Yibing Li, and Fang Ye, "Performance Enhancement of a USV INS/CNS/DVL Integration Navigation System Based on an Adaptive Information Sharing Factor Federated Filter," *Sensors*, Vol. 17, No. 2, 2017.

[5] James C. Kinsey, Ryan M. Eustice, and Louis L. Whitcomb, "A Survey of Underwater Vehicle Navigation: Recent Advances and New Challenges," *IFAC Conference of Manoeuvering and Control of Marine Craft*, 2006.

[6] Maritime surface vessels commonly use Differential GPS to provide positioning, navigation, and timing information. Differential GPS uses a ground station as a reference to correct for measurement errors. See James R. Clynch, "A Short Overview of Differential GPS," Naval Postgraduate School, Department of Oceanography, December 2001.

[7] UUVs can leverage USV-provided GPS connectivity to navigate. See A. Vasilijević, D. Nađ, F. Mandić, N. Mišković, and Z. Vukić, "Coordinated Navigation of Surface and Underwater Marine Robotic Vehicles for Ocean Sampling and Environmental Monitoring," *IEEE/ASME Transactions on Mechatronics*, Vol. 22, No. 3, 2017.

[8] Kinsey, Eustice, and Whitcomb, 2006.

[9] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV Navigation and Localization: A Review," *IEEE Journal of Oceanic Engineering*, Vol. 39, No. 1, 2014.

[10] J. Snyder, "Doppler Velocity Log (DVL) Navigation for Observation-Class ROVs," *OCEANS 2010 MTS/IEEE Seattle*, 2010.

drifting errors that are similar to INS's errors, leading to significant navigational inaccuracies over time.[11] Platforms that correct INS drifting errors with DVL measurements must operate in close vicinity to the seafloor to establish a bottom lock,[12] and this requirement renders such platforms impractical aboard USVs operating in deeper waters.

Acoustic modems can be hosted aboard UUVs and USVs and can enable collaborative navigation with other vessels. Modems also obviate the need for fixed networks (fundamental for long baseline, short baseline, and ultrashort baseline systems), which can extend mission range.[13] Cooperative teams of UUVs can be heterogeneous or homogenous; heterogeneous teams tend to have a few vessels with high-resolution navigational payloads equipped to provide accurate positional data to other UUVs in the swarm.[14] As discussed later, sources of error in acoustic navigation and communication can stem from several sources, including the physical limitations of the speed of sound in water and fluctuations in water temperature and density.[15]

Ranging and imaging sonar systems can also be leveraged to aid maritime vessel navigation by detecting features (e.g., the seabed, objects) that are then used as reference points.[16]

Radar

Radar imagery can be combined with sea charts to determine the positioning of a surface vessel. Comparing radar imagery with sea charts creates a real-time visualization of the environment, which is not always possible with satellite imagery because it is not continuously updated; for instance, coastlines change depending on the tide and might look different at any given time. Furthermore, using radar to enhance positioning information is not viable in deep-sea applications because land formations (i.e., coastlines) are used as reference points.[17] Radar systems are discussed in further detail in the section on sensors later in this appendix.

---

[11] Paull et al., 2014.

[12] L. Medagoda, J. C. Kinsey, and M. Eilders, "Autonomous Underwater Vehicle Localization in a Spatiotemporally Varying Water Current Field," *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015.

[13] Paull et al., 2014.

[14] The vessels with superior navigational capabilities are referred to as *communication and navigation aids*. See Gao Rui and M. Chitre, "Cooperative Positioning Using Range-Only Measurements Between Two AUVs," *OCEANS '10 IEEE Sydney*, 2010.

[15] Kinsey, Eustice, and Whitcomb, 2006.

[16] Paull et al., 2014. Ranging sonars include echosounders, profilers, and multibeam sonars. Imaging sonars include sidescan, forward-looking, and synthetic-aperture sonars.

[17] H. Ma, E. Smart, A. Ahmed, and D. Brown, "Radar Image-Based Positioning for USV Under GPS Denial Environment," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 19, No. 1, 2018.

Depth

Depth can be calculated by measuring ambient pressures in the water column. Pressure sensors provide UUVs with instantaneous depth measurements, which can safeguard the vehicle from operating at depths that would damage its functionality. Strain gauges are one type of sensor that measures pressure by observing the strain exerted on a metal alloy. These measurements can be accurate to within 0.1 percent; calibration is needed for higher-resolution results. Quartz crystals can also be employed to determine depth to an accuracy of 0.01 percent; their resonant frequencies are correlated to ocean pressures.[18]

Orientation

Compasses and other payloads, such as magnetic, roll-and-pitch, and angular-rate sensors, can be employed to determine platform orientation. Gyrocompasses, a specific type of compass traditionally used on large surface vessels, have seen improvements that enable them to be employed on smaller military craft. INSs typically contain gyroscopes to provide north-seeking capabilities. Magnetic sensors can provide a low-cost solution to determine vehicle heading but are highly susceptible to systemic errors.[19] Roll-and-pitch sensors can be accelerometers, pendulum tilt sensors, or fluid-level sensors. The purpose of these sensors is to determine a platform's orientation by referencing gravitational forces. Low-cost sensors are susceptible to performance degradation when surges in acceleration are observed, although higher-cost sensors are designed to mitigate measurement errors. Angular-rate information is ascertained through gyroscopes. High-fidelity gyroscopes, such as optical gyroscopes, are costly and have relatively large power requirements, thus inhibiting their feasibility on smaller UUVs.[20]

Light and Optical

Images from stereo and monocular cameras can be used to aid the navigation and positioning of maritime vessels. These systems are more effective as a navigational aid in feature-rich environments. For example, CNSs use the stars as reference points to navigate a platform. These electro-optical systems can be used aboard USVs but are not capable of providing reliable measurements to UUVs because of the light attenuation in the water column. Poor weather conditions, such as dense fog or cloudiness, can inhibit a CNS from delivering precise positioning information. Seawater also affects camera performance in underwater operations because it limits range, causes scatter-

---

[18] Kinsey, Eustice, and Whitcomb, 2006.

[19] Errors can stem from the orientation of the sensor or magnetic field disturbances induced by the platform or locale.

[20] Kinsey, Eustice, and Whitcomb, 2006.

ing, and presents a low-light environment. Applications suited to subsurface camera–aided navigation include pipeline tracking and hull inspection.[21]

Sensors

Sonar

Sonar has historically been a major player in the maritime sensing domain and remains so for unmanned vehicles. The term *sonar* is a combination of the words *sound*, *navigation*, and *ranging*, and it refers to the use of underwater acoustic signals to navigate or characterize an area. Sonar systems either listen for externally generated acoustic signals or transmit their own acoustic signal pulses toward an area of interest and listen for reflected return signals from targets.[22] For systems that actively generate sensing pulses, distortions to the reflected pulse are used to calculate parameters, such as the distance from the target and the target's relative speed compared with the sonar source.[23] Such information may be used for navigation; target identification; target tracking; or, with a more advanced system, 3D imaging of a target.

Active and passive sonar systems are best suited for different applications. The active source pulse generated by active sonar systems requires some amount of electric power to generate; this should be a design consideration for power-limited systems.[24] Active sonar power requirements scale with desired system range; that is, pulses must be sufficiently powerful to travel twice the distance from the source to the target—once from the source to the target and once as a reflection from the target to the source. Active sonar systems may also require high power if attempting to detect features below absorbing layers of material—notably, objects like mines buried below the ocean floor or submarines that may be coated with sound-absorbing material to reduce sonar reflections. Generating an active sonar pulse signal is an overt action that would alert others nearby to the presence of the active sonar. Such disclosure is not conducive to covert operations.[25] Active sonar is used in a variety of Navy applications today, including MCM operations.

Although passive sonar systems are more covert than active systems, they are less useful for navigation purposes because their receivers rely on external source signals

---

[21] Paull et al., 2014.

[22] Passive and active sonars, respectively.

[23] Craig M. Payne, *Principles of Naval Weapon Systems*, Annapolis, Md.: Naval Institute Press, January 2010.

[24] For instance, the Navy should weigh the power needs when designing battery-powered UUVs with limited onboard power capacity and multiple other important system power demands, including data processing, navigation, and propulsion. Exact sonar system power requirements vary widely and depend on such parameters as sonar range, return signal fidelity, and system frequency.

[25] Payne, 2010.

that are not generated by key navigation obstacles, including undersea geological features. To maximize sensitivity and collection of ambient acoustic signatures, passive sonar systems may be configured as towed arrays. Passive towed-array sonars are an arrangement of sonar elements pulled behind a ship or submarine to listen for a broad spectrum of potential targets. Towed arrays are not constrained by the ship's or submarine's hull shape and may therefore be more sensitive than size-limited systems.[26] Passive sonar is commonly used in various Navy ASW applications today.

Various configurations of active sonar systems can enable improved sensing capability. The most basic active sonar, the single-beam sonar, relies on a transducer producing a single sound pulse in the direction of interest and waiting for the initial return signal to determine distance, relative speed (compared with the sonar source), and other parameters about the target. Though simple and relatively inexpensive, single-beam sonar systems are limited because their viewing area is limited. Multibeam sonars use constructive and destructive interference of multiple collocated sonar array elements to improve beam width and range relative to single-beam sonars. These wider beams with improved range allow higher resolution and larger coverage areas and, therefore, improved targeting and bathymetric fidelity. Thus, in general for all sonar systems, the broader the search area, the less capable the system is of developing high-resolution data on a single target point. And the more array elements that the multibeam system uses, the higher the fidelity of the returned signal; however, having more sensing elements requires more energy, requires more space aboard the vehicle, and is more expensive than variants with fewer elements.[27]

A modification to a multibeam sonar system enables a system known as a *sidescan sonar*. A sidescan sonar is a type of imaging sonar that uses pairs of hydrophones directed toward opposite sides of a ship or submarine to improve the fidelity of received acoustic pulse data.[28] In addition to improving the understanding of received pulse directivity, sidescan sonars listen for continuous received pulse returns. The temporal element of sidescan sonar systems requires that the sensing elements be traveling through the water at a constant speed. By recording length and strength of received pulses, an understanding of seafloor geometry and constituent material may be developed through data processing.[29]

---

[26] Payne, 2010. Passive sonar systems may also be configured as stationary systems. The U.S. Navy's Sound Surveillance System, a worldwide array of hydrophones connected to land-based facilities via communication cables, enabled the Navy to gather undersea acoustic data for many years. See Edward C. Whitman, "SOSUS: The 'Secret Weapon' of Undersea Surveillance," *Undersea Warfare*, 2005.

[27] L-3 Communications SeaBeam Instruments, *Multibeam Sonar Theory of Operation*, East Walpole, Mass., 2000.

[28] The opposite sides of the ship or submarine that are targeted are either integral to the vessel or towed behind the vessel, as in a towed array.

[29] L-3 Communications SeaBeam Instruments, 2000.

Sidescan sonars are known as a type of two-dimensional imaging sonar in that they may generate a more complex image of an ocean floor than other, more basic sonar systems can. Although this level of data fidelity is beneficial, its complexity requires several sensing elements, onboard processing capability, and collaboration between vehicle propulsion and sensing systems, and it is more costly than simpler, single-element sonar systems. Higher-frequency sonar systems with even more-complex onboard processing may even generate 3D images. This level of complexity and cost would be challenging to incorporate into certain platforms but, unlike other optical sensors, is capable of developing high-resolution images of areas in water conditions with low optical visibility. Notably, of the sonar payloads carried aboard unmanned vehicles identified in our catalog, multibeam sonars were the most common, and imaging sonars were identified in a few cases.

Other notable sonar configurations include bistatic or multistatic sonar and synthetic-aperture sonar. The active sonar systems described earlier use a collocated transmitter and receiver pair to measure underwater acoustic signatures. Bistatic or multistatic sonar systems use separately located transmitters and receivers to reduce transmission losses and errors associated with self-generated transmit signal noise.[30] Such an arrangement could also be used to reduce the detectability of the receiving vehicles, which would not need to transmit an active pulse to collect acoustic data. Synthetic-aperture sonar, like sidescan sonar, relies on relative motion between the active transmitter and targets. Using multiple active pings toward a target area and data analysis of returns, synthetic-aperture sonar can produce high-resolution images of a target area with relatively few sensing elements.[31] However, synthetic-aperture sonar has a very strict speed limit; it can move forward only half an array length per ping.

