

Department of Defense
RESEARCH & ENGINEERING

TECHNICAL ASSESSMENT: AUTONOMY



Office of Technical Intelligence
Office of the Assistant Secretary of Defense
for Research & Engineering
February 2015



Distribution A: Approved for public release; distribution is unlimited



About the Office of Technical Intelligence

The Office of Technical Intelligence (OTI) provides the U.S. Department of Defense Research and Engineering community and partners holistic, defense-relevant insights into emerging and potentially disruptive technology to enable U.S. and mitigate adversary technological surprise. To do so, OTI identifies emerging and potentially disruptive science and technology, recommends efficient research and development (R&D) strategies, and coordinates intelligence collection, analysis and dissemination to inform research and engineering activities. OTI accomplishes these missions through three complimentary efforts: technology watch and horizon scanning, technical assessments, and tailored intelligence support and coordination.

OTI technology watch and horizon scanning efforts are developing methods to identify nascent and disruptive science, technology, and capabilities through the exploitation of tailored approaches and tools, including analysis of scientific literature, patents, and worldwide investment using both open source and internal data.

OTI technical assessments provide decision-relevant R&D strategy inputs on emerging and potentially disruptive technologies to the Research and Engineering community by exploring opportunities and threats the technologies could enable, conducting data-driven analyses of drivers to forecast future trends and identify unique DoD needs, recommending specific investment and policy approaches, and developing and seed funding projects to leverage those opportunities.

OTI intelligence support activities are focused on coordinating efforts across the Research and Engineering community, ensuring timely and valuable analysis reaches users, and providing mechanisms to enhance communication between policymakers, researchers, and analysts.

About this Assessment

Interviews Conducted: 11/13/2014 – 1/16/2015

Investment Data Finalized: 1/23/2015

Table of Contents

Executive Summary.....	iii
Introduction	1
Enablers of Autonomous Systems	1
Challenges from the Security Environment & Benefits of Autonomy	4
Identifying Defense R&D and Policy Priorities.....	9
Efficient Investment and Policy Approaches	13
Conclusion.....	21
Appendix A: Commercial Investment Insights: Autonomy	22

Executive Summary

U.S. and foreign technology and capability development is pushing existing human-machine systems to the edge of their abilities by introducing extreme timescales, high levels of complexity, severe risk to warfighters, and increasing costs. While these trends and the challenges they pose to the U.S. Department of Defense (DoD) do not appear likely to abate, autonomy has the potential to enable U.S. forces to break out of current limitations by allowing systems to understand the environment, to make decisions, and to act more effectively and with greater independence from humans. In doing so, autonomy can augment or replace humans to enhance performance, to reduce risk to warfighters, and to decrease costs.

This assessment identifies research and development (R&D) and policy opportunities to position DoD to more effectively leverage autonomy. Based on an analysis of the security environment, opportunities presented by autonomy, and private sector investment, there are four major gaps in DoD efforts to date:

1. There is no unified analytic framework to examine needs and opportunities for autonomy across DoD tasks and missions.
2. Few DoD efforts are conducting R&D, carrying out experimentation, or developing approaches to testing for systems to operate against intelligent adversaries.
3. While there is substantial interest in autonomy to enhance capabilities and to decrease risk to warfighters, there is relatively little focus on leveraging autonomy explicitly to decrease costs.
4. There is insufficient R&D, experimentation, and policy for developing architectures, concepts of operations, and test, evaluation, verification, and validation approaches to ensure future systems are affordable and can operate effectively as a joint force.

The assessment addresses these gaps by providing a capability-focused analytic framework that applies across mission spaces and by making recommendations in the key technology and policy areas that are critical to ensure that the U.S. maintains a superior and affordable force:

Technology: Autonomy relies on three multidisciplinary technical fields: perception, cognition, and action, which cover areas from sensors to artificial intelligence and robotics. There are opportunities to leverage private sector investment where applications overlap and in technology for more permissive environments, but DoD has unique, critical needs for technologies to enable operations in complex, adversarial environments.

Recommendation: Leverage private sector activity in low-cost aerial systems, data analytics, cyber defense, human-machine interaction, and efficiency-related technologies, and focus DoD perception, artificial intelligence, and robotics R&D on developing autonomy for platforms intended to operate in complex environments with special consideration for adversarial behaviors.

Modularity & Interoperability: Enabling the reuse and reconfiguration of both hardware and software from different systems in a modular fashion will play an important role in determining R&D costs of new platforms. This will also affect the flexibility of U.S. forces to tailor hardware and software to mission needs. Likewise, interoperability standards for communication can play an important role in enabling synergistic effects across platforms and domains. These needs are not unique to autonomous systems, but they will be critical to maximizing the effectiveness of their widespread implementation.

Recommendation: Fund research to develop a forward looking open architecture and interoperability standards for autonomous systems and mandate cross-Service compliance.

Resiliency: The fact that autonomous systems ‘think’ differently than humans will open up new vulnerabilities, as the U.S. and adversaries can take advantage of the shortcomings of machine perception and cognition. These systems will not necessarily be more vulnerable than human-controlled systems, but they will be vulnerable in different ways. This is particularly challenging because widespread adoption of a large number of autonomous systems with similar or identical perception and cognition systems raises the potential for one or a small number of weaknesses to endanger a large proportion of the force, as with agricultural monocultures and disease susceptibility.

Recommendation: Study the unique vulnerabilities of autonomous systems, and intentionally design heterogeneity into classes of U.S. systems to mitigate system-wide vulnerabilities.

Concepts of Operations Development & Experimentation: Autonomy will enable and benefit from new concepts of operations (CONOPS) by making it possible for systems to operate in environments, at levels of performance, and in new configurations which have not been possible to date. Autonomy will also impact the behavior of humans and raise new ethical and legal challenges. As with other technologies, the most effective CONOPS will not be immediately evident and will require experimentation to identify. At the same time, almost all DoD efforts are focused on enhancing capabilities and warfighter protection, but autonomy also presents major cost-saving opportunities in areas such as logistics, maintenance, and data analysis.

Recommendation: Fund intensive, adversarial experimentation in realistic environments to inform CONOPS development, paying special attention to artificial intelligence’s propensity towards unconventional approaches to problems; opportunities to employ larger numbers of lower-cost systems; the interaction between humans and autonomous systems; and ethical-legal considerations. Complement these efforts with programs to experiment with and implement autonomy explicitly to reduce costs.

Test, Evaluation, Verification, & Validation (TEV&V): Because of the complexity of autonomous systems, especially those that can learn, there will also be major challenges to carrying out effective TEV&V. Complex systems will render current approaches that rely on holistic testing infeasible because it will be impossible to test all possible circumstances, especially in laboratories where the environment cannot sufficiently replicate the real world. As a result, there are open needs for metrics, standards, methodologies, and appropriate environments to conduct TEV&V for autonomous systems.

Recommendation: Fund research into metrics, standards, methodologies, and approaches for TEV&V, including modeling, simulation, and licensing, and establish a dedicated air-land-sea autonomy range to ensure a realistic environment for TEV&V, which can also support experimentation.

Introduction

U.S. and foreign technology and capability development is pushing existing human-machine systems to the edge of their abilities by introducing extreme timescales, high levels of complexity, severe risk to warfighters, and high costs. While these trends and the challenges they pose to the U.S. Department of Defense (DoD) do not appear likely to abate, autonomy has the potential to enable U.S. forces to break out of current limitations by allowing systems to understand the environment, to make decisions, and to act more effectively and with greater independence from humans. In doing so, autonomy can augment or replace humans to enhance performance, to reduce risk to warfighters, and to decrease costs. Because of the importance of its potential impacts, autonomy has drawn substantial attention in the defense community, but many DoD efforts are still nascent and, as of yet, uncoordinated.

The goal of this assessment is to identify research and development (R&D) and policy opportunities for DoD to more efficiently and effectively leverage the benefits of autonomy. It starts by defining the autonomy space in terms of the technology and policy areas that influence the performance and cost of autonomous systems. It then identifies opportunities for DoD to benefit from autonomy, the characteristics of the environments in which those systems will operate, where private sector activities will fulfill DoD needs, and the areas in which DoD will require targeted investment due to unique challenges. The final section synthesizes the technology and policy drivers with DoD needs, private sector activity, and the state of technology development to identify key gaps and to make recommendations that can better position DoD to reap a full measure of the benefits of autonomy.