Radar

The term *radar* is a combination of the words *radio*, *detection*, and *ranging* and refers to the use of transmitted electromagnetic waves and received reflections off a target to assist with navigation or targeting. Similar to active sonar systems, radar systems generate pulses of electromagnetic energy toward a direction of interest. Reflected signals received by the system are processed and displayed to operators. As with active sonar, a radar reflection off a target is limited by the target's ability to absorb the transmitted signal.[32]

Transmitters that use pulses of energy followed by inactive periods for reception of reflection are known as pulse radars. These systems use breaks in pulse transmission to prevent interference of transmitted and reflected signals, which would prevent

---

[30] Payne, 2010.

[31] D. Marx, M. Nelson, E. Chang, W. Gillespie, A. Putney, and K. Warman, "An Introduction to Synthetic Aperture Sonar," *Proceedings of the Tenth IEEE Workshop on Statistical Signal and Array Processing*, 2000.

[32] Payne, 2010.

the proper identification of targets. Continuous wave radars transmit with no breaks, so reflections off of stationary targets cannot be distinguished from transmitted signals. However, relative motion between the radar and target causes signal distortions following the reflection, which enables these signals to be distinguished from the transmitted signals.[33]

Radar is particularly useful because electromagnetic waves transmitting through air are not limited by time of day or visibility conditions, although certain types of radar perform better in adverse weather conditions than others do. However, radar systems can be utilized only for UAVs, USVs, or surfaced UUVs because electromagnetic waves may be transmitted only through the air. This is because the wavelength of electromagnetic energy used by radar systems is easily absorbed by seawater, which prevents reflection off a target and reception of returned signals. Notably, radar range increases with height above the ocean level. Limited elevation above the water level results in transmitted pulses traveling only a short distance and then being absorbed by wavetops. On the opposite end of the scale from radars located just above the ocean surface are airborne radars, which have a much larger range because they can transmit beyond the horizon seen at sea level.[34]

Other notable radar parameters for unmanned vehicles include frequency, beam width, antenna gain, and power output. Radar frequency, beam width, and antenna gain are functions of transmit antenna size and shape. Such size and shape are particularly important parameters for unmanned vehicles, which are often constrained by platform size. While higher frequencies allow for smaller antennas, smaller antennas generally illuminate smaller areas and therefore have narrower beams.[35] As with sonar systems, radar systems with a broader search area tend to be less effective at developing high-resolution images of a more detailed target area.

Higher-frequency transmissions are more subject to losses resulting from ambient conditions, such as foul weather. Antenna gain is a characterization of an antenna's ability to concentrate transmissions or receive signals in a particular direction. This ability is also highly dependent on antenna shape. Last, as with sonar systems, radar power is a major factor in determining radar system range. Increased transmission power increases the power of reflected signals, increasing the likelihood that reflected signals are received by the originator. However, increased power comes at a platform-level cost and may not be feasible in all applications.[36]

---

[33] Payne, 2010.

[34] Payne, 2010.

[35] Workarounds for antennas with narrower beams exist, such as rotating the antenna-transmitting element in a circle to illuminate an area 360 degrees around the transmitting element. Phased-array radars are an alternative to mechanically rotating the radar element. Phased-array radars use designed interference from multiple collocated transmitting elements to steer a radar beam in a direction of interest. See Payne, 2010.

[36] Payne, 2010.

Taking into account these standard parameters, there are three notable types of radar that are worth mentioning in the context of unmanned maritime vehicles: S band, X band, and dual band. The terms *S band* and *X band* are IEEE designations that refer to the frequency band of the associated radar system; X-band radars use a higher frequency range.[37] As described earlier, using a higher frequency range allows for a physically smaller antenna and, potentially, a higher-resolution return signal but risks higher-propagation losses, especially in poor weather conditions. Because of their smaller size, X-band radars are of particular interest for smaller platforms. A *dual-band* radar is a very specialized radar system that can simultaneously use the S and X bands. Such a system allows a vessel to take advantage of the higher resolution of the X band with the improved poor-weather performance of the S band. Currently, the CVN-78 is the U.S. Navy's only vessel operating a dual-band radar.[38] Radars of many frequencies are used throughout the Navy, including on unmanned vehicles, for various purposes, such as ISR, decoy missions, fire control, targeting, and navigation.

Environmental

In the academic and commercial worlds, unmanned maritime vehicles are broadly used for gathering various types of environmental data. For instance, sea surface temperature, subsurface temperature, surface and subsurface salinity, sea level, ocean color, ocean current, and ice levels are examples of key ocean environmental parameters measured and tracked by the European Commission's Copernicus Marine Environment Monitoring Service. Understanding the temporal behavior of these and other oceanic environmental parameters is important for meteorology and various maritime industries, including fishing and maritime shipping.[39]

Such environmental parameters have significant effects on operations in a region, can rapidly change with weather conditions, and are very location-dependent. For example, during turbulent sea conditions, sea bottom sediment may mix with water in shallow coastal areas, limiting visibility for optical sensors aboard vehicles collecting data in a given geographic area. As we discuss later, a detailed understanding of a region's environmental conditions is an important aspect of military operations in a region and may be enabled by unmanned maritime vehicles carrying one of a variety of sensors to record environmental conditions. Environmental sensors, such as those used to measure the parameters mentioned here, are relatively small sensors that rely on little

---

[37] The S band uses a frequency range of 2 GHz to 4 GHz, and the X band uses a range of 8 GHz to 12 GHz. See National Academies of Sciences, Engineering, and Medicine, *Handbook of Frequency Allocations and Spectrum Protection for Scientific Uses*, 2nd ed., Washington, D.C.: National Academies Press, 2015.

[38] Raytheon, "Dual Band Radar," webpage, undated-b.

[39] Karina von Schuckmann, Pierre-Yves Le Traon, Enrique Alvarez-Fanjul, Lars Axell, et al., "The Copernicus Marine Environment Monitoring Service Ocean State Report," *Journal of Operational Oceanography*, Vol. 9, Suppl. 2, September 2016.

external power, which means that their presence aboard maritime autonomous vehicles may have relatively small effects on power needs.

Light and Optical

Optical sensing under water is complicated by the fact that optical wavelengths are scattered, diffused, and distorted by completely clear water—and even more so by the cloudy, turbid water of the ocean. Imaging with traditional electro-optical cameras under water is further complicated by the low-light underwater environment and multiple layers of light refraction that take place among the sensor, the glass covering the sensor, and the water between the sensor face and the targeted object. Some of these challenges may be corrected by sensor calibration, but this requires knowledge of water conditions surrounding the sensor, as well as sufficient processing capability to correct for these errors on the fly. Other methods to improve sensing associated with electro-optical cameras include ensuring close proximity of the sensor to the targeted object and recording multiple images of the target in varying light conditions or from varying locations.[40] Optical sensors are used for mine identification in MCM operations today. Collecting clear optical images is a time-intensive operation that requires significant platform power usage.

Two commonly used mapping and imaging methods for various platforms are (1) light detection and ranging and (2) laser line scanning. Much like sonar and radar, these methods rely on measured return signals from an initial pulse transmission to collect data about a target's location. These sensors come in two primary types: time of flight and triangulation, each distinguished by how return pulses are received and translated into usable information. Time-of-flight systems determine the length of time required for a pulse to travel from transmitter to receiver in order to develop an image of the target. Triangulation systems rely on multiple laser devices dispersed to a common feature to develop a sense of distance and depth from the target. While laser technology is frequently used for coastal seabed imaging, it is not often used for wholly underwater measurements. This is because light is often reflected and refracted by water, so reception of the transmitted pulses is not reliable.[41]

---

[40] Exposing target objects to light from multiple angles can assist in recreating a 3D model of the target object. This may be performed by multiple optical sensors capturing an image or by moving one sensor around to capture multiple images of the target. See Miquel Massot-Campos and Gabriel Oliver-Codina, "Optical Sensors and Methods for Underwater 3D Reconstruction," *Sensors*, Vol. 15, No. 12, 2015.

[41] Massot-Campos and Oliver-Codina, 2015.

Communications

Radio and Satellite

The principles of radio and satellite communication are the same as those described for radar systems, except for the power required for communication transmissions and the frequency band used for these communication modes. Providing sufficient communication transmission power is important to ensuring that each transmission reaches its desired destination. That said, because communication signals are not expected to travel twice the distance to the target,[42] the required power level for these signals is lower than that for radar. Communication frequency bands may be divided into a few categories, as described in Table B.1.

A key distinction in the communication bands in Table B.1 is the propagation mode of each band (third column). *Propagation mode* refers to the way that the electromagnetic wave interacts with the earth's surface and atmosphere as it travels from the source to its destination. Ground waves follow the earth's surface and do not typically penetrate the atmosphere. This behavior allows them to travel long distances, although they are more subject to various kinds of losses and have less bandwidth than higher-frequency signals do. Sky waves are of an appropriate frequency to use a reflection interaction with the atmosphere to travel long distances. Although these interactions behave somewhat predictably with ambient conditions (weather, time of day, and sun activity), such conditions are outside the control of operators and may not always be condu-

Table B.1
Radio and Satellite Communication Bands

Frequency Range	Frequency Band Designation	Propagation Mode
30–3,000 Hz	Extremely low	Ground wave
3–30 kHz	Very low	Ground wave
30–300 kHz	Low	Ground wave
300–3,000 kHz	Medium	Sky wave
3–30 MHz	High	Sky wave
30–300 MHz	Very high	Space wave
300–3,000 MHz	Ultra high	Space wave
3–30 GHz	Super high	Space wave
30–300 GHZ	Extremely high	Space wave

SOURCE: Payne, 2010.

---

[42] That is, they are required to travel only once toward the target, rather than once toward the target and once back to the origin after reflecting off the target (like in radar transmissions).

cive to desired operations. With insufficient reflection off the atmosphere, sky waves become *space waves*—that is, waves that can penetrate the atmosphere and be received by orbiting satellites. Satellite communication takes advantage of this phenomenon to allow terrestrial entities to communicate with satellites, which, in turn, can relay the communication around the world. Space waves may also be used for exclusively terrestrial communications, although they are unable to reflect off the atmosphere to increase their range, instead passing through the atmosphere into space.[43]

In terms of design considerations, like with radar systems, lower-frequency bands require physically larger antennas, which might not be feasible for space-limited unmanned vehicles. As a result, unmanned maritime vehicles may prefer higher-frequency communication bands, which are more limited in their maximum ranges. Additionally, radio and satellite communication frequency bands do not propagate under water and are therefore available only to UAVs, USVs, and surfaced UUVs. Thus, while radio and satellite communications may be used somewhat covertly, additional mechanisms, such as encryption, need to be used to ensure that these signals are not intercepted by adversaries. In addition to intercepting radio or satellite transmissions, adversaries may jam signals in given frequency bands. Jamming involves transmitting powerful, spurious signals in a given frequency band to prevent receivers from identifying signals that they are expecting.[44]

Acoustic

As a result of the challenges in transmitting electromagnetic signals under water, acoustic communications via acoustic modem are the primary method of communication with and between deployed undersea vehicles. Acoustic communications behave similarly to active sonar systems. Low-frequency acoustic signals propagate better under water, but, as described earlier, low-frequency signals have low bandwidth and therefore low data rates. Acoustic signals transmitted under water also are subject to various loss mechanisms, must contend with significant ambient noise, travel at relatively low speeds, and have the potential to alert nearby adversaries to the presence of the vehicle transmitting the communication signal.[45] Acoustic communications are used for nearly all of the Navy’s unmanned undersea kill chains today.

Tethered

Tethered communications involve a hardwired connection between an unmanned vehicle and a control station, limiting the range that the unmanned vehicle may travel away from the station. On the positive side, hardwired communications allow a communi-

---

[43] Payne, 2010.

[44] Payne, 2010.

[45] M. Stojanovic and J. Preisig, “Underwater Acoustic Communication Channels: Propagation Models and Statistical Characterization,” *IEEE Communications Magazine*, Vol. 47, No. 1, 2009.

cation link that is less lossy, has a higher data rate, and is faster than other communication mechanisms described in this section. As a result, hardwired communications are especially valuable for remotely operated, tethered salvage UUVs that use manipulator arms (discussed later) to perform their missions. These UUVs require near-real-time transmission of a high-quality video feed to the vehicle operator to allow proper operation of the vehicle's manipulator arms. Fiber optic and ethernet are the two most commonly reference types of tethered communications that we identified during our survey of common payloads and existing unmanned maritime platforms.[46]

Weapons

Minesweeping Gear

Clearing minefields using surface vessels is typically achieved by employing minesweeping gear. Contact sweeping gear is towed behind the vessel and cuts mine tethers, and the mine can then be shot and detonated from a safe distance once it reaches the surface. Acoustic sweeping gear can detonate acoustic mines by simulating sound waves emitted from ship propellers and projecting them over longer distances. Similarly, magnetic sweeping gear passes electrical current through a cable, creating a magnetic field that is strong enough to influence the firing mechanism of nearby mines.[47] A degaussing system can also be used aboard a ship to reverse the ship's magnetic field. Electrical cables are wrapped around the ship's hull to achieve this effect.[48]

Electronic Warfare Systems

Electronic warfare systems can be hosted on ships to provide electronic support, detection, and countermeasure protection against adversaries.[49] The Navy has used the AN/SLQ-32 system since the 1970s but has invested in the Surface Electronic Warfare Improvement Program to enhance capabilities of the system.[50] Specific capabilities that have been augmented or introduced include antiship missile defense, counter-

---

[46] Another variant of tethered communications is the use of undersea telecommunication infrastructure to communicate between an operational undersea platform and a home station. While this is a more complex arrangement, it is a higher-data-rate method that eliminates the limitation of range associated with most cabled communication methods. In our review of maritime communication methods, we found that this method had been used for manned undersea vehicles in the past, but we did not find it to be associated with any unmanned platforms.

[47] Algerines, "The 'Art' of Minesweeping," webpage, May 2013.

[48] Cryogenic Society of America, "HTS Degaussing System," *Cold Facts*, Vol. 25, No. 2, Spring 2009.

[49] Raytheon, "AN/SLQ-32(V) Shipboard EW System," webpage, undated-a; and Lockheed Martin, "Naval Electronic Warfare," webpage, undated.

[50] The Surface Electronic Warfare Improvement Program is a block upgrade program; three blocks are currently set, and a fourth is being discussed. See U.S. Navy, "Surface Electronic Warfare Improvement Program (SEWIP)," webpage, January 30, 2017.

targeting and countersurveillance, upgraded electronic support antennas and receivers, and an open combat system interface. Electronic attack improvements are tabled for improvement in the future.[51]

Weapon Suites

We identified several weapon suites implemented on USVs to provide force protection and antipiracy capabilities. These systems can be equipped with assault rifles; rocket, missile, and grenade launchers; or machine and Gatling guns.[52] Autonomy in weapon suites is severely limited; all systems that we found required remote operation (or having a human in the loop to at least make decisions on engagement).

Miscellaneous

Manipulator Arms

Historically used on UGVs, manipulator arms are now being developed for underwater use.[53] The applications currently envisioned are (1) explosive ordnance disposal and (2) inspection, repair, and maintenance (particularly useful to the oil and gas industry). For explosive ordnance disposal, the goal is to provide standoff capabilities to divers, in much the same way that UGVs equipped with manipulator arms assist bomb technicians in disabling ordnance on land.

Towing Mechanisms

Towing mechanisms on USVs have been used to deploy sensor arrays in the water, on the surface, and in the air. Sensors towed in the water are typically some type of sonar array that provides the USV with enhanced detection capabilities. The Towed Airborne Lift of Naval Systems project is an example of an aerially towed array in which sensors used for ISR can be deployed at altitudes of 500 to 1,500 ft above sea level, greatly increasing line of sight.[54]

---

[51] U.S. Navy, 2017.

[52] See Berenice Baker, “No Hands on Deck—Arming Unmanned Surface Vessels,” Naval Technology, November 22, 2012.

[53] RE2 Robotics, “RE2 Robotics Wins Navy Contract to Develop Underwater Manipulator Arms,” press release, February 7, 2017.

[54] DARPA, “TALONS Tested on Commissioned U.S. Navy Vessel for First Time,” press release, August 15, 2017.

APPENDIX C

In-Depth Analysis: Platforms

To identify existing UUVs and USVs, we investigated platforms being developed or used by academia, industry, or DoD (particularly DARPA and the Navy). We developed a catalog to track our findings and list key platform parameters, such as length, speed, and endurance. The parameters we recorded as part of our catalog are listed in the box below.

We performed this investigation through a review of open-source literature regarding unmanned maritime platforms. This literature included published journal articles discussing use of a particular platform by a research organization, online brochures and data sheets for platforms being sold by a variety of industry organizations, and unclassified documentation of DoD unmanned platform capabilities. Notably, this public documentation did not always list all of our parameters of interest. Despite this limitation, we reviewed the catalog data for trends and patterns to determine gaps in which additional Navy investment could further the unmanned maritime platform design space. In our review, we identified various platforms that are either tethered or limited in their autonomy. We still added these platforms to the catalog because we did not want to omit systems before fully investigating them.

Parameters of Interest for Unmanned Maritime Platforms

Class (USV, UUV, etc.)
Model name
Beam width (m)
Payload power (kW)
Endurance (days)
Launch and recovery method
Onboard sensors
Algorithms
Country of origin
Source (industry, DoD, academia)
Purpose or mission
Height (m)
Nominal speed (knots)
Range (miles)
Level of autonomy
Power source
Communication methods
Manufacturer
Length (m)
Dry weight (kg)
Max speed (knots)
Max operating depth (m)
Tethered or nontethered
Actuators
Technology readiness level

Our search resulted in the identification of 178 UUV platforms and 89 USV platforms, for a total of 267 platforms. Table C.1 displays the number of platforms that each investment source is predominantly developing or using.

To identify the trends related to platform design parameters, we attempted to independently assess the full data set of platforms based on each parameter of interest. As mentioned in Chapter Two, because of the inconsistent nature of the open-source data that we evaluated, we were unable to obtain data for each parameter for all platforms in our catalog. Because of data availability, the parameters that we were able to assess for trends were platform length, weight, speed, endurance, and power source.

Table C.1
Investment Sources for Unmanned Maritime Platforms

Investment Source	USV Platforms	UUV Platforms
Academia	2	31
Industry	84	135
DoD[a]	3	12

[a] The DoD category consists of 14 platforms developed by the Navy and one by DARPA.

APPENDIX D

In-Depth Analysis: Patents

In this appendix, we describe patents as they relate to autonomous maritime technology, with two purposes. First, we aim to characterize the patented technological domain of autonomous maritime systems. This domain can be understood as a subset of the overall technological domain in that not all technological innovations are patented. Some inventors choose to protect their intellectual property via secrecy. Such inventions are not captured by patent analysis. Although we recognize this caveat, the use of patent data offers the means to, among other things, quantify long-term technological trends, identify major national and organizational actors, identify with precision the technological subfields in which innovative activity is concentrated, and determine the particular technologies and organizations that have driven technological change in the field.

Second, we aim to describe a subset of patented autonomous maritime technologies that have explicit military application. While the dual-use character of much of the technology associated with autonomous systems makes a discrete demarcation between civilian and military ends difficult, certain technologies can plainly be categorized as military. By examining these technologies in detail, we aim to identify technological trends of particular relevance to naval warfare.

In service of these objectives, we constructed two data sets. The first comprises the universe of patents for autonomous systems that are meant for use in maritime settings. It consists of 3,813 patents filed from 1974 to 2018 by 2,895 organizations in 44 countries.[1] The second data set comprises autonomous maritime system patents that contain, in their documentation, an explicit reference to a military application. It consists of 247 patents filed from 1989 to 2018 by 290 organizations in 22 countries.

---

[1] Technically, the results refer to patent families (a patent family is the set of patents granted in various countries for a single underlying innovation). For analysis such as this one, using patent families as units of innovation is preferred over using individual patents because using patent families avoids double counting a single innovation that has been filed in more than one jurisdiction.

Data Collection

To arrive at the first sample of patents, we began by searching the Derwent Innovations Index for a series of keywords related to autonomous maritime technologies.[2] We downloaded the full results as individual text files, each of which contained the complete documentation for the patent in question. Conveniently, the text files were fielded and could thus can be parsed based on field indicators (e.g., TI = patent title, PN = priority number). We used Vantage Point, a text-mining software application, to parse and clean the text files and remove duplicate entries.

For the second sample, we sought to identify a subset of patents from the first sample that have overt military application. That is, in defining the scope of the second sample, we sought to include only patents for maritime autonomous technologies in which the associated patent documentation referred to at least one intended military application. We found that filtering the first sample using Derwent Class Code W07 (electrical military equipment and weapons) achieved this end.[3]

Autonomous Maritime Patents

In this section, we examine the patented technological domain for autonomous maritime systems by presenting three sets of analyses. First, we present data, at various levels of aggregation, on patent output trends over time. Second, we define the major players in the technological space by presenting cumulative patent output data at the organization and country levels. Finally, we present data on the particular technological subfields in which innovation has been concentrated.

Time Trends

Examining autonomous maritime patent output with respect to time reveals dramatic growth. The first patent fitting our search criteria was filed in 1970 in the United Kingdom for a UUV used to conduct underwater surveying (patent number GB1500684). In 2016, the most recent year for which a full cohort of data is available, 680 patents matching the search criteria were granted. Figure D.1 plots the annual global output

---

[2] The exact search was

> TS = (unmanned maritime OR unmanned underwater OR autonomou* underwater OR autonomou* maritime OR USV OR unmanned boat OR autonomou* boat OR autonomou* ship water OR autonomou* ship sea OR autonomou* submarine OR unmanned submarine OR unmanned undersea OR autonomou* undersea).

The search was conducted on May 1, 2018.

[3] Derwent Class Codes are curated by subject-matter experts at Thomson Reuters. For further justification regarding the use of code W07 to identify military technologies, see Jon Schmid, "The Diffusion of Military Technology," *Defence and Peace Economics*, February 17, 2017.

Figure D.1
Global Output of Autonomous Maritime Patents, 1970–2016

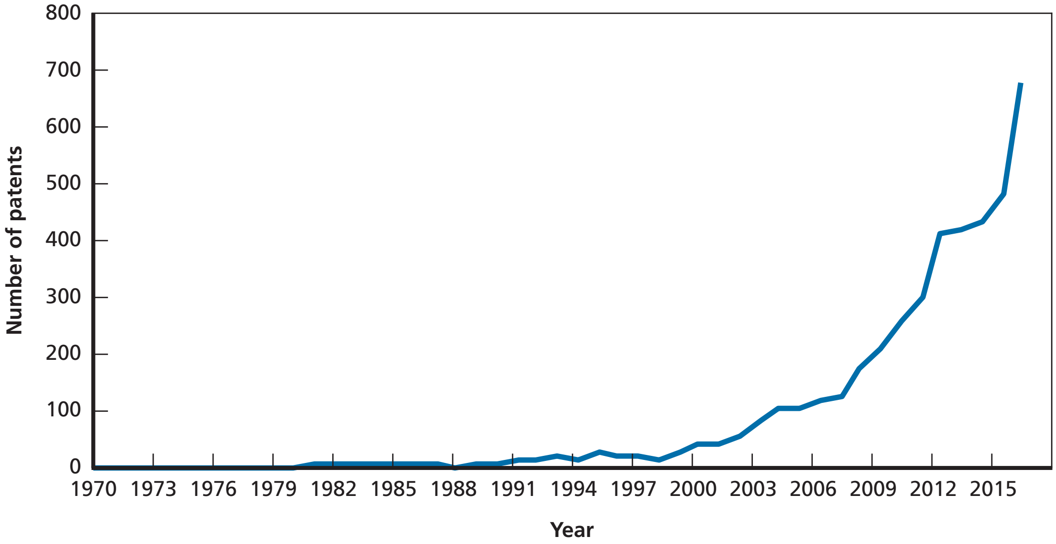

SOURCE: Clarivate Analytics, Derwent Innovations Index, accessed May 1, 2018a.

of autonomous maritime patents from 1970 to 2016. The average annual growth rate over this period was 19 percent.