Enablers of Autonomous Systems

This section describes the factors that influence the performance and cost of autonomous systems. They represent a broad range of technology, design, and policy factors that enable the development of individual systems, architectures to integrate multiple systems, and approaches to employing autonomous systems as part of the broader force. This discussion frames the field of autonomy and provides a rubric, which the final section of this report uses to highlight gaps and opportunities for R&D and policy activities.

Individual Systems

At the most basic level, autonomy draws on three broad, multidisciplinary technical fields: perception, cognition, and action.¹ Perception enables systems to understand their environment through a combination of sensors, which provide data, and algorithms to turn those data into contextual understanding. For example, a ground robot attempting to traverse a field needs to observe and then understand the terrain in front of it before it can decide whether it is passable. Higher-level perception

¹ While some in the field draw a distinction between automation as more rigid and autonomy as able to operate under higher levels of complexity or uncertainty, we do not differentiate between the two because the distinction falls away when viewing them as enablers of DoD capabilities. Even if there is a difference of degree between them, they both perform the same function of enabling non-human decisions and actions.

includes understanding more complex and valuable attributes, such as individuals' intentions or whether systems pose a threat.

Cognition represents the ability of a machine to make decisions and relies on a system's computational hardware and artificial intelligence algorithms. These decisions can range from simple decision trees, such as in a vending machine – accept and verify payment, register selection, and dispense item – to more complex analysis which involves many factors that are unknown ahead of time. Returning to the example of the ground robot, once it has perceived the terrain in its environment, the robot must make judgments as to the advantages and disadvantages of certain routes and choose a course of action that best allows it to accomplish its mission. The level of complexity of the environment, the task at hand, and the broader mission all influence the level of difficulty of making effective decisions.

Action is a broad category which describes the ability of autonomous systems to carry out decisions, whether through physical or digital means. While autonomous systems are traditionally conceived of in terms of individual robotic systems, they can also take the form of multi-robot systems in which each acts as part of an 'organism,' and importantly, autonomy will enable systems without robotics, which are likely to be as or more common than robotic systems.² Software that analyzes information, such as programs to process large amounts of sensor data in order to present curated selections to humans or to monitor a computer network to block malicious code, are examples of digital action systems without a robotic component.

Perception, cognition, and action needs will vary depending on the goals of the system, the mission, and the operating environment. For instance, a system designed for operation in a relatively well-understood environment will require less perception, but it may still require complex decision making if facing a challenging problem, and action may range from displaying information to a human pilot to executing complex flight maneuvers based on whether a human is connected to the system. The degree of autonomy may also vary over time in a single system, as it comes in and out of contact with operators or requests support. The dynamic relationship between technical needs and capability benefits demonstrates that *the level of autonomy of a system is not a goal in and of itself.*³ Rather, *autonomy is a capability driver that DoD can design into military systems* in conjunction with human roles to produce a more effective and affordable force.

Viewing the capabilities of autonomous systems within the context of human-machine systems recognizes a critical role for humans. In order to better capitalize upon the relative strengths of humans and machines, autonomous systems will operate on a spectrum from *tightly coupled* to *loosely coupled*. Where

² Here, the term robot is meant to encompass any mechanical system without a human 'onboard.' This could include forms such as Boston Dynamics' Atlas, a humanoid robot that played a large role in the DARPA Robotics Challenge, as well as systems like the MQ-1 Predator unmanned aerial vehicle

³ This is important to note because a number of organizations analyze autonomy through the lens of the level of autonomy of systems. For example, the U.S. Department of Transportation (DoT) uses levels of autonomy to classify automobile technologies. This approach may be appropriate for DoT because cars are fairly standardized and operate in a relatively constrained environment, which means that their level of autonomy provides a good proxy for capability. However, DoD enjoys neither of these simplifying conditions for most mission areas. As such, comparing levels of autonomy may be misleading, and military benefits should be the key metric.

human performance provides benefits relative to machine perception, cognition, or action and where risk to warfighters is acceptable, autonomous systems will benefit from being more tightly-coupled with humans. Often, tightly coupled systems will feed information directly to humans, such as automatically cueing a warfighter to a threat or analyzing large amounts of imagery and presenting only images of note to a human analyst. For missions where humans are less effective or which are too dangerous and remote control is not feasible, more loosely-coupled systems will operate with less input from human operators; these more closely fit the traditional conception of autonomous systems – those that operate while out of touch with humans, such as a strike platform in a communications-denied environment.

Even the most loosely coupled systems will still interact with humans at various points in missions though, whether while being programmed, during mission execution, or upon return. As a result, both tightly and loosely coupled systems will require the ability to interact with humans effectively. The benefits of autonomy will come not from fully replacing human decision making, but from augmenting it or elevating it to higher levels of abstraction, such as moving the role of humans from mission execution to planning.⁴ As such, human-machine interaction is another key technology area supporting autonomy.

Architectures

Beyond the technologies that enable autonomy for individual systems, there are architectural factors that will substantially affect the costs and benefits of autonomous systems. While not unique to autonomous systems, the degree to which systems are modular – allowing for the reuse and reconfiguration of hardware and software from different platforms and vendors – will play an important role in determining acquisition costs of families of platforms. This will also affect the flexibility of U.S. forces to tailor hardware and software to mission needs. Without common software and hardware standards that enable this, DoD will continue to pay to replicate R&D efforts.

Likewise, interoperability standards for communication can play an important role in enabling synergistic effects from autonomous systems. This will allow greater complexity of operations as autonomous systems and humans can coordinate activity and solve problems collaboratively. Examples from current DoD systems demonstrate that imposing interoperability standards before developing systems is key to prevent major challenges and costs for retrofitting or redesigning systems.

The fact that autonomous systems ‘think’ differently than humans will open up new potential challenges to the resiliency of military forces, as adversaries may be able to take advantage of the shortcomings of machine perception and cognition. This is not to say that autonomous systems will necessarily be more vulnerable than human-controlled systems, but they will be vulnerable in different ways. If there is widespread adoption of a large number of autonomous systems with similar or identical perception and cognition systems, this raises the potential for one or a small number of weaknesses to endanger a large

⁴ Highlighting the role for manned autonomous systems is also important because discussion of autonomous systems is often intermixed with discussion of unmanned systems. This can be misleading, as manned systems can integrate substantial autonomy, for example to manage defensive systems, and unmanned systems may be almost totally human-controlled, as is the case with most of today’s unmanned aerial and ground systems – sometimes appropriately referred to as remotely-piloted vehicles.

proportion of the force. This is typically a less serious problem with manned systems because each human operator is idiosyncratic.

The risk of deploying a large number of similar systems is akin to the risk of monoculture in farming. Most large-scale farmers grow only a single variety of the same plant at the same time because it is typically easier to manage and the most productive under given conditions; however, this also means that the appearance of a harmful pest or disease can threaten the entire crop at once, instead of only a small proportion if different varieties are not equally susceptible. Open architecture and interoperability standards may make this challenge more severe if they propagate vulnerabilities throughout the force.

Employment

Concepts of operations (CONOPS) and test, evaluation, verification, and validation (TEV&V) will also play key roles in determining the value of autonomous systems and the rate at which they enter the force. Autonomy will open up new types of missions and the potential for systems to operate in environments, at levels of performance, or in configurations which have not been feasible to date. By adding machine perception and decision-making abilities to current and future systems, it will also be possible to operate on different timescales, to coordinate more closely, and to approach challenges using machine instead of human strengths and approaches. Thus, new CONOPS may dramatically increase the effectiveness of employing autonomous systems; however, if experience with other technologies holds, the most effective CONOPS will not be immediately evident and will require experimentation to identify.

Because of the complexity of autonomous systems, especially those that can learn, there will also be major challenges to carrying out effective TEV&V. Complex systems will render current approaches that rely on holistic testing infeasible because it will be impossible to test all possible environments and situations, especially in a laboratory, where the artificiality of the environment can introduce significant artifacts into the process. As a result, the development of metrics, standards, and methodologies to conduct TEV&V for autonomous systems will be key to enable their timely integration into the force.

While this section provides a rubric to understand the range of inputs that determine the impact of autonomous systems on a military force, DoD's specific needs will vary based on particular applications. The following section analyzes the challenges posed by the operational environment to identify the applications where DoD can apply autonomy to generate the greatest benefits.