These data can also be used to generate forecasts. Figure D.2 uses the observed data to specify an exponential function to model global patent output from 1990 to 2020. The function can be used to generate annual forecasts for future years. The predicted annual global output values shown in Figure D.2 were 774 for 2017, 916 for 2018, 1,085 for 2019, and 1,284 for 2020.[4]

Decomposing the global growth trend by country reveals that the accelerated growth observed in recent years was driven largely by Chinese output. In the most recent decade for which a full complement of data is available (2007–2016), China's annual patent output grew at a rate of 85 percent, while the output of the United States and the rest of the world grew at 18 percent and 17 percent, respectively. Figure D.3 plots the annual patent output for China, the United States, and the rest of the world from 2000 to 2016.[5]

---

[4] The estimated function is $6.836^{e0.169t}$, where t equals the number of years beyond $t = 0$.

[5] China's accelerated output appears to have continued in 2017 and the first half of 2018. While the data for this period are incomplete because of delays in the aggregation of national-level data, the existing data indicate that, during that period, China produced more autonomous maritime patents (608) than the United States and the rest of the world combined (424).

Figure D.2
Global Output of Autonomous Maritime Patents, Observed Versus Predicted, 1990–2020

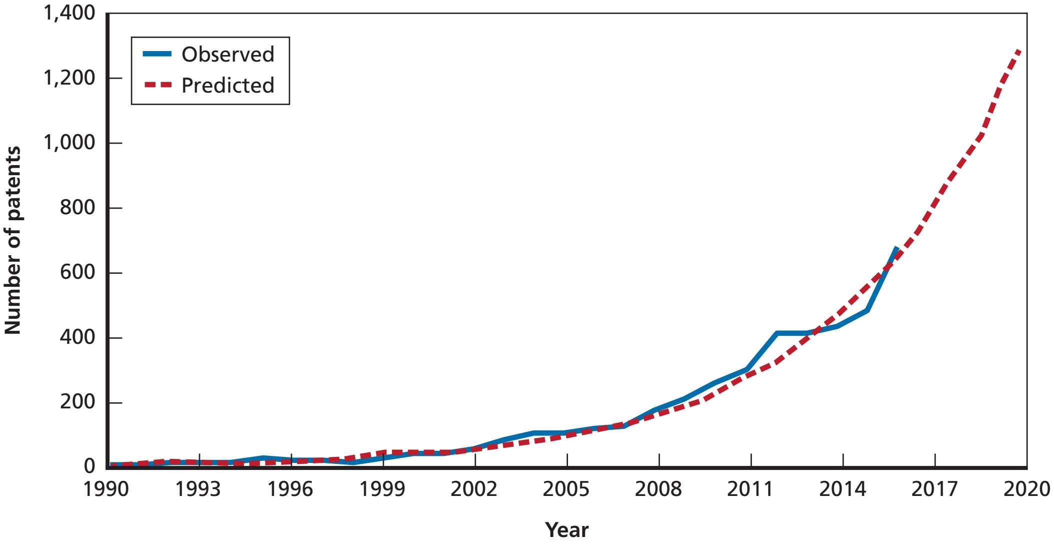

SOURCE: Clarivate Analytics, 2018a.

Figure D.3
Output of Autonomous Maritime Patents for China, the United States, and the Rest of the World, 2000–2016

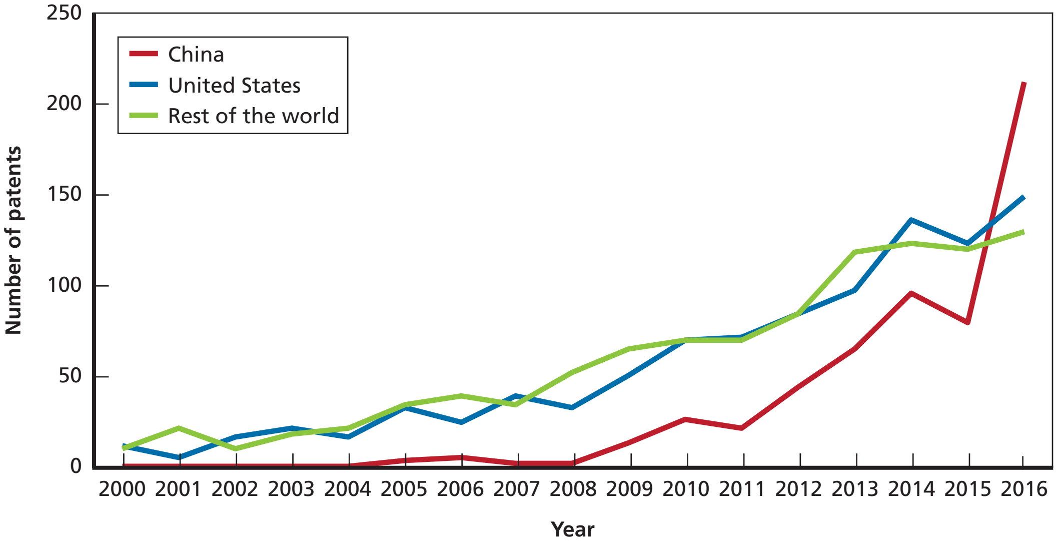

SOURCE: Clarivate Analytics, 2018a.

Major Players

Table D.1 ranks the top 15 countries in terms of cumulative autonomous maritime patent output. The table indicates that the United States and China have been responsible for the bulk of technological progress in this field, representing more than 70 percent of all patented innovation.

Table D.2 provides a list of the top 30 organizations to which autonomous maritime patents are assigned. These organizations, known as assignees, hold the rights to the intellectual property that underlies the patents. Although exceptions exist, an assignee is typically the organization that was responsible for developing the patented innovation.

Table D.1
Top 15 Countries for the Output of Autonomous Maritime Patents, 1970–2016

Country	Number of Autonomous Maritime Patents	Percentage of Total
United States	1,279	38.80
China	1,184	35.92
South Korea	174	5.28
Germany	145	4.40
Japan	144	4.37
Russia	131	3.97
France	103	3.13
United Kingdom	66	2.00
Canada	14	0.42
Spain	12	0.36
India	12	0.36
Netherlands	6	0.18
Australia	5	0.15
Taiwan	5	0.15
Brazil	4	0.12

NOTE: The numbers in this table are based on basic patent country of origin. Patents attributed to the Soviet Union and East Germany are added to Russia and Germany, respectively. Excluded are patents that were filed with the World Intellectual Property Organization or the European Patent Organization. Excluding these data likely results in an underestimate of the output of European countries, which are more likely to file patents at these multilateral agencies. In the sample considered here, 401 patents listed the World Intellectual Property Organization as their basic country, and 115 listed the European Patent Organization.

Table D.2
Top 30 Assignees of Autonomous Maritime Patents, 1970–2016

Organization	Number of Autonomous Maritime Patents	Country of Origin	Organization Type
U.S. Navy	183	United States	Government
Google	163	United States	Firm
Harbin Engineering University	136	China	University
Boeing	91	United States	Firm
Lockheed Martin	69	United States	Firm
Atlas Elektronik	66	Germany	Firm
DJI	69	China	Firm
Shenyang Institute of Automation	51	China	Government
Powervision Tech	50	China	Firm
Honeywell	38	United States	Firm
Mitsubishi	38	Japan	Firm
Raytheon	37	United States	Firm
Zhejiang University	33	China	University
Jiangsu University of Science and Technology	30	China	University
Shanghai University	30	China	University
Wuhan University of Technology	28	China	University
Northwestern Polytechnic University	26	China	University
Korean Institute of Ocean Science and Technology	25	Korea	Government
CGG	23	France	Firm
BAE Systems	22	United Kingdom	Firm
Daewoo Shipbuilding	21	South Korea	Firm
Wuhan Intelligent Equipment	21	China	Firm
Veniam	17	United States	Firm
China Shipbuilding Co.	14	China	Firm (state-owned)
Sifang Jibao Co.	14	China	Firm
Huazhong University of Science and Technology	14	China	University
Shanghai Maritime University	14	China	University
South China University of Technology	14	China	University
Nanjing University of Information Science and Technology	13	China	University
Shanghai Jiao Tong University	13	China	University

Comparing the top assignees for the United States and for China reveals an interesting distinction in how innovation proceeds in each country. In the United States, with the exception of the large amount of innovation conducted by the U.S. Navy, the vast majority of innovation in this technological field is advanced by firms.[6] In China, on the other hand, universities play a more significant part; indeed, 11 of the country's top 16 patent assignees are universities, including the assignee with the largest single number of patents. While private firms are strongly represented in the top five of China's patent assignees, universities have a larger role there than they do in the United States.

Technological Focus

Table D.3 provides the top ten four-digit International Patent Classification (IPC) codes for the patents in our first sample (all autonomous maritime patents). IPC codes classify patents based on their technological content and thus offer a means to identify the technological fields in which patenting is concentrated. Examining the most-com-

Table D.3
Top Ten IPC Codes for Autonomous Maritime Patents, 1970–2016

IPC	Description	Number of Patents
B63G	Offensive or defensive arrangements on vessels; minelaying; minesweeping; submarines; aircraft carriers	637
G05D	Systems for controlling or regulating nonelectric variables	626
B63B	Ships or other waterborne vessels; equipment for shipping	588
B63C	Launching, hauling out, or dry-docking of vessels; lifesaving in water; equipment for dwelling or working under water; means for salvaging or searching for underwater objects	450
G01S	Radio direction-finding; radio navigation; determining distance or velocity by use of radio waves; locating or presence-detecting by use of the reflection or reradiation of radio waves; analogous arrangements using other waves	407
G01C	Measuring distances, levels, or bearings; surveying; navigation; gyroscopic instruments; photogrammetry or videogrammetry	364
B63H	Marine propulsion or steering	290
B64C	Airplanes; helicopters	251
G06F	Electric digital data-processing	213
H04B	Transmission	172

---

[6] The technological contribution of Google to this field was also particularly evident. Google developed 11 of the top 50 most highly cited patents. The importance of patent citation is discussed later in this appendix.

mon IPC codes assigned to the patents in the overall autonomous maritime sample provides additional support for two observations made in the body of this report. First, with the exception of code B63G, none of the top ten IPC classes is related to weapon or military technologies. This observation supports the contention made earlier that much of the progress in UUV and USV development has been advanced by the commercial sector. Second, the allocation of patents across IPC classes confirms the high rate of innovation in both platforms (B63G, B63B, B64C) and information-based processes (G01S, G01C, G06F, H04B) that was described in Chapter Two of this report.

IPC codes are assigned to patents based on the code's relevance to the intellectual content that underlies the patent. A patent can be assigned multiple IPC codes if the underlying innovation is relevant to more than one classification. Thus, a patent that has been assigned two IPC codes covers an innovation that relates to subject matter from both classifications. This feature of the patent classification process allows us to increase the granularity with which we consider the technological fields in which innovation is concentrated.

Figure D.4 provides the autocorrelation matrix of four-digit IPC codes for the autonomous maritime patent data set. Table D.4 provides the descriptions of these codes. Code pairs with high correlations (indicated by darker shades of red) represent technological subfields that frequently appear together on a single patent. The matrix reveals several code pairs to be particularly common in the autonomous maritime patent data set.

The highest pairwise correlation is between B64C and B64D. Examining individual patents—there are 77 in the sample—reveals that these patents tend to fall into two groups. First are technologies that simultaneously span the sea and air domains. Examples of such technologies include an amphibious UAV (patent number CN106585948) and a UAV-based active sonar buoy for potential use in ASW (patent number WO2011137335). The second type of patent in this group is for technologies of prospective application to either maritime or aerial vehicles. These technologies are exemplified by various power control systems (e.g., patent numbers WO2017113338 and 2016061726) meant for use in either maritime or aerial autonomous systems.