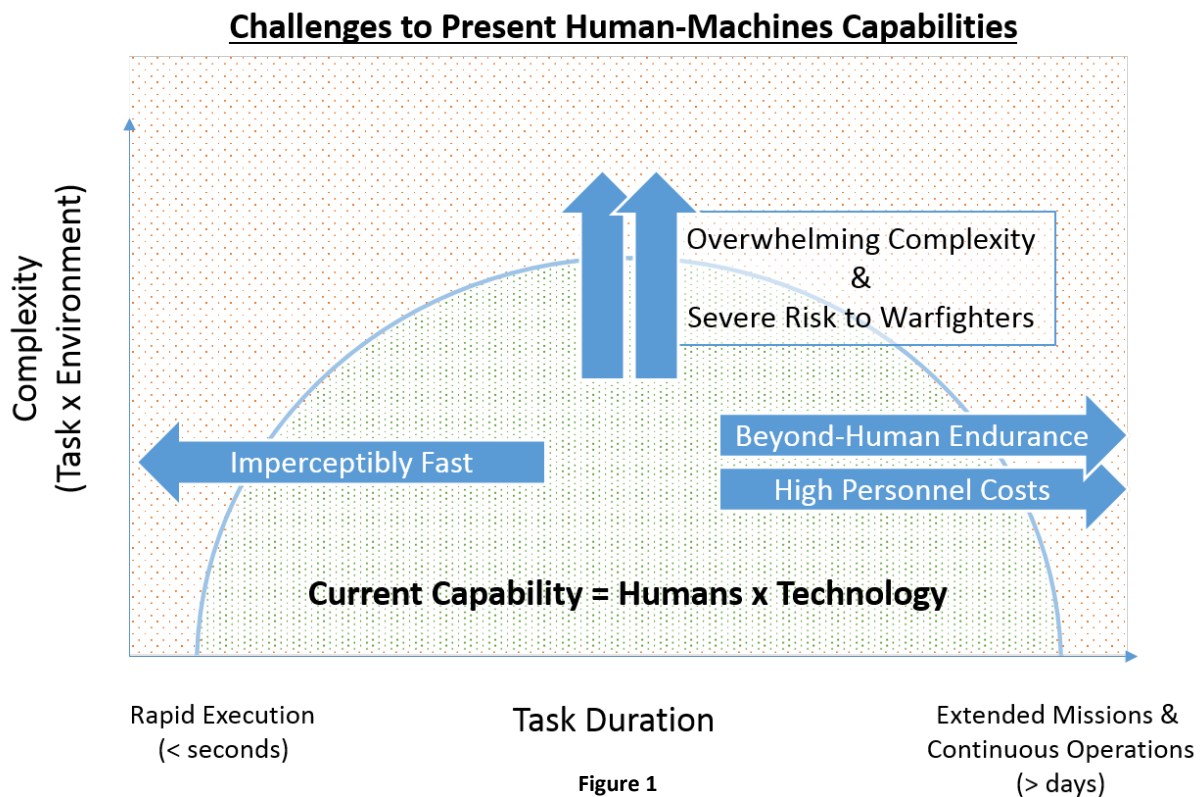
Challenges from the Security Environment & Benefits of Autonomy

The current and likely future operational environments pose a range of challenges to U.S. capabilities. U.S. and foreign technology and capability development is pushing existing human-machine systems to the edge of their abilities, especially in terms of extreme timescales, high levels of complexity, risk to warfighters, and cost. For example, some aspects of cyber operations require millisecond or faster reactions, and machines, but not humans, can endure very long duration missions, as may be beneficial for intelligence, surveillance, and reconnaissance (ISR) activities. Even on timescales where humans operate effectively, collection systems are generating more data than analysts can manage, and anti-

access and area-denial (A2/AD) capabilities are increasing the complexity of planning and executing operations to the point where current systems are moving toward their limit of effectiveness.

Even though certain missions may be technically feasible, the proliferation of advanced military systems to potential adversaries may make undertaking them undesirable due to the risk to U.S. military personnel, such as the threat of advanced air defense systems potentially driving commanders to keep manned platforms beyond their range. Similarly, where it has been possible to overcome some of the bounds on human performance by using more humans, such as by rotating crews to cover extended 24-hour operations, rising costs and limited budgets are rendering this unsustainable.⁵

While these challenges are understood in the defense community, there is no overarching framework in which to analyze the needs for and benefits of autonomy. This impedes collaboration between efforts in various mission areas and across Services. The following graphic visualizes a framework that describes these challenges and applies across mission spaces by analyzing tasks in terms of complexity and duration. The ‘inverted U’ represents the current bounds of human-machine performance as limited by technical, human performance, budgetary, and ethical factors, and the arrows demonstrate the factors that are challenging the effectiveness of current systems. The following section applies this framework to assess the potential benefits of autonomy.



⁵ Human performance modification technologies provide an additional avenue through which to achieve some performance goals, but they are beyond the scope of this assessment.

Benefits of Autonomy

While the trends pushing U.S. warfighters, systems, and budgets towards their limits do not appear likely to abate, autonomy promises to enable U.S. forces to break out of their current boundaries. By augmenting or replacing humans, autonomous systems have the potential to drive more effective decision making on broader timescales under higher levels of complexity with lower risk to warfighters and to decrease costs. The following sections describe the potential for autonomy to drive *performance enhancement*, *risk reduction*, and *cost reduction* beyond what is possible with today's human-machine systems.

Performance Enhancement

At the edge of or beyond the 'inverted U,' missions and tasks require performance on extreme timescales or under high complexity and typically require operations against intelligent adversaries. These represent challenges to current systems and opportunities to employ autonomy to increase performance.⁶

For some aspects of cyber operations and hypersonic warfare, there are windows of opportunity to act that are clearly too short for effective human intervention. In this case, adding autonomy to a system will enable computers to provide otherwise unachievable decision speed and quality. If one side more effectively leverages autonomy in this manner, it will provide the potential to 'get inside' the adversary's decision-action cycle and influence or defeat the adversary's courses of action.⁷ Air and missile defense systems have been some of the early adopters of autonomous capabilities due to the short timelines required for successful intercepts.

At the other end of the task-duration spectrum, autonomy will enable very long missions that would otherwise be impossible due to human endurance, power requirements, and communications challenges that would not allow for continuous human presence or monitoring. Using ISR as an example, it will be possible to place autonomous systems in the environment that collect information for periods of weeks to years. Very long-duration missions have substantial promise because they open up opportunities to characterize normal background patterns of activity with greater depth compared to the relative snapshots provided by many ISR systems today, enabling better anomaly detection and planning.