The second-highest pairwise correlation is between G06T and G06K. Examining the patents that were assigned codes from both of these IPC classes reveals a high rate of innovation related to the capturing, detecting, tracking, and transmitting of image information to and from a maritime setting. Examples of such patents include a method for detecting objects from sonar image by unmanned underwater craft (patent number DE102013102650), a means of tracking targets from unmanned boats (patent number CN106981071), and methods to optimize images captured during underwater surveys (patent number WO2014060562).

Figure D.4
Autocorrelation Matrix of IPC Codes for Autonomous Maritime Patents

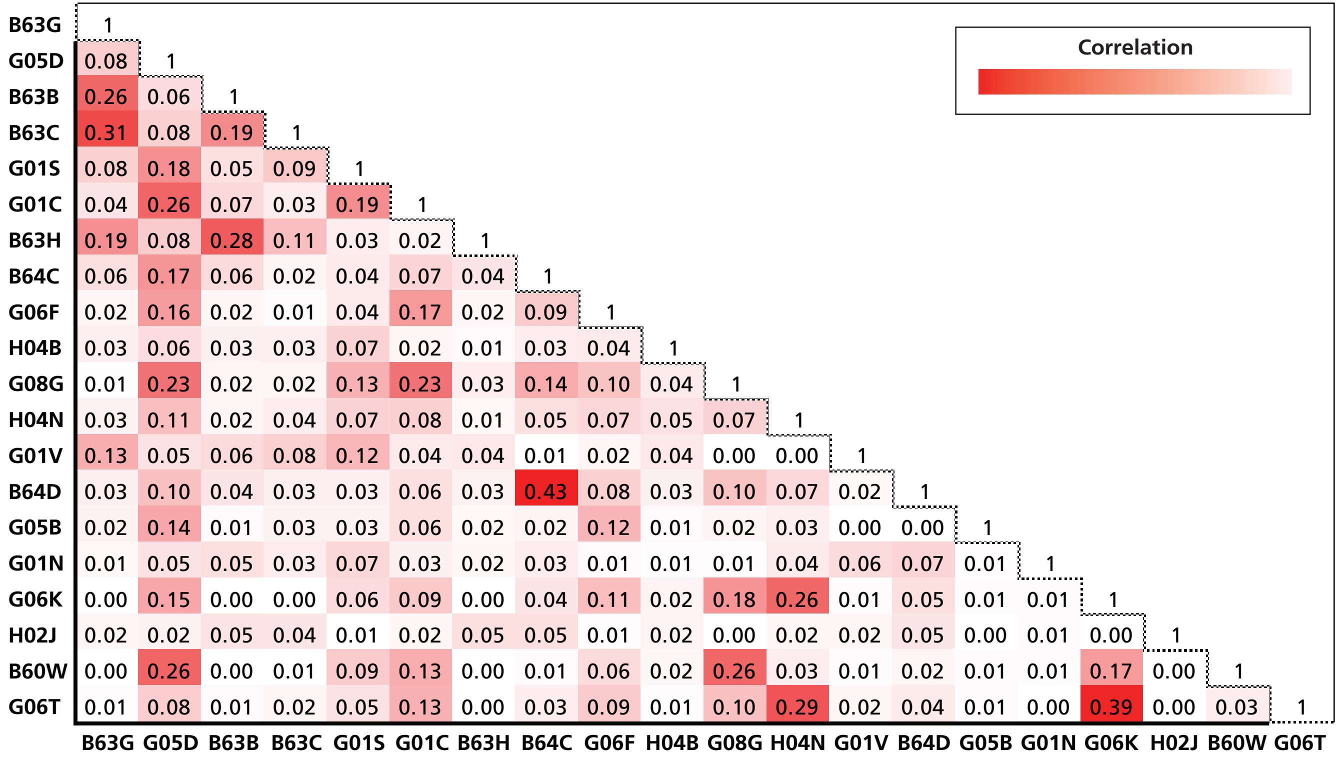

	B63G	G05D	B63B	B63C	G01S	G01C	B63H	B64C	G06F	H04B	G08G	H04N	G01V	B64D	G05B	G01N	G06K	H02J	B60W	G06T
B63G	1																			
G05D	0.08	1																		
B63B	0.26	0.06	1																	
B63C	0.31	0.08	0.19	1																
G01S	0.08	0.18	0.05	0.09	1															
G01C	0.04	0.26	0.07	0.03	0.19	1														
B63H	0.19	0.08	0.28	0.11	0.03	0.02	1													
B64C	0.06	0.17	0.06	0.02	0.04	0.07	0.04	1												
G06F	0.02	0.16	0.02	0.01	0.04	0.17	0.02	0.09	1											
H04B	0.03	0.06	0.03	0.03	0.07	0.02	0.01	0.03	0.04	1										
G08G	0.01	0.23	0.02	0.02	0.13	0.23	0.03	0.14	0.10	0.04	1									
H04N	0.03	0.11	0.02	0.04	0.07	0.08	0.01	0.05	0.07	0.05	0.07	1								
G01V	0.13	0.05	0.06	0.08	0.12	0.04	0.04	0.01	0.02	0.04	0.00	0.00	1							
B64D	0.03	0.10	0.04	0.03	0.03	0.06	0.03	0.43	0.08	0.03	0.10	0.07	0.02	1						
G05B	0.02	0.14	0.01	0.03	0.03	0.06	0.02	0.02	0.12	0.01	0.02	0.03	0.00	0.00	1					
G01N	0.01	0.05	0.05	0.03	0.07	0.03	0.02	0.03	0.01	0.01	0.01	0.04	0.06	0.07	0.01	1				
G06K	0.00	0.15	0.00	0.00	0.06	0.09	0.00	0.04	0.11	0.02	0.18	0.26	0.01	0.05	0.01	0.01	1			
H02J	0.02	0.02	0.05	0.04	0.01	0.02	0.05	0.05	0.01	0.02	0.00	0.02	0.02	0.05	0.00	0.01	0.00	1		
B60W	0.00	0.26	0.00	0.01	0.09	0.13	0.00	0.01	0.06	0.02	0.26	0.03	0.01	0.02	0.01	0.01	0.17	0.00	1	
G06T	0.01	0.08	0.01	0.02	0.05	0.13	0.00	0.03	0.09	0.01	0.10	0.29	0.02	0.04	0.01	0.00	0.39	0.00	0.03	1

NOTE: Code pairs with high correlations (indicated by darker shades of red) represent technological subfields that frequently appear together on a single patent.

Table D.4
IPC Code Descriptions

IPC	Description
B63G	Offensive or defensive arrangements on vessels; minelaying; minesweeping; submarines; aircraft carriers
G05D	Systems for controlling or regulating nonelectric variables
B63B	Ships or other waterborne vessels; equipment for shipping
B63C	Launching, hauling out, or dry-docking of vessels; lifesaving in water; equipment for dwelling or working under water; means for salvaging or searching for underwater objects
G01S	Radio direction-finding; radio navigation; determining distance or velocity by use of radio waves; locating or presence-detecting by use of the reflection or reradiation of radio waves; analogous arrangements using other waves
G01C	Measuring distances, levels, or bearings; surveying; navigation; gyroscopic instruments; photogrammetry or videogrammetry
B63H	Marine propulsion or steering
B64C	Airplanes; helicopters
G06F	Electric digital data-processing
H04B	Transmission
G08G	Traffic control systems
H04N	Pictorial communication
G01V	Geophysics; gravitational measurements; detecting masses or objects
B64D	Equipment for fitting in or to aircraft; flying suits; parachutes; arrangements or mounting of power plants or propulsion transmissions in aircraft
G05B	Control or regulating systems in general; functional elements of such systems; monitoring or testing arrangements for such systems or elements
G01N	Investigating or analyzing materials by determining their chemical or physical properties
G06K	Recognition of data; presentation of data; record carriers; handling record carriers
H02J	Circuit arrangements or systems for supplying or distributing electric power; systems for storing electric energy
B60W	Conjoint control of vehicle subunits of different type or different function; control systems specially adapted for hybrid vehicles; road vehicle drive control systems for purposes not related to the control of a particular subunit
G06T	Image data processing or generation, in general

Patent Citation Analysis

Patent applicants are required to cite as "prior art" all patented innovations that were critical inputs to the applicant innovation. Highly cited patents thus refer to innovations that have played a substantial role in subsequent technological progress.[7] By considering the most highly cited patents within the autonomous maritime sample, we aim to identify the patented technologies and assignees that have made outsized contributions to the state of the art of autonomous maritime systems and their related technologies.

Table D.5 provides the five autonomous maritime patents that have been most frequently cited by other patents within the data set. That is, the table depicts the patents that have had the largest impact on the field of autonomous maritime systems. This information is useful in defining the innovations and organizations most responsible for advancing technological change in the field. The citation data used here are from the Derwent Patents Citation Index. This data source aggregates the citations received by patents filed at six major patent agencies: United States, Japan, the United Kingdom, Germany, the Patent Cooperation Treaty, and the European Patent Office. Citation counts were matched to patents using each patent's Derwent Primary Accession Number (a unique record identifier).

Table D.5
Top Five Most Highly Cited Patents for Autonomous Maritime Systems, 1970–2016

Patent Title (U.S. Patent Number)	Times Cited in the Autonomous Maritime Field	Assignee
Apparatus and method for deploying, recovering, servicing, and operating an autonomous underwater vehicle (US8109223)	41	Subsea 7 Ltd (UK firm)
Submarine deployed ocean bottom seismic system (US6474254B1)	24	Western Geco LLC (U.S. firm)
Sonotube compatible unmanned aerial vehicle and system (US6056237A)	17	Individual assignee
Underwater power and data relay (US6223675B1)	15	Coflexip SA (French firm)
Autonomous command and control unit for mobile platform (US6122572A)	13	State of Israel

SOURCE: Clarivate Analytics, Derwent Patents Citation Index, accessed May 1, 2018b.

---

[7] The notion that highly cited patents refer to particularly important innovations is supported by empirical findings that indicate that citation counts correlate strongly with the opinions of knowledgeable peers about the technological impact of a given patent (see M. B. Albert, D. Avery, F. Narin, and P. McAllister, "Direct Validation of Citation Counts as Indicators of Industrially Important Patents," *Research Policy*, Vol. 20, No. 3, June 1991) and the patent's market value (see Cristina Odasso, Giuseppe Scellato, and Elisa Ughetto, "Selling Patents at Auction: An Empirical Analysis of Patent Value," *Industrial and Corporate Change*, Vol. 24, No. 2, 2015).

Autonomous Maritime Military Patents

In this section, we discuss the same topics as in the previous section, except that the sample described here comprises patents for autonomous maritime technologies with an overt military application. We use this (considerably smaller) data set to describe temporal trends, identify major players, and identify the technological subfields in which innovative activity has been concentrated. It is worth noting that patents that are not classified as having an overt military application may nevertheless be used for military ends. The sample used here thus likely underestimates the amount of intellectual property in the field that has military application.

Time Trends

Figure D.5 plots the annual global output of autonomous maritime patents with military applications from 1989 to 2016. The first patent fitting these criteria was filed in 1989. This patent (patent number DE3826653) was for an underwater minesweeping apparatus developed by Rheinmetall GmbH, a German military technology development firm. In 2016, 15 patents fitting the search criteria were granted. Although output growth for the subset of patents with overt military application has been more variable than that of the broader field, the overall trend over time has been positive. The average annual growth rate between 1989 and 2016 was 36 percent. However, the small numbers of cases each year and the volatility of the time series limit the generality of interpretation. Numbers have certainly increased over time, but annual variation remains significant.