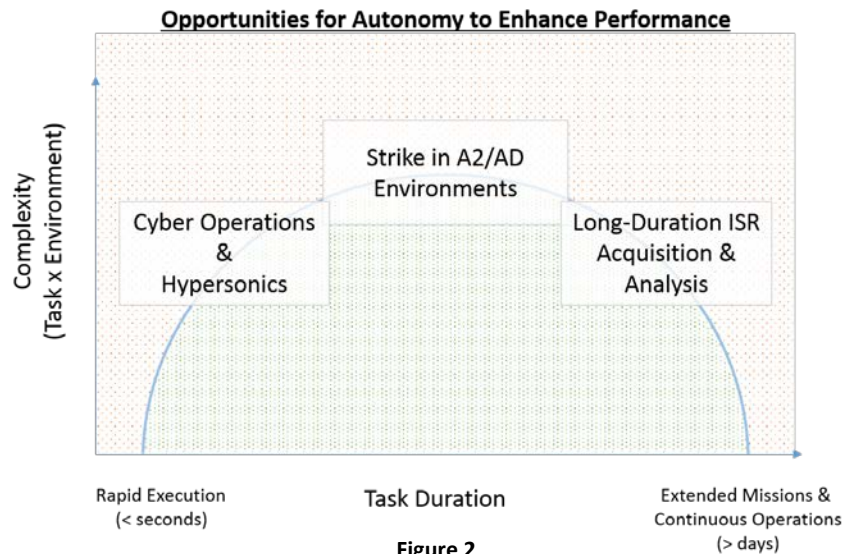


Figure 2

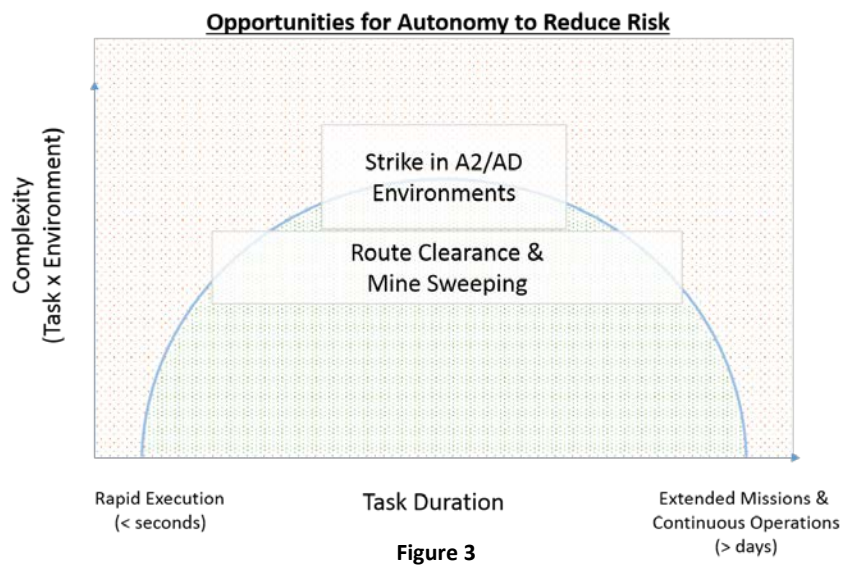
⁶ In each section, we discuss examples of missions and tasks, but these are not meant to be comprehensive.

⁷ Sometimes referred to as the observe, orient, decide, and act (OODA) loop.

Along with improved performance on extreme timescales, autonomy can offer benefits in missions involving high levels of complexity. For operations in a sophisticated A2/AD environment, strike assets are likely to face a combination of high-performance integrated air-defense systems, jamming, mobile targets, and a range of other challenges. To overcome this, successful attacks may require complex tasking and precise coordination. Due to the number of factors that affect the success of a mission under these circumstances, humans alone will lose the ability to perform well unaided, and autonomous systems can support or supplant human activities. They will do this by, for example, identifying and fulfilling information sharing needs to support situational awareness; performing ad-hoc network creation for shared communications; and providing decision support for targeting, navigation, and weapons delivery. Implementing autonomy will also present opportunities to expand the use of unmanned systems, which will enable designs with greatly increased aero- or hydrodynamics and that could operate at much higher performance, such as G-loads, because there is no need to account for the human form or frailties.

Risk Reduction

Risk considerations often place restrictions on the operational capabilities of deployed forces. Especially in complex threat environments, commanders may choose to limit the use of manned systems to protect U.S. military personnel. This risk calculus most strongly impacts missions that sit at high-levels of complexity in intermediate timescales, such as those in Figure 3. These missions are challenging, but manned systems have traditionally undertaken them. Like with performance enhancement, these environments also tend to be highly adversarial and require operations against intelligent adversaries.



Like with performance enhancement, these environments also tend to be highly adversarial and require operations against intelligent adversaries.

Loosely-coupled autonomous systems offer one means of achieving missions with greater protection because they can remove the need for a human presence. The military already uses unmanned ground vehicles as a partial measure in this regard to be able to approach potential improvised explosive devices (IED) closely, although IED disposal still requires local human operators because current systems have very limited perception and cognition abilities.

Beyond simply removing a human from danger, autonomy can provide substantial benefits for manned platforms in more tightly coupled applications. Autonomous systems that enable faster decisions and actions will typically lead to shorter exposures to threats, and autonomy that improves performance of defensive systems on a platform will enhance their ability to counter threats. For example, improved real-time route planning could avoid threats, and autonomous defensive systems could react to new ones as

they appear. It may also be possible to increase protection by using autonomy to increase the survivability of platforms. For example, NASA has a program to add autonomous software components to the fly-by-wire system of an F-15 in order to allow a human to maintain control when the aircraft is damaged.

Cost Reduction

Independent of improving capabilities and protection, autonomy has the potential to enable major cost savings. Figure 4 highlights several opportunities for DoD to cut costs in various areas. Here, the primary application of autonomy is for long-duration and continuous-operation missions, which tend to require large numbers of people to execute. In comparison to applications of autonomy for performance enhancement and risk reduction, most environments in which systems would need to operate to reduce costs and the tasks they would complete are more structured and less adversarial than the battlefield.

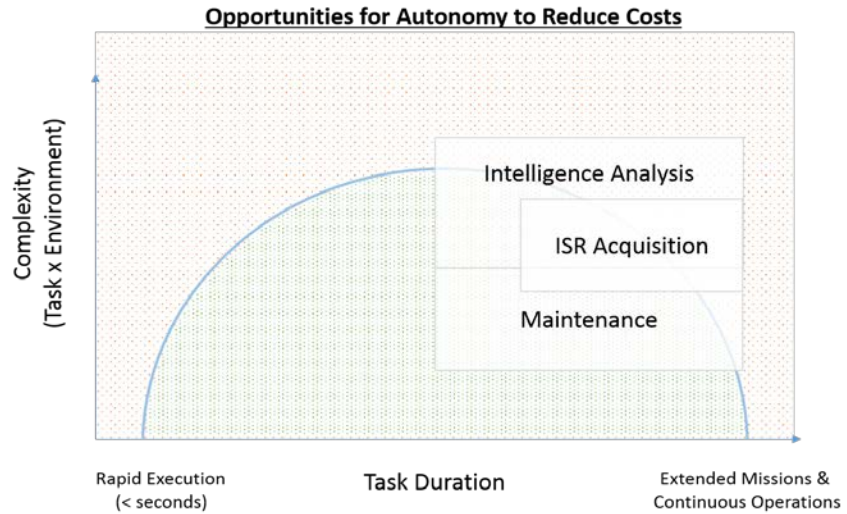


Figure 4

For example, the volume of information produced by intelligence collection systems over time can be enormous, requiring a large number of analysts to process effectively, and the high cost of training, salaries, and benefits translates this into a large burden for DoD. Autonomous analytic systems have the potential to sift through large amounts of data to cue analysts to important information, decreasing the number of humans required, especially for low-level tasks.

Maintenance is another area in which the application of autonomous systems could substantially decrease costs. At present, maintenance schedules are based on expectations of the lifetime of parts and the effects of wear and tear or routine inspections. A system capable of self-diagnosing and communicating its maintenance needs could more accurately inform these schedules as well as aid maintenance personnel to perform tasks more quickly and effectively. This would allow the military to accrue cost savings from decreases in personnel needs, wasted parts, and damage to systems from mistakes, while maximizing operational availability. Likewise, logistics, especially surrounding storing, finding, and shipping equipment from supply depots, and various other support missions could greatly benefit from the application of autonomy to reduce their footprint, while maintaining readiness.

As with performance, autonomy can also enable cost savings through new system design parameters. While manned submarines are effective ISR platforms, they are expensive. One factor in their high cost is that they must be large enough to accommodate and sustain the crew. Due to bandwidth restrictions, developing remotely-operated submarines is a serious challenge, but autonomy opens up the potential

for underwater systems which do not require a crew – or take a much smaller crew – and which therefore might be smaller and less expensive.

Complex Interactions

Implementing of autonomy will also cause interactions between cost, risk, and capability. For instance, increasing the loiter time of an autonomous ISR platform may improve performance, but it might also increase the costs of intelligence analysis if humans need to filter the data feeds. These interactions can also be beneficial. For example, many flight-line operations on aircraft carriers are still managed manually, and the Naval Safety Center reports that most major aviation mishaps are caused by human error.⁸ If, instead, a system comprised of algorithms and sensors could monitor activity on the deck, manage movement, and provide warnings autonomously, the ship could mitigate the risks to personnel, decrease costs from accidents, and improve the overall performance of flight deck operations.

Implementing autonomy to remove warfighters from danger will also enable increased performance through new CONOPS. Without a human onboard, for instance, an autonomous system could choose to fly directly into the engagement envelope of an air-defense system without concern for personal safety. Thus, the interactions between performance, protection, and cost will require careful analysis to enable the greatest synergies and to balance the benefits and drawbacks of implementing autonomy.

Identifying Defense R&D and Policy Priorities

The opportunities in performance, risk, and cost suggest that DoD can reap tremendous benefits from developing a new generation of systems leveraging autonomy. This will require substantial R&D and policy efforts, however, because, as described earlier, autonomous systems require component technologies, architectures, CONOPS development, and testing in order to build effective platforms, to integrate multiple platforms, to understand how to employ them effectively, and to field robust systems. This is even more challenging when considering that DoD must successfully address each of these substantial challenges in the face of budget constraints. As a result, it is critical to take an efficient approach to R&D and acquisition. To do so, DoD should leverage external investment where possible. The following section analyzes commercial interest in autonomy to identify overlapping and unique needs to inform investment prioritization.

Opportunities to Leverage the Private Sector & Unique DoD Needs

The private sector can be a key source for technology and capability development where overlapping needs exist, as it is developing a range of technologies relevant to autonomy for commercial applications. It also often develop technology at lower cost and with faster design cycles, which can provide benefits if DoD can identify and leverage these opportunities. As such, this section analyzes convergent and divergent applications for DoD and the private sector to identify where DoD should collaborate or purchase commercial technology and where DoD must invest because commercial technologies are unlikely to solve defense challenges.

⁸ Navy Safety Center Annual Mishap Overview FY 13, p. 8.

To understand trends in the commercial space, the Office of Technical Intelligence interviewed individuals involved in commercial applications of autonomy and analyzed data on venture capital and corporate investment in private companies in relevant technology areas.⁹ Because private companies tend to be at an earlier stage in their development, trends in formation, foci, and ability to attract investment are signals about the future contours of the commercial space. Additional analysis of private sector investment trends is included as Appendix A.

Convergent Applications

In terms of performance enhancement, private sector needs largely converge with DoD interests in the areas of large-scale data processing and cyber defense. Considering the value that companies already generate from analyzing large amounts of user data and extracting actionable information – such as from web browsing and buying behavior – there is substantial commercial interest in high-end information analytic systems which include autonomous components. Current private sector trends towards targeted advertising and assistance in diagnosing disease mirror DoD needs in terms of delivering real-time or near real-time analysis. Likewise, in the cyber defense field, the private sector is investing in real-time autonomous authentication and detection in order to prevent unauthorized intrusions into networks, which have direct DoD applications.

The commercial interest in these areas suggests that the private sector is likely to develop technologies that are valuable to DoD. An analysis of venture capital data shows that over 100 private companies are working in data analysis areas relevant to DoD needs, and these have attracted almost \$900 million in private investment from 2010-2014. For cyber defense, private companies founded between 2010-2014 have received more than \$200 million in private investment. Given the overlaps in these application areas, these data suggest that there is enough private sector investment in order to develop technologies of substantial interest to DoD.

While there are fewer parallels between DoD and the private sector in risk reduction, needs do converge for maintenance operations in harsh environments. Industries such as oil and gas, power generation and distribution, and telecommunications conduct maintenance underground, undersea, and in other dangerous environments which can pose threats to human workers. These threats will likely incentivize the development of both tightly and loosely coupled autonomous systems to augment or to replace humans in dirty and dangerous jobs, which are similar to DoD applications, such as ship maintenance. Because of this, DoD will likely have some opportunities to leverage commercial technology where the threat originates less from adversarial activity and more from challenging environments.

In terms of cost reduction, defense and commercial opportunities for autonomy are very similar. Increasing efficiency is a core focus of companies, and application areas such as information analysis,

⁹ The Office of Technical Intelligence leveraged software and databases from Quid, Inc. to support this analysis. By comparing the similarity of company descriptions from investment documents and websites with search terms of interest, Quid's platform identified relevant companies that have received private investment and metrics about their formation and investment events. The data supporting the study covers the period from the first quarter 2010 until the fourth quarter 2014 and is most representative of activity in the U.S. and Europe because of disparities in the use of venture capital and reporting of investment events in other parts of the world.

logistics, and maintenance share major similarities. If instead of focusing on enhanced performance, information analytic technologies are used to decrease manpower requirements, they have the potential to decrease costs. Logistics is another area that is ripe for DoD to leverage commercial development, as DoD has enormous supply operations. Analysis shows that 71 companies working on technologies related to autonomy and cost reduction received funding at some point in the last five years, totaling around \$250 million in private investment.

Not surprisingly, applications in this space are varied, from inventory management software and warehouse control tools to dispensers used in the pharmaceutical industry. This suggests that successful integration could streamline many aspects of the Department's 'tail,' decreasing manpower needs and waste and generating savings. Moving along the logistics chain, autonomous driving and convoy technologies are slowly rolling out for commercial trucking, which has the potential to decrease manpower requirements for DoD, although the usefulness of commercial-grade technology may be limited in adversarial conditions.

For employing individual robotic platforms, there is the potential for overlap in lower-capability, lower-cost, unmanned aerial vehicles and sensors. The private sector is developing inexpensive aerial systems to decrease the cost of aerial surveying and collecting crop data, in addition to low-cost sensors for automotive applications, mobile devices, and video games. If DoD develops CONOPS for lower-performance systems, there is an opportunity to leverage a large amount of private investment, as the unmanned aerial system sector has attracted approximately \$500 million alone in the past 5 years. However, this overlap largely does not extend to high-performance systems and sensors.

One area that has substantial overlap across the categories of performance, risk, and cost is human-machine interaction technology. Despite different missions and operational environments, DoD and commercial needs for human-machine interaction technology are similar because both are seeking relatively intuitive mechanisms for interaction. As such, heavy private sector investment should benefit DoD. Gesture recognition technology is an area of current investment and company formation, with at least 14 companies founded between 2010-2014 and private investment of \$130 million. Beyond private company work, a notable example of advancement in this area is Microsoft's Kinect technology. While originally designed for the Xbox video game platform, it is now being used or developed for retail environments, operating rooms, and physical therapy clinics.¹⁰ Natural language processing is another area which can enable fluid interaction between humans and machine, and this is a dynamic space in the private sector, with approximately 50 companies founded between 2010-2014 and overall investment since of almost \$850 million.

Beyond the implications for technologies, commercial and defense needs converge in similar areas for CONOPS development and TEV&V, although not for systems architectures. As described above, data analytic systems will have share similar CONOPS in terms of identifying trends and anomalous data, although differences in the degree to which the environment is adversarial will affect the degree to which commercial TEV&V is relevant. Cyber defense and high-frequency trading are inherently adversarial

¹⁰ For further information, see: <http://www.microsoft.com/en-us/kinectforwindows/meetkinect/default.aspx>

activities, so these will share similarities with DoD needs to develop robust CONOPS and TEV&V, and logistics and maintenance needs convergence tightly. CONOPS and TEV&V for systems to operate in warehouses and depots should also be markedly similar, which should further increase the attractiveness of commercial technology in these areas. If present systems are any guide, however, proprietary communication, data standards, code, and hardware will limit the ability to leverage technology and data across systems, which is likely to be a challenge for DoD in leveraging commercial technology while managing system costs.

Unique DoD Applications

The primary difference between DoD and private sector applications is DoD's need to carry out sophisticated operations in complex, adversarial environments. This is particularly true for self-contained autonomous systems, which generate capabilities through on-board systems. At this level, most physical systems for defense applications will require significantly higher performance than their commercial counterparts, as well as the ability to manage limitations on size, weight, power, and bandwidth. For instance, the private sector has few incentives to develop the high-end sensors required to enable autonomous, high-speed, high-maneuverability operations, such as is envisioned by DARPA's Fast Lightweight Autonomy program, which seeks to create an autonomous system that can fly like a bird or an insect. Likewise, while commercial applications of autonomy for marketing and sales will develop perception systems to observe customers as they shop, requirements will be less stringent than for DoD applications in a warzone.

The unstructured nature of the battlefield creates substantially greater complexity than commercial environments. For an unmanned ground vehicle, even identifying whether it can drive through a field that is partially obscured by vegetation and has mud, holes, and rocks adds a substantial measure of complexity compared to driving along paved roads. Some missions, such as undersea or in an A2/AD environment, may preclude communication with network resources, requiring more extensive on-board capabilities while also facing power limitations. Beyond these difficulties, the prospect of operating against intelligent adversaries creates potentially major challenges and is worthy of special emphasis. For example, any system with a collision avoidance algorithm provides adversaries an opportunity to influence its behavior by intentionally putting it at risk of collision. For a maritime system, that might mean that consistently putting a ship in a situation where it does not have the right of way would allow an adversary to direct it to turn in ways that would disrupt its mission or even to guide it into a port.