Figure D.5
Global Output of Autonomous Maritime Patents with a Military Application, 1989–2016

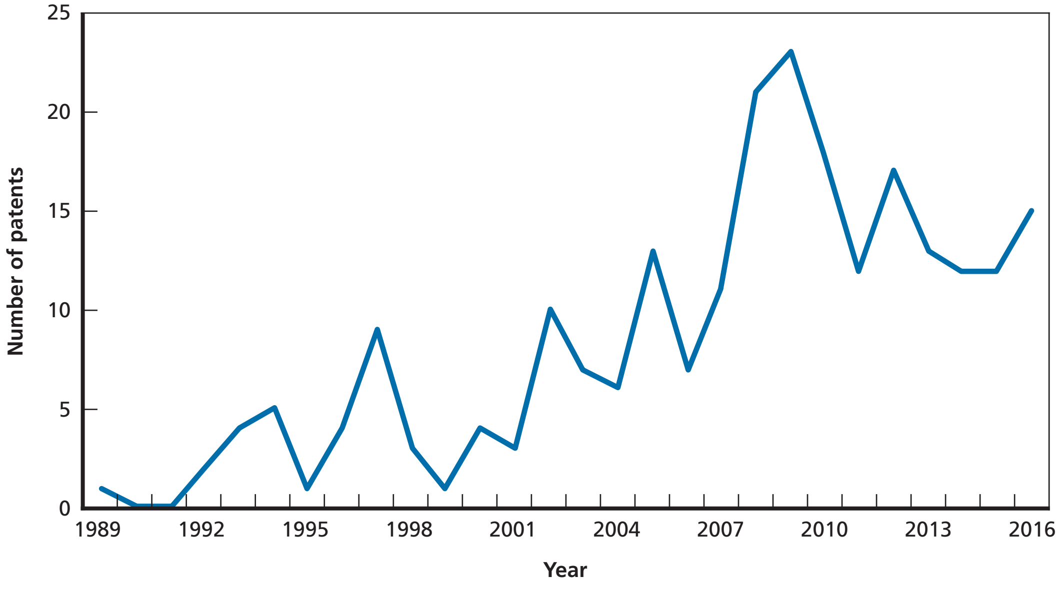

Major Players

Table D.6 lists the ten countries with the most cumulative autonomous maritime patents with a military application. The United States is responsible for the bulk of these innovations and has made a large and growing contribution to the broader field of autonomous maritime systems. In contrast, China's relatively low contribution to the global output of this subset of patents is somewhat conspicuous. However, we caution against assuming that China is uninterested in the military applications of unmanned maritime systems. Our classifying a patent as having a military application requires that the documentation for the patent in question explicitly mention such an application. If potential military applications are regularly omitted from patent documents for which such an application is intended or likely, the totals provided in the table would underestimate national levels of output of these technologies.

Table D.7 lists the primary organizations to which patents in the second sample are assigned. Two organizations, one for its contribution and one for its absence, are worth highlighting. First, the contribution of the U.S. Navy in developing these technologies is substantial. Of the 247 patents with military application that have been granted, the U.S. Navy is listed as assignee for 50 (more than 20 percent). Indeed, the top individual inventor for the sample was Navy engineer Christopher Hillenbrand, who is listed as inventor on eight patents (twice as many as any other individual).

Table D.6
Top Ten Countries for the Output of Autonomous Maritime Patents with a Military Application, 1989–2016

Country	Number of Autonomous Maritime Patents with a Military Application	Percentage of Total
United States	147	72.06
Germany	16	7.84
China	13	6.37
Japan	9	4.41
France	5	2.45
United Kingdom	4	1.96
South Korea	4	1.96
Russia	3	1.47
Canada	2	0.98
Australia	1	0.49

Table D.7
Top 12 Assignees of Autonomous Maritime Patents with a Military Application, 1989–2016

Organization	Number of Autonomous Maritime Patents with a Military Application	Country of Origin	Organization Type
U.S. Navy	50	United States	Government
Lockheed Martin	15	United States	Firm
Boeing	12	United States	Firm
Raytheon	12	United States	Firm
Atlas GmbH	7	Germany	Firm
Honeywell	7	United States	Firm
Draper Lab.	4	United States	Firm
Saab	4	Sweden	Firm
DJI	4	China	Firm
Geneva Aerospace	3	United States	Firm
Lawrence Livermore National Laboratory	3	United States	Government
Ocom Technologies Ltd.	3	China	Firm

Second, the absence of Google from Table D.7 is noteworthy. Google was the assignee associated with the largest number (163) of civilian autonomous maritime patents (see Table D.2). None of these patents, however, lists military applications among its potential uses.[8]

Technological Focus

Table D.8 provides the top ten IPC codes for the data set comprising technologies with an overt military application. Unsurprisingly, when compared with the primary technological fields represented in the civilian patent data, a greater number of top IPC codes for the second sample are explicitly related to weapon systems. Specifically, four (B63G, F42B, F41F, and F41H) of the top ten focus on weapon technologies. As was the case in the broader nonmilitary sample, we found a fairly balanced division of innovative output between platforms (B63G, B64C) and information-based processes (G06F, G05D, G01S, G01C).

---

[8] This observation is likely explained by two factors. First, Google's autonomous system research and development efforts thus far have been largely focused on civilian-facing applications, such as driverless automobiles. Second, Google is unlikely to enumerate within its patent documents the potential military applications of the innovations because of public relations considerations.

Table D.8
Top Ten IPC Codes for Autonomous Maritime Patents with a Military Application, 1989–2016

IPC	Description	Number of Patents
B63G	Offensive or defensive arrangements on vessels; minelaying; minesweeping; submarines; aircraft carriers	75
G06F	Electric digital data-processing	37
G05D	Systems for controlling or regulating nonelectric variables	33
G01S	Radio direction-finding; radio navigation; determining distance or velocity by use of radio waves; locating or presence-detecting by use of the reflection or reradiation of radio waves; analogous arrangements using other waves	30
G01C	Measuring distances, levels, or bearings; surveying; navigation; gyroscopic instruments; photogrammetry or videogrammetry	27
F42B	Explosive charges	26
B64C	Airplanes; helicopters	23
F41F	Apparatus for launching projectiles or missiles from barrels	19
H04B	Transmission	19
F41H	Armor; armored turrets; armored or armed vehicles; means of attack or defense	18

References

Akeila, E., Z. Salcic, and A. Swain, “Reducing Low-Cost INS Error Accumulation in Distance Estimation Using Self-Resetting,” *IEEE Transactions on Instrumentation and Measurement*, Vol. 63, No. 1, 2014, pp. 177–184.

Albert, M. B., D. Avery, F. Narin, and P. McAllister, “Direct Validation of Citation Counts as Indicators of Industrially Important Patents,” *Research Policy*, Vol. 20, No. 3, June 1991, pp. 251–259.

Algerines, “The ‘Art’ of Minesweeping,” webpage, May 2013. As of May 17, 2018:
http://www.minesweepers.org.uk/sweeping.htm

American Society for Cybernetics, homepage, 2016. As of August 31, 2018:
http://asc-cybernetics.org/

Amodei, Dario, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané, “Concrete Problems in AI Safety,” ArXiv, June 21, 2016. As of October 5, 2018:
https://arxiv.org/pdf/1606.06565v1.pdf

Amoiralis, Eleftherios, Marina Tsili, Vassilios Spathopoulos, and A. Hatziefremidis, “Energy Efficiency Optimization in UAVs: A Review,” *Materials Science Forum*, Vol. 792, August 2014, pp. 281–286.

Atanasov, N., J. Le Ny, K. Daniilidis, and G. J. Pappas, “Decentralized Active Information Acquisition: Theory and Application to Multi-Robot SLAM,” *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4775–4782.

Baker, Berenice, “No Hands on Deck—Arming Unmanned Surface Vessels,” Naval Technology, November 22, 2012. As of May 17, 2018:
https://www.naval-technology.com/features/featurehands-on-deck-armed-unmanned-surface-vessels/

Bolme, David S., J. Ross Beveridge, Bruce A. Draper, and Yui Man Lui, “Visual Object Tracking Using Adaptive Correlation Filters,” *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2010, pp. 2544–2550.

Carlini, Nicholas, and David Wagner, “Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods,” in Association for Computer Machinery, *AISec ’17: Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, New York, 2017, pp. 3–14.

Chen, Ling, Sen Wang, Klaus McDonald-Maier, and Huosheng Hu, “Towards Autonomous Localization and Mapping of AUVs: A Survey,” *International Journal of Intelligent Unmanned Systems*, Vol. 1, No. 2, 2013, pp. 97–120.

Clarivate Analytics, Derwent Innovations Index, accessed May 1, 2018a. As of May 1, 2018:
http://apps.webofknowledge.com/DIIDW_GeneralSearch_input.do?product=DIIDW&search_mode=GeneralSearch&SID=7AoZbHAMmzhNuI8BapS&preferencesSaved=

———, Derwent Patents Citation Index, accessed May 1, 2018b. As of October 15, 2018:
https://clarivate.com/products/derwent-patents-citation-index/

Clarke, E. M., T. A. Henzinger, and H. Veith, *Handbook of Model Checking*, Cham, Switzerland: Springer International Publishing, 2016.

Clynch, James R., "A Short Overview of Differential GPS," Naval Postgraduate School, Department of Oceanography, December 2001. As of April 19, 2018:
http://www.oc.nps.edu/oc2902w/gps/dgpsnote.html

Cockburn, Andrew, "Kill Chain: The Rise of the High-Tech Assassins," *Huffington Post*, Thought Matters blog, December 6, 2017. As of August 31, 2018:
https://www.huffingtonpost.com/thought-matters/kill-chain-the-rise-of-th_b_9540292.html

Committee on Autonomous Vehicles in Support of Naval Operations, *Autonomous Vehicles in Support of Naval Operations*, Washington, D.C.: National Academies Press, 2005. As of August 31, 2018:
https://www.nap.edu/read/11379/chapter/1

Common Objects in Context, homepage, undated. As of May 17, 2018:
http://cocodataset.org/#home

Conway, John H., "The Angel Problem," in Richard Nowakowski, ed., *Games of No Chance*, Vol. 29, Cambridge, UK: Cambridge University Press, 1996, pp. 3–12. As of August 31, 2018:
http://library.msri.org/books/Book29/files/conway.pdf

Cortes, Jorge, and Magnus Egerstedt, "Coordinated Control of Multi-Robot Systems: A Survey," *SICE Journal of Control, Measurement, and System Integration*, Vol. 10, No. 6, 2017, pp. 495–503.

Cryogenic Society of America, "HTS Degaussing System," *Cold Facts*, Vol. 25, No. 2, Spring 2009. As of May 17, 2018:
https://www.cryogenicsociety.org/resources/cryo_central/hts_degaussing_systems/

DARPA—*See* Defense Advanced Research Projects Agency.

Defense Advanced Research Projects Agency, "TALONS Tested on Commissioned U.S. Navy Vessel for First Time," press release, August 15, 2017. As of May 17, 2018:
https://www.darpa.mil/news-events/2017-08-15

Defense Science Board, *Task Force Report: The Role of Autonomy in DoD Systems*, Washington, D.C.: U.S. Department of Defense, July 2012. As of April 19, 2018:
https://fas.org/irp/agency/dod/dsb/autonomy.pdf

———, *Technology and Innovation Enablers for Superiority in 2030*, Washington, D.C.: U.S. Department of Defense, October 2013. As of October 5, 2018:
https://www.acq.osd.mil/dsb/reports/2010s/DSB2030.pdf

———, *Summer Study on Autonomy*, Washington, D.C.: U.S. Department of Defense, June 2016. As of April 19, 2018:
https://www.hsdl.org/?view&did=794641

Dektor, S., and S. Rock, "Improving Robustness of Terrain-Relative Navigation for AUVs in Regions with Flat Terrain," *2012 IEEE/OES Autonomous Underwater Vehicles (AUV)*, 2012, pp. 1–7.

DoD—*See* U.S. Department of Defense.

DSB—*See* Defense Science Board.

Durrant-Whyte, Hugh, and Tim Bailey, "Simultaneous Localization and Mapping: Part I," *IEEE Robotics and Automation Magazine*, Vol. 13, No. 2, 2006, pp. 99–110.

Džeroski, Saso, and Bernard Ženko, "Is Combining Classifiers with Stacking Better Than Selecting the Best One?" *Machine Learning*, Vol. 54, No. 3, March 1, 2004, pp. 255–273.