The Google Car project nicely illustrates the differences between DoD and commercial needs. Despite the fact that analysts often hold up the Google Car as the paragon of autonomous systems development, even in defense circles, it has limited onboard capabilities. While Google has achieved impressive accomplishments, the car cannot even identify a traffic signal by itself, instead relying on exquisite maps, which are only possible because roads are relatively static. It is also designed to act predictably because that is beneficial to safety in a relatively cooperative environment -- incentives for efficiency, reliability, low-cost, and user-friendliness will push the private sector to optimize its systems in ways that would create major vulnerabilities for DoD. As a consequence, the private sector is unlikely to address DoD's needs for high-performance or protection in complex, adversarial environments because

These performance requirements apply to the physical aspects of systems as well as to the perception and cognition algorithms. Operating in complex, adversarial environments imposes significant burdens on robots by necessitating maneuverability, speed, or hardiness. The divergence between commercial needs and DoD needs will be especially wide for sea systems, as there is some commercial development of aerial and ground robots, but much less for maritime applications. Nonetheless, the private sector's major robotics focus areas, to include aerial systems and medical robotics, will still not provide high-end capabilities or the ability to operate in unstructured environments. This is not to say that the private sector will not develop components that benefit DoD R&D though, especially considering the magnitude of investment in the robotics field – \$3.3 billion in investment from 2010-2014.

The same differences in applications and operating environments that affect technology applicability also impact considerations of architectures, CONOPS development, and TEV&V. Most companies will develop proprietary system architectures that will be, at most, interoperable with parts from the same brand. Even if there are coalitions of companies that decide to make robots and other systems interoperable, it is unlikely that DoD would seek to match such a standard, as this would increase the likelihood of interference and potentially the ease of development of counters to military systems. In terms of CONOPS, commercial applications of autonomy will typically seek to implement the most predictable, efficient CONOPS possible, which will be poorly suited for operations in adversarial environments. Similarly, for TEV&V, the defense requirement for systems to operate in highly-complex, adversarial environments will create complexities that require testing in these environments to gather high-fidelity data, differentiating defense from commercial needs.

Thus, looking across DoD applications, the private sector will be well positioned to support DoD needs in terms of cost reduction, information analysis, and cyber defense, but DoD will need focused investments to develop systems to operate in complex, adversarial environments. However, just because private sector investment will develop relevant technologies does not mean that DoD will benefit. The Office of the Secretary of Defense and the Services will need to enact concerted efforts to identify, evaluate, acquire, and tailor commercial technologies. This is a particularly significant challenge for cost applications, as most DoD researchers and policymakers focus on applications of autonomy to enhance performance and protection to the neglect of pure-cost saving applications.

Efficient Investment and Policy Approaches

Based on the preceding discussions of component technologies, architectures, and implementation factors and the opportunities for DoD applications of autonomy, the balance of the paper discusses R&D and policy needs and makes recommendations to better position DoD to reap a full measure of the benefits of autonomy in a cost-effective manner. One issue that cuts across many of the recommendations is the consideration of adversarial influence. As discussed above, adversaries will attempt to exploit autonomous systems by structuring the environment and interacting with autonomous systems, but few current efforts are conducting R&D, carrying out experimentation, or developing approaches for testing systems to enable effective operation against intelligent adversaries.

Technology

DoD has critical needs for technologies which can enable operations in complex, adversarial environments and that are well beyond commercial needs. The following discussion identifies priority areas across the perception, cognition, and action fields.

Perception requires a combination of sensors and algorithms to interpret sensor data. Because autonomy will enable fast and highly maneuverable systems and very long-duration systems which will not have access to infrastructure, DoD should focus investment in the sensor arena on high-performance approaches which can enable systems like those contemplated in DARPA's Fast Lightweight Autonomy Program and very low-power sensors that can enable long-duration, continuous operation missions, such as using underwater gliders for ISR.

Recommendation 1: Develop high-performance and low-power sensors to support short-timescale decisions and long-duration systems.

Researchers have already identified general mechanisms to train perception algorithms, relying on machine learning, but these yield relatively brittle systems in the complexity of the real world, and very little research to-date is examining perception in adversarial environments. As a result, DoD should focus its investment in perception on technologies suited to complex, changing environments with special considerations for spoofing, jamming, use of camouflage, and other adversarial behaviors. Two technical opportunities to enable more robust perception algorithms are sensor fusion and algorithms that can understand relationships. Moving from single-sensor approaches, such as relying only on a camera, to multiple sensor modalities, which might also include radar and GPS, for example, has the potential to make systems more robust by increasing the ability to differentiate between items that might appear similar or are difficult to identify when only using one modality. Another promising approach is to provide the system a mechanism for understanding relationships between parts of a whole. For example, if the system understands that a human has hands, arms, legs, a torso, and a head, is likely to wear shoes and other items of clothing, and has a typical body temperature, it is more likely to identify that there is a human standing behind a car when only the head and shoes are visible and that a cardboard cutout is not a living human.

Recommendation 2: Focus perception R&D on technologies suited to complex environments with special consideration for adversarial behaviors.

Cognition requires processing that can operate under appropriate size, weight, and power constraints and suitable artificial intelligence algorithms. In terms of computing, systems with broadband communications or which generate large amounts of power will be able to leverage cloud computing resources or intensive onboard processing, so DoD efforts should focus on developing processors for systems with limited power and communications availability. Research into neuromorphic computing is one promising avenue to achieve lower power consumption and high performance. Based on the structures of animal brains, neuromorphic approaches can offer large computational resources at a fraction of typical energy costs by mimicking brain function where pathways required for computation are only activated as needed.

Recommendation 3: Support continued research into low-power, high-performance onboard computing to support long-duration autonomous systems.

Moving to artificial intelligence, there is currently no general mechanism for creating these algorithms. Because of its wide variety of needs, DoD should invest in basic research in this area. Neuroscience is one area where DoD should examine opportunities to draw inspiration to develop more effective artificial intelligence systems, especially considering the large amount of investment in neuroscience at present. Before the basic research matures, DoD will still need to develop specialized algorithms, especially those that can operate in adversarial environments. Like perception, there is very little development work for operating in adversarial environments. In order to achieve this, DoD should fund research into approaches that render artificial intelligence less susceptible to intelligent influence. One area worthy of consideration is applied game theory, which may allow systems to effectively respond to adversary actions.

Recommendation 4: Support basic research into general methods for developing artificial intelligence algorithms, including studies examining opportunities to leverage advances in neuroscience.

Recommendation 5: Fund research into methods to develop artificial intelligence algorithms that are resistant to malicious human influence.

In the action field, there may be opportunities to leverage low-cost aerial systems when moderate performance is acceptable. DoD should also look to the private sector for data analytic and some cyber defense applications, as well as most human-machine interaction and efficiency-related technologies. For many robotics applications, however, broad investment needs for autonomy do not differ greatly from those of other high-performance systems designed for military use, where DoD requirements for ruggedness and performance typically exceed the capabilities offered by commercial platforms. In this context, DoD should pay special attention to opportunities to develop loosely coupled systems that leverage new designs, configurations, or levels of performance because they can operate unmanned or with reduced manning.

Recommendation 6: Leverage private sector R&D in low-cost aerial systems, data analysis software, cyber defense, human-machine interaction, and efficiency-related technologies.

Recommendation 7: When developing robotics for loosely coupled systems, analyze opportunities for designs that are unencumbered by the human form or frailties.

Architectures

In addition to preparing autonomous systems for a complex, adversarial environment, DoD should take steps to ensure the best return on its investment as it develops and begins to acquire new classes of autonomous systems. To do so, DoD should ensure that software and hardware components are modular so they can be reused across systems to decrease costs and enable rapid upgrades. It should also focus on designing systems that can communicate with each other to enhance capabilities and on implementing

design approaches that safeguard against system-wide vulnerabilities. This ‘portfolio’ approach will ensure the greatest return on DoD investments, while mitigating major risks.

The main obstacle to introducing modularity into the force is reliance on proprietary hardware and software that are tethered to the original manufacturer. This creates legal complications involving intellectual property and practical challenges involving design that impose high costs for platform modifications. To mitigate these challenges, DoD should require an open architecture for systems, which is a series of protocols that define how different aspects of the system interface and communicate with each other from a hardware and software point of view. Adopting an open architecture would mean that any developer could create new modules – perception, cognition, or action – and integrate them relatively quickly into existing systems, including with parts from other designers – to develop new or upgraded systems. This would decrease upgrade and development costs and timelines.

To support this, DoD should fund research to develop a forward looking architecture that is able to support a wide range of component types and data flows with enough flexibility to support components and types of information that designers have not yet envisioned. While there are substantial benefits from this approach, it is important to acknowledge potential drawbacks. Because an open architecture will specify interface characteristics, systems will be less tightly integrated than if they had no restrictions. Nonetheless, the benefits from an open architecture should greatly outweigh this cost.