Eckstein, Megan, and Sam LaGrone, "Trio of Studies Predict the U.S. Navy Fleet of 2030," *USNI News*, February 14, 2017. As of August 31, 2018:
https://news.usni.org/2017/02/14/trio-of-studies-look-to-the-u-s-navy-fleet-of-2030

Elbanhawi, M., and M. Simic, "Sampling-Based Robot Motion Planning: A Review," *IEEE Access*, Vol. 2, 2014, pp. 56–77.

Fax, J. A., and R. M. Murray, "Information Flow and Cooperative Control of Vehicle Formations," *IEEE Transactions on Automatic Control*, Vol. 49, No. 9, 2004, pp. 1465–1476.

Feigenbaum, Edward A., and Julian Feldman, eds., *Computers and Thought*, London: McGraw-Hill, 1963.

Freedberg, Sydney J., Jr., "Navy Scraps RMMV Mine Drone, Accelerates CBARS," *Breaking Defense*, February 26, 2016.

Gang, Byeong Gyu, and Sejin Kwon, "Design of an Energy Management Technique for High Endurance Unmanned Aerial Vehicles Powered by Fuel and Solar Cell Systems," *International Journal of Hydrogen Energy*, Vol. 43, No. 20, May 2018, pp. 9787–9796.

Gillmer, Thomas C., and Bruce Johnson, *Introduction to Naval Architecture*, Annapolis, Md.: Naval Institute Press, 1982.

Goodfellow, Ian, Jonathon Shlens, and Christian Szegedy, "Explaining and Harnessing Adversarial Examples," ArXiv, March 20, 2015.

Gortney, W. E., and C. D. Haney, *Introduction to the Readiness Kill Chain*, Norfolk, Va.: U.S. Fleet Forces Command, April 2013. As of October 5, 2018:
http://www.public.navy.mil/usff/documents/rkc_booklet_final-w-signatures_11apr13.pdf

Guo, Y., M. Bennamoun, F. Sohel, M. Lu, and J. Wan, "3D Object Recognition in Cluttered Scenes with Local Surface Features: A Survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 36, No. 11, 2014, pp. 2270–2287.

Helmer, Scott, David Meger, Marius Muja, James J. Little, and David G. Lowe, "Multiple Viewpoint Recognition and Localization," in Ron Kimmel, Reinhard Klette, and Akihiro Suimoto, eds., *Computer Vision—ACCV 2010: 10th Asian Conference on Computer Vision*, Part 1, Berlin: Springer-Verlag, 2011, pp. 464–477.

Henriques, J. F., R. Caseiro, P. Martins, and J. Batista, "High-Speed Tracking with Kernelized Correlation Filters," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 37, No. 3, 2015, pp. 583–596.

Huang, Hui-Min, ed., *Autonomy Levels for Unmanned Systems (ALFUS) Framework*, Vol. 1: *Terminology*, Gaithersburg, Md.: National Institute of Standards and Technology, Special Publication 1011-I-2.0, October 2008. As of August 31, 2018:
https://www.nist.gov/sites/default/files/documents/el/isd/ks/NISTSP_1011-I-2-0.pdf

IEEE—*See* Institute of Electrical and Electronics Engineers.

ImageNet, "ImageNet Large Scale Visual Recognition Challenge (ILSVRC)," webpage, 2015. As of May 17, 2018:
http://www.image-net.org/challenges/LSVRC/

Institute of Electrical and Electronics Engineers, "Data Fusion Contest," webpage, undated As of May 17, 2018:
http://www.grss-ieee.org/community/technical-committees/data-fusion/data-fusion-contest/

Interpretable Machine Learning, "2V2 Debate: Caruana, Simard vs. Weinberger, LeCun, Interpretable ML Symposium, NIPS 2017," video, December 2017. As of August 27, 2018: https://www.youtube.com/watch?v=2hW05ZfsUUo

Jontz, Sandra, "DARPA Christens New Sea Robot Vessel Sea Hunter," *SIGNAL Magazine*, April 7, 2016. As of May 16, 2018: https://www.afcea.org/content/Article-darpa-christens-new-sea-robot-vessel-sea-hunter

Kaplan, Adam, *Path Planning and Energy Management of Solar-Powered Unmanned Ground Vehicles*, thesis, Ames, Iowa: Iowa State University, 2015. As of October 5, 2018: https://lib.dr.iastate.edu/cgi/viewcontent.cgi?article=5839&context=etd

Kearns, Kris, "DoD Autonomy Roadmap," briefing slides, National Defense Industrial Association 19th Annual Science and Engineering Technology Conference, March 21, 2018.

Keller, John, "Navy Asks Metron for Autonomy and Control Software for Future Large-Displacement UUV," *Military and Aerospace Electronics*, May 22, 2013. As of April 19, 2018: http://www.militaryaerospace.com/articles/2013/05/Metron-LDUUV-autonomy.html

Kinsey, James C., Ryan M. Eustice, and Louis L. Whitcomb, "A Survey of Underwater Vehicle Navigation: Recent Advances and New Challenges," *IFAC Conference of Manoeuvering and Control of Marine Craft*, 2006, pp. 1–12.

Klahr, Philip, and Donald A. Waterman, *Artificial Intelligence: A RAND Perspective*, Santa Monica, Calif.: RAND Corporation, P-7172, January 1986. As of August 31, 2018: https://www.rand.org/pubs/papers/P7172.html

Kristan, Matej, Ales Leonardis, Jiri Matas, Michael Felsberg, et al., "The Visual Object Tracking VOT2017 Challenge Results," *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*, 2017, pp. 1949–1972.

Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *NIPS '12: Proceedings of the 25th International Conference on Neural Information Processing Systems*, Vol. 1, New York: Curran Associates Inc., 2012.

L-3 Communications SeaBeam Instruments, *Multibeam Sonar Theory of Operation*, East Walpole, Mass., 2000. As of April 19, 2018: https://www3.mbari.org/data/mbsystem/sonarfunction/SeaBeamMultibeamTheoryOperation.pdf

Lagor, Francis D., Kayo Ide, and Derek A. Paley, "Incorporating Prior Knowledge in Observability-Based Path Planning for Ocean Sampling," *Systems and Control Letters*, Vol. 97, November 2016, pp. 169–175.

"Large Displacement Unmanned Underwater Vehicle Innovative Naval Prototype (LDUUV INP)," *NavalDrones*, April 13, 2015. As of April 19, 2018: http://www.navaldrones.com/LDUUV-INP.html

Leang, Isabelle, Stéphane Herbin, Benoît Girard, and Jacques Droulez, "On-Line Fusion of Trackers for Single-Object Tracking," *Pattern Recognition*, Vol. 74, 2017.

Leonard, Naomi Ehrich, and Alex Olshevsky, "Cooperative Learning in Multiagent Systems from Intermittent Measurements," *SIAM Journal on Control and Optimization*, Vol. 53, No. 1, 2015, pp. 1–29.

Lockheed Martin, "Naval Electronic Warfare," webpage, undated. As of May 17, 2018: https://www.lockheedmartin.com/en-us/capabilities/electronic-warfare/naval-ew.html

Luo, Juan, and Gwun Oubong, "A Comparison of SIFT, PCA-SIFT and SURF," *International Journal of Image Processing*, Vol. 3, No. 4, 2009, pp. 143–152.

Luo, Wenhan, Junliang Xing, Anton Milan, Xiaoqing Zhang, Wei Liu, Xiaowei Zhao, and Tae-Kyun Kim, “Multiple Object Tracking: A Literature Review,” ArXiv, May 22, 2017.

Ma, H., E. Smart, A. Ahmed, and D. Brown, “Radar Image-Based Positioning for USV Under GPS Denial Environment,” *IEEE Transactions on Intelligent Transportation Systems*, Vol. 19, No. 1, 2018, pp. 72–80.

Marr, Bernard, “The 4 Ds of Robotization: Dull, Dirty, Dangerous and Dear,” *Forbes*, October 16, 2017.

Marx, D., M. Nelson, E. Chang, W. Gillespie, A. Putney, and K. Warman, “An Introduction to Synthetic Aperture Sonar,” *Proceedings of the Tenth IEEE Workshop on Statistical Signal and Array Processing*, 2000, pp. 717–721.

Massot-Campos, Miquel, and Gabriel Oliver-Codina, “Optical Sensors and Methods for Underwater 3D Reconstruction,” *Sensors*, Vol. 15, No. 12, 2015, pp. 31525–31557.

Mastellone, Silvia, Dušan M. Stipanović, Christopher R. Graunke, Koji A. Intlekofer, and Mark W. Spong, “Formation Control and Collision Avoidance for Multi-Agent Non-Holonomic Systems: Theory and Experiments,” *International Journal of Robotics Research*, Vol. 27, No. 1, 2008, pp. 107–126.

McCarthy, John, Marvin Minsky, Claude Elwood Shannon, and Nathaniel Rochester, *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, Hanover, N.H.: Dartmouth College, 1955.

Medagoda, L., J. C. Kinsey, and M. Eilders, “Autonomous Underwater Vehicle Localization in a Spatiotemporally Varying Water Current Field,” *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 565–572.

Medby, Jamison Jo, and Russell W. Glenn, *Street Smart: Intelligence Preparation of the Battlefield for Urban Operations*, Santa Monica, Calif.: RAND Corporation, MR-1287-A, 2002. As of October 15, 2018:
https://www.rand.org/pubs/monograph_reports/MR1287.html

Multiple Object Tracking Benchmark, homepage, undated. As of May 17, 2018:
https://motchallenge.net

National Academies of Sciences, Engineering, and Medicine, *Handbook of Frequency Allocations and Spectrum Protection for Scientific Uses*, 2nd ed., Washington, D.C.: National Academies Press, 2015.

Neurath, Marie, *Machines Which Seem to Think*, London: Parrish, 1954.

Nilsson, Nils J., *The Quest for Artificial Intelligence: A History of Ideas and Achievements*, Cambridge, UK: Cambridge University Press, 2009.

Odasso, Cristina, Giuseppe Scellato, and Elisa Ughetto, “Selling Patents at Auction: An Empirical Analysis of Patent Value,” *Industrial and Corporate Change*, Vol. 24, No. 2, 2015, pp. 417–438.

Office of Naval Research, *Naval S&T Strategy: Innovations for the Future Force*, Washington, D.C., 2015. As of October 5, 2018:
http://www.dtic.mil/dtic/tr/fulltext/u2/1036755.pdf

Olfati-Saber, R., “Flocking for Multi-Agent Dynamic Systems: Algorithms and Theory,” *IEEE Transactions on Automatic Control*, Vol. 51, No. 3, 2006, pp. 401–420.

O’Rourke, Ronald, *China Naval Modernization: Implications for U.S. Navy Capabilities—Background and Issues for Congress*, Washington, D.C.: Congressional Research Service, RL33153, August 21, 2018. As of October 5, 2018:
https://fas.org/sgp/crs/row/RL33153.pdf

Ouimet, Michael, and Jorge Cortés, "Robust Coordinated Rendezvous of Depth-Actuated Drifters in Ocean Internal Waves," *Automatica*, Vol. 69, July 2016, pp. 265–274.

Papernot, Nicolas, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami, "The Limitations of Deep Learning in Adversarial Settings," ArXiv, November 24, 2015.

Paull, L., S. Saeedi, M. Seto, and H. Li, "AUV Navigation and Localization: A Review," *IEEE Journal of Oceanic Engineering*, Vol. 39, No. 1, 2014, pp. 131–149.

Payne, Craig M., *Principles of Naval Weapon Systems*, Annapolis, Md.: Naval Institute Press, January 2010.

Program Executive Office Littoral Combat Ship Public Affairs, "Large Displacement Unmanned Underwater Vehicle Program Achieves Acquisition Milestone," press release, Washington, D.C.: U.S. Navy, September 3, 2015. As of April 19, 2018:
http://www.navy.mil/submit/display.asp?story_id=90932

Public Law 114-328, National Defense Authorization Act for Fiscal Year 2017, December 23, 2016.

Raytheon, "AN/SLQ-32(V) Shipboard EW System," webpage, undated-a. As of May 17, 2018:
https://www.raytheon.com/capabilities/products/slq32

———, "Dual Band Radar," webpage, undated-b. As of May 17, 2018:
https://www.raytheon.com/capabilities/products/dbr

RE2 Robotics, "RE2 Robotics Wins Navy Contract to Develop Underwater Manipulator Arms," press release, February 7, 2017. As of May 17, 2018:
http://www.resquared.com/re2-robotics-develops-underwater-arms/

Rui, Gao, and M. Chitre, "Cooperative Positioning Using Range-Only Measurements Between Two AUVs," *OCEANS '10 IEEE Sydney*, 2010, pp. 1–6.

Safari, Sajjad, Faridoon Shabani, and Dan Simon, "Multirate Multisensor Data Fusion for Linear Systems Using Kalman Filters and a Neural Network," *Aerospace Science and Technology*, Vol. 39, December 2014, pp. 465–471.

Saligrama, Venkatesh, ed., *Networked Sensing Information and Control*, 1st ed., Boston, Mass.: Springer Science and Business Media, 2008.

Savitz, Scott, "Rethink Mine Countermeasures," *Proceedings Magazine*, Vol. 143, July 2017.

Savitz, Scott, Irv Blickstein, Peter Buryk, Robert W. Button, Paul DeLuca, James Dryden, Jason Mastbaum, Jan Osburg, Phillip Padilla, Amy Potter, Carter C. Price, Lloyd Thrall, Susan K. Woodward, Roland J. Yardley, and John M. Yurchak, *U.S. Navy Employment Options for Unmanned Surface Vehicles (USVs)*, Santa Monica, Calif.: RAND Corporation, RR-384-NAVY, 2013. As of August 31, 2018:
https://www.rand.org/pubs/research_reports/RR384.html

Scheidegger, Samuel, Joachim Benjaminsson, Emil Rosenberg, Amrit Krishnan, and Karl Granstrom, "Mono-Camera 3D Multi-Object Tracking Using Deep Learning Detections and PMBM Filtering," ArXiv, February 27, 2018.

Schmid, Jon, "The Diffusion of Military Technology," *Defence and Peace Economics*, February 17, 2017, pp. 1–19.

Schwager, M., J. J. Slotine, and D. Rus, "Decentralized, Adaptive Control for Coverage with Networked Robots," *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 3289–3294.

Scott, William Lewis, and Naomi Ehrich Leonard, “Optimal Evasive Strategies for Multiple Interacting Agents with Motion Constraints,” *Automatica*, Vol. 94, August 2018, pp. 26–34.

Shamma, Jeff S., ed., *Cooperative Control of Distributed Multi-Agent Systems*, New York: Wiley-Interscience, 2008.

Sheridan, Thomas B., *Humans and Automation: System Design and Research Issues*, Hoboken, N.J.: John Wiley, 2002.

Snyder, J., “Doppler Velocity Log (DVL) Navigation for Observation-Class ROVs,” *OCEANS 2010 MTS/IEEE Seattle*, 2010, pp. 1–9.

Solomon, Jonathan F., *Defending the Fleet from China’s Anti-Ship Ballistic Missile: Naval Deception’s Roles in Sea-Based Missile Defense*, thesis, Washington, D.C.: Georgetown University, April 15, 2011. As of April 19, 2018:
http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.454.8264&rep=rep1&type=pdf

Sontag, Sherry, Christopher Drew, and Annette Lawrence Drew, *Blind Man’s Bluff: The Untold Story of American Submarine Espionage*, New York: PublicAffairs Press, 1998.

Stojanovic, M., and J. Preisig, “Underwater Acoustic Communication Channels: Propagation Models and Statistical Characterization,” *IEEE Communications Magazine*, Vol. 47, No. 1, 2009, pp. 84–89.

Su, Housheng, Xiaofan Wang, and Guanrong Chen, “Rendezvous of Multiple Mobile Agents with Preserved Network Connectivity,” *Systems and Control Letters*, Vol. 59, No. 5, May 2010, pp. 313–322.

Sun, D., C. Wang, W. Shang, and G. Feng, “A Synchronization Approach to Trajectory Tracking of Multiple Mobile Robots While Maintaining Time-Varying Formations,” *IEEE Transactions on Robotics*, Vol. 25, No. 5, 2009, pp. 1074–1086.

Szegedy, Christian, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus, “Intriguing Properties of Neural Networks,” ArXiv, February 19, 2014.

Tang, Zhijun, and U. Ozguner, “Motion Planning for Multitarget Surveillance with Mobile Sensor Agents,” *IEEE Transactions on Robotics*, Vol. 21, No. 5, 2005, pp. 898–908.

Tarraf, Danielle C., ed., *Control of Cyber-Physical Systems*, Cham, Switzerland: Springer International Publishing, March 2013.

Test and Evaluation, Verification, and Validation Working Group, *Technology Investment Strategy 2015–2018*, Washington, D.C.: U.S. Department of Defense, Autonomy Community of Interest, May 2015. As of October 5, 2018:
http://www.dtic.mil/dtic/tr/fulltext/u2/1010194.pdf

Umpleby, Stuart, “Definitions of Cybernetics,” *Larry Richards Reader 1997–2007*, 1998.

U.S. Department of Defense, *Reliance 21: Operating Principles*, Washington, D.C., January 2014. As of April 19, 2018:
https://www.acq.osd.mil/rd/publications/docs/Reliance_21_Op_Principles_Jan_2014.pdf

“U.S. DoD Finds More LCS Mission Package Problems,” *Maritime Executive*, November 11, 2016. As of August 31, 2018:
https://www.maritime-executive.com/article/us-dod-finds-more-lcs-mission-package-problems#gs.8MZN16A

U.S. Navy, *Capability Development Document for Large Diameter Unmanned Undersea Vehicle System*, Washington, D.C., February 2014.

———, “Surface Electronic Warfare Improvement Program (SEWIP),” webpage, January 30, 2017. As of May 17, 2018:
http://www.navy.mil/navydata/fact_display.asp?cid=2100&tid=475&ct=2

Vasilijević, A., D. Nađ, F. Mandić, N. Mišković, and Z. Vukić, “Coordinated Navigation of Surface and Underwater Marine Robotic Vehicles for Ocean Sampling and Environmental Monitoring,” *IEEE/ASME Transactions on Mechatronics*, Vol. 22, No. 3, 2017, pp. 1174–1184.

Vidal, R., O. Shakernia, H. J. Kim, D. H. Shim, and S. Sastry, “Probabilistic Pursuit-Evasion Games: Theory, Implementation, and Experimental Evaluation,” *IEEE Transactions on Robotics and Automation*, Vol. 18, No. 5, 2002, pp. 662–669.

Visual Object Tracking Challenge, homepage, undated. As of May 17, 2018:
http://www.votchallenge.net

von Schuckmann, Karina, Pierre-Yves Le Traon, Enrique Alvarez-Fanjul, Lars Axell, et al., “The Copernicus Marine Environment Monitoring Service Ocean State Report,” *Journal of Operational Oceanography*, Vol. 9, Suppl. 2, September 2016, pp. s235–s320.

Wakayama, C., D. Grimmett, and R. Ricks, “Active Multistatic Track Initiation Cued by Passive Acoustic Detection,” *2012 15th International Conference on Information Fusion*, 2012, pp. 2577–2584.

Wang, Qiuying, Xufei Cui, Yibing Li, and Fang Ye, “Performance Enhancement of a USV INS/CNS/DVL Integration Navigation System Based on an Adaptive Information Sharing Factor Federated Filter,” *Sensors*, Vol. 17, No. 2, 2017.

Wang, Timothy, Romain Jobredeaux, Marc Pantel, Pierre-Loic Garoche, Eric Feron, and Didier Henrion, “Credible Autocoding of Convex Optimization Algorithms,” *Optimization and Engineering*, Vol. 17, No. 4, December 2016, pp. 781–812.

Wang, Y., E. Garcia, F. Zhang, and D. Casbeer, *Cooperative Control of Multi-Agent Systems: Theory and Applications*, New York: John Wiley & Sons, 2017.

Whitman, Edward C., “SOSUS: The ‘Secret Weapon’ of Undersea Surveillance,” *Undersea Warfare*, 2005. As of May 17, 2018:
http://www.public.navy.mil/subfor/underseawarfaremagazine/Issues/Archives/issue_25/sosus.htm

Wiener, Norbert, *Cybernetics or Control and Communication in the Animal and the Machine*, Cambridge, Mass.: MIT Press, 1961.

Wongpiromsarn, T., U. Topcu, and A. Lamperski, “Automata Theory Meets Barrier Certificates: Temporal Logic Verification of Nonlinear Systems,” *IEEE Transactions on Automatic Control*, Vol. 61, No. 11, 2016, pp. 3344–3355.

Woodman, Oliver J., *An Introduction to Inertial Navigation*, Cambridge, UK: University of Cambridge Computer Laboratory, Technical Report No. 696, August 2007.

Wynn, Russell B., Veerle A. I. Huvenne, Timothy P. Le Bas, Bramley J. Murton, Douglas P. Connelly, Brian J. Bett, Henry A. Ruhl, Kirsty J. Morris, Jeffrey Peakall, Daniel R. Parsons, Esther J. Sumner, Stephen E. Darby, Robert M. Dorrell, and James E. Hunt, “Autonomous Underwater Vehicles (AUVs): Their Past, Present and Future Contributions to the Advancement of Marine Geoscience,” *Marine Geology*, Vol. 352, June 2014, pp. 451–468.

Xiao, Feng, Long Wang, and Tongwen Chen, “Connectivity Preservation for Multi-Agent Rendezvous with Link Failure,” *Automatica*, Vol. 48, No. 1, January 2012, pp. 25–35.

Yan, Jin, Liao Yan, A. A. Minai, and M. M. Polycarpou, “Balancing Search and Target Response in Cooperative Unmanned Aerial Vehicle (UAV) Teams,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, Vol. 36, No. 3, 2005, pp. 571–587.

Yi, Kwang Moo, Eduard Trulls, Vincent Lepetit, and Pascal Fua, “LIFT: Learned Invariant Feature Transform,” ArXiv, July 29, 2016.

Yoon, M. K., B. Liu, N. Hovakimyan, and L. Sha, “VirtualDrone: Virtual Sensing, Actuation, and Communication for Attack-Resilient Unmanned Aerial Systems,” *2017 ACM/IEEE 8th International Conference on Cyber-Physical Systems (ICCPS)*, 2017, pp. 143–154.

Zhang, H., F. L. Lewis, and A. Das, “Optimal Design for Synchronization of Cooperative Systems: State Feedback, Observer and Output Feedback,” *IEEE Transactions on Automatic Control*, Vol. 56, No. 8, 2011, pp. 1948–1952.

Zhao, Z. C., X. S. Wang, and S. P. Xiao, “Cooperative Deception Jamming Against Radar Network Using a Team of UAVs,” *2009 IET International Radar Conference*, 2009, pp. 1–4.

The U.S. Navy is interested in developing autonomous capabilities to execute tasks that are increasingly hazardous for humans and to enhance warfighting capabilities. In this report, RAND researchers explore current and potential military applications of autonomous systems, focusing especially on unmanned undersea vehicles and unmanned surface vehicles. The analysis centered on four areas: the current state of the art of autonomous technology, current kill chains and capabilities, future fleet architecture and its autonomous capabilities, and autonomy in alternative concepts of operation. The authors conclude that, although technological advances have occurred, autonomous systems capable of responding to unexpected changes in the environment do not yet exist and may not for some time. Any development of these capabilities will require substantial military investment because commercial systems are unlikely to meet the Navy's needs. The authors recommend that the Navy revisit assumptions about technological progress in autonomy, align the development of autonomy with the development of other capabilities that may be limiting factors, develop new concepts of operation to take advantage of autonomy's key characteristics, and reevaluate force requirements in light of the state of autonomous technology.

RAND NATIONAL DEFENSE RESEARCH INSTITUTE

www.rand.org

RR-2751-NAVY