Recommendation 8: Fund research to develop a forward-looking open architecture for autonomous systems and mandate cross-Service compliance with these standards.

Moving beyond the individual-system level, a number of the proposed benefits of battlefield autonomous systems depend on coordination between systems and between systems and humans, and the greater the number that can potentially work together, the greater the opportunities to generate complex and emergent behaviors. As such, it will be highly valuable to enable information flows between machines and between machines and humans. This does not mean that all systems must use the same operating system or use identical sensor modalities, but their output to the user or to other autonomous systems must be intelligible, regardless of the originating platform. While programs such as the Joint Architecture for Unmanned Systems offer models for establishing interoperability standards, these standards are not typically embraced across domains and Services. DoD should further develop and enforce interoperability standards.

Recommendation 9: Fund research to develop forward-looking interoperability standards for autonomous systems and mandate cross-Service compliance with these standards.

Finally, behaviors of autonomous systems may make them vulnerable to adversarial influence, meaning that understanding these mechanisms is an important part of competing with adversary systems and developing DoD systems. DoD should also consider the vulnerabilities of deploying a large number of similar systems, thereby creating the weaknesses of monoculture. For systems where humans are ultimately responsible for decision making, variation between individuals means that even systems of the same type tend to react differently to stimuli and that it is more difficult to influence a whole class of systems effectively. However, if there is widespread use of similar or identical perception and cognition

in autonomous systems, this raises the potential for one or a small number of vulnerabilities to endanger a large proportion of the force. To combat this, DoD should introduce heterogeneity to increase system-wide resilience. Systems may introduce heterogeneity through designs featuring some degree of random behavior or by variations in the structure and code of systems.

Systems of systems comprised of heterogeneous classes will tend to be more resilient, but this will also force designers to choose intentionally not to optimize certain components to maintain variability. Thus, this tradeoff must be considered, but in general, the dangers of deploying a monoculture are greater than the drawbacks from heterogeneity. Benefits of heterogeneity might also be one of the future value propositions for maintaining higher-levels of human influence in the system. Despite the flaws and limits of human cognition, humans are fundamentally idiosyncratic, making groups relatively adaptable and robust. DoD should study the optimal makeup of human-systems teams to identify complimentary applications of human and autonomous perception and cognition.

Recommendation 10: Intentionally design heterogeneity into classes of systems to mitigate the likelihood of system-wide vulnerabilities, and study the optimal role of humans to increase resilience.

Employment

Implementing autonomous systems into the broader military force will require experimentation and testing to understand the most effective ways to use them and to ensure their appropriate operation, but there has been relatively little to date. The difficulty of predicting the behavior of advanced systems in a complex environment means that focused CONOPS development and TEV&V of autonomous systems will be critical to ensuring autonomy's transition from technical possibility to effective operational capability.

To best take advantage of the unique interactions and capabilities offered by autonomous systems, DoD will need to develop new CONOPS. New capabilities will enable missions over extremely long durations, with forces sometimes operating without communication or logistical support. Likewise, the potential for autonomous systems to collect data for long periods of time to inform planning and to act on a sub-second timescale poses unique opportunities and challenges for tailored approaches to missions and command and control. In addition, research in games, such as chess, go, and military simulations, demonstrates that autonomous systems sometimes analyze situations very differently than humans, which increases the likelihood of arriving at conclusions that are unexpected from a human perspective. With little research to date, the implications of new capabilities and approaches are still unclear, and DoD should fund intellectual experimentation, modeling and simulation, and field experimentation to explore new CONOPS, as new approaches to applying technologies are just as important as the technologies themselves.

Experimentation can provide the additional benefit of informing system design. For example, in the period between World War I and World War II, experimentation with carrier aviation demonstrated that launching large waves of aircraft was important, and this influenced carrier design in ways that were crucial to U.S. naval performance in World War II. One particular area worth considering for autonomous systems is the potential value of larger numbers of lower-cost systems. This is particularly applicable to

the air domain, where commercial investment in aerial systems is likely to deliver moderate capabilities at relatively low cost.

Developing effective CONOPS is more than just capitalizing on new capabilities. It also requires recognition of new limitations. The interaction of CONOPS with technological factors will be critical and require substantial experimentation. For example, one factor that is not immediately obvious outside of the robotics field is that small systems tend to have substantially shorter range because they cannot carry as much fuel or energy. Ethical factors will also have an important influence on CONOPS. In November 2012, DoD released Directive 3000.09, *Autonomy in Weapons Systems*, which lays out requirements for the development and use of autonomous systems in response to challenges to our current ethical-legal approaches when systems make increasingly complex decisions based on increasingly complex perception. DoD must increase its study of the technical-operational and ethical-legal implications of autonomous systems to inform both what is possible and what is appropriate. Considering the critical role that CONOPS play and the wide range of influences that will affect the optimal use of autonomous systems, DoD should supplement internal work with prize competitions to engage a broad set of communities about how to most effectively employ autonomous systems in different operational environments.

Recommendation 11: Fund experimentation to develop new CONOPS for autonomous systems, while paying special attention to opportunities to leverage artificial intelligence’s propensity towards unconventional approaches to problem solving, opportunities to capitalize on larger numbers of lower-cost systems, and ethical-legal considerations.

Recommendation 12: Develop prize competitions for new CONOPS to gather insights from a broad set of communities.

Aside from developing CONOPS to enhance capabilities, the Services should give greater focus to implementing CONOPS and appropriate technologies specifically designed to reduce costs. Reducing costs from logistics, maintenance, and information analysis and increasing system readiness can free up resources to invest in additional capabilities, but at present, the Services are highly focused on opportunities to implement autonomy within the scope of performance enhancement and warfighter protection. Cost-focused applications face much lower technical barriers, as the environments tend to be more structured and predictable, so they are also likely to pay relatively quick dividends, especially because the private sector is developing technology DoD can leverage.

Recommendation 13: Take advantage of opportunities to leverage commercial R&D in autonomy to reduce costs, particularly in logistics, maintenance, and information analysis.

Given their potential sophistication, autonomous systems will also require the development of new metrics, standards, and methodologies for TEV&V. A range of DoD organizations have already identified TEV&V as a key challenge for autonomous systems, including the Autonomy Community of Interest and DoD’s Test Resource Management Center, which is set to begin an in-depth study of this area.

Nonetheless, there is relatively little ongoing research and development in this area considering the magnitude of the challenge. The traditional design of experiments methodology, which systematically documents cause and effect relationships, will not be suitable for the complexity of systems with high-end perception and cognition operating in complex environments – especially learning systems that by definition change over time – because it will not be possible to observe all possible permutations of inputs and behaviors. These factors will make it difficult to predict exactly how these systems will work, and they may make it even more difficult to predict the various ways they will fail.

Moving forward, one framework DoD should investigate is a licensing approach. Using humans as an example, it is impossible to exhaustively test all of the failure states of a person, but by putting human operators through rigorous trials related to mission needs, we build trust in their capabilities and eventually deem them worthy of a license, such as certification in a particular aircraft type. For autonomous systems, the goal should not be to establish certainty of the system's behavior in all situations, but to build trust that it will act reasonably, will have certain limits on inappropriate behavior, and will fail relatively predictably. Financial firms are one area to investigate for insights into licensing systems that will influence the environment, as high-frequency trading systems necessarily influence the market once connected, but they must be tested before doing so. Of course, these systems have also been involved in major market failures, such as the 'Flash Crash' of 2010, but these types of events are likely to provide valuable insights as well.

Recommendation 14: Fund research and development in metrics, standards, and methodologies for TEV&V, including an examination of licensing as an approach to bring complex autonomous systems into the force.

Recommendation 15: Engage with the finance industry to examine how corporations conduct TEV&V for trading systems before connecting them to exchanges.

In order to put autonomous systems through rigorous trials and enable effective TEV&V, it will be critical to test autonomous systems in realistic environments, as laboratory settings will not replicate the complexities of the real world. Using the DARPA Robotics Challenge as an example, teams tested their robots extensively in large indoor settings, adding fidelity by, for example, bringing in lights to simulate glare from the sun. However, because the teams did not test their robots outside, they failed to account for wind, which kept closing doors after the robots opened them during the real competition. No laboratory setting can fully replicate the chaos of the real world, especially a battlefield environment. As such, DoD should ensure that we challenge systems during development and TEV&V. In addition to assessing behavior in complex environments, testing should include requirements to identify how difficult it is to predict the behavior of systems based on their code in case systems are captured or code is exfiltrated in cyber attacks, as well as how difficult it is to predict behavior based on observations of systems in the environment, of which adversaries will likely try to take advantage.

The issue of testing is especially relevant if humans are meant to interact with these systems. As described above, autonomous systems will influence human behavior, so testing must ensure that neither the human nor the system causes failures in the other. For instance, the issue of trust is often a source of

failure. A lack of trust may cause humans to disregard the machine, and too much trust may cause them to overlook errors. In order to test these interactions, DoD should require that TEV&V include human-machine teams, operating first in laboratory environments for safety and then in complex environments for realism.

Recommendation 16: Fund extensive testing in realistic environments, which include humans, to ensure that systems operate effectively and are relatively resistant to adversarial behavior.

Given these expansive testing needs and the unique challenges of TEV&V for autonomy, DoD should establish a dedicated air-land-sea range to enable TEV&V in complex, realistic environments. The more complex autonomous systems become and the more systems interact on the battlefield, the more we will need to test them in realistic settings to understand their behavior. As such, ranges will be a key resource for systems development and TEV&V. While expensive, this would ensure availability of range time and the opportunity to investigate interaction between systems developed by the different Services. It would also have extremely valuable spillover benefits for experimentation in support of CONOPS development and system design.

Recommendation 17: Establish a dedicated air-land-sea range for TEV&V of and experimentation with autonomous systems.

In addition to field testing, the development of more extensive modeling and simulation will be key to accelerating testing and conserving resources. Because it is fast, but provides lower fidelity, modeling and simulation is most useful in an iterative process with field experimentation. However, modeling and simulation capabilities for autonomous systems are currently limited, especially concerning the behaviors of multiple systems in complex environments. As a result, DoD should invest in improving these methods, with a special focus on identifying the areas where it provides high fidelity and where it loses predictive power. These development efforts should also be used to benefit CONOPS development and system design.

Recommendation 18: Invest in modeling and simulation to improve the speed and efficiency of TEV&V.

Cross Cutting Areas

Due to the broad range of areas that influence autonomous systems development and testing, research from many different fields across academia and industry will be applicable. This ranges from technological development, where biological sciences can provide insights to guide work in the areas of perception and cognition, to implementation issues, where psychology can inform work on trust and interaction and the videogame industry can offer resources for modeling and simulation. While many of these fields fall outside of traditional autonomy research, their contributions are likely to prove valuable in accelerating future developments. DoD should continue and expand engagement with non-traditional partners through grants, awards, and challenge prizes to leverage their work in overcoming the hurdles to fielding highly effective autonomous systems.

Recommendation 19: Continue and strengthen engagement with fields not traditionally involved in autonomy research to find new approaches to and solutions for DoD challenges and opportunities in autonomy.

Conclusion

Autonomy can provide DoD tremendous value by enhancing the performance of military systems, decreasing risk to U.S. warfighters, and generating cost savings. DoD R&D efforts should focus on developing systems to operate in complex, adversarial environments and leveraging commercial technology for information analysis, cyber defense, and cost-saving applications. To support these technology development efforts, DoD should develop policy and conduct research and experimentation to support open architectures, interoperability, resilience, CONOPS development, and TEV&V to ensure that U.S. systems maximize capability and flexibility, while minimizing risk and cost. If successful, developments in these areas will enable DoD to overcome challenges posed by the security environment, paving the way for continued U.S. military superiority.

Appendix A: Commercial Investment Insights: Autonomy

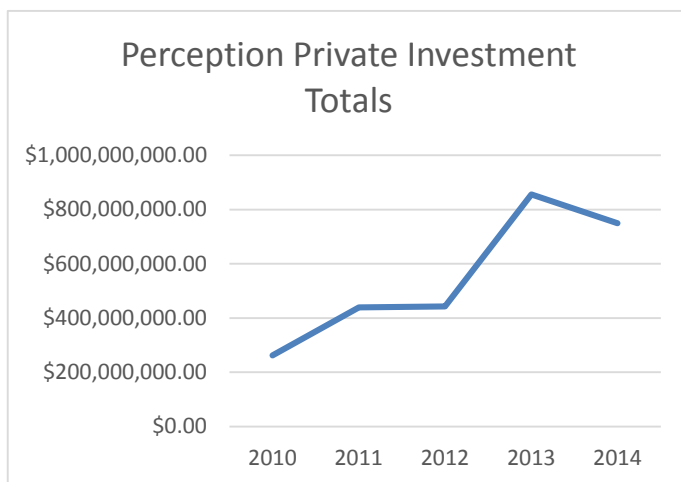
The Office of Technical Intelligence carried out a study of the private company landscape in fields relevant to autonomy in order to gain insight into where DoD can capitalize on private sector technology development. Because private companies tend to be at an earlier stage in their development, trends in formation, foci, and ability to attract investment are signals about future trends in the commercial field and provide insights to inform DoD R&D strategy.

For this study, OTI conducted interviews of subject matter experts in the area of autonomy and leveraged a commercially available private sector analytic tool, from Quid Inc., to analyze the structure and relative size of fields relevant to autonomy. We parsed this space into three sectors that represent key technical areas for autonomy: perception, cognition, and robotics.

- Perception – companies in this network represent the sensing aspect of autonomy, which allows a system to perceive and understand its operational environment. Perception is typically created from a combination of sensors, which provide data inputs, and algorithms to turn that information into contextual understanding.
- Cognition – companies grouped into this network represent the thinking aspect of autonomy, ranging from understanding how different factors in an environment interact with each other to making decisions without the help of a human operator. Cognition results from a combination of computation and artificial intelligence algorithms.
- Robotics – companies in this network represent the physical aspect of autonomy. Although an autonomous system need not include a robotic component, robotics enable a system to act or move within its operational environment.

Findings

A critical finding of this analysis is that there has been and continues to be significant investment in all technology areas relevant to autonomy – both in numbers of companies and investment dollars. However,



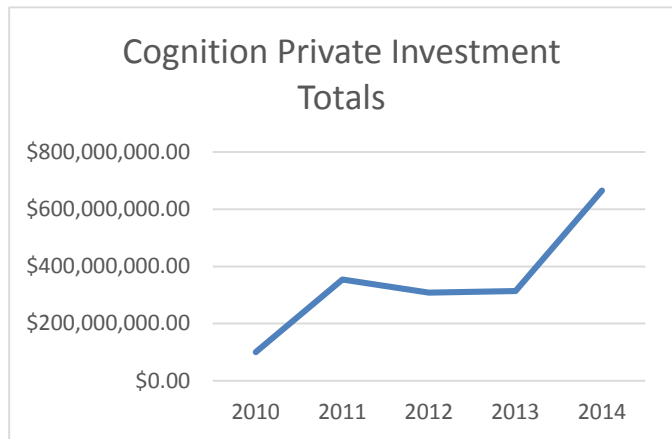
scoping this space proves difficult is heavily influenced by the end technology considered autonomy and the capabilities DoD aims to develop. Regardless, investment across each of these areas has increased since 2010.

Perception: The perception industry has experienced significant growth since 2010, as well as many notable commercial successes, such as Microsoft’s Kinect technology. Clearly defined areas where the private sector is placing money include gesture recognition and natural language

processing technologies. Since the start of 2010, our analysis identified approximately \$850 million in venture capital funding flowing to companies developing natural language processing technology.

Companies developing gesture recognition technology garnered approximately \$130 million in private investment during this same period. Perhaps most importantly, as the Autonomy Technical Assessment outlines in more detail, DoD needs in these two areas overlap significantly with those in the commercial space, creating a prime opportunity to leverage commercial advances towards DoD applications.


Cognition: The cognition space has experienced a similar trajectory to the perception space – venture capital funding levels have grown significantly since 2010, and new company formation remains high. DoD requirements for cognition vary significantly - from logistics planning to threat identification – but are not always convergent with commercial applications. However, a space where DoD and private sector needs align is data analysis – an expensive, yet necessary part of both arenas. For example, commercial efforts towards targeted advertising and process optimization and efficiency have parallel applications within DoD and the private sector.



Companies working in data analysis attracted almost \$900 million in investment since 2010 – a sizeable portion of the larger cognition area.

Interviews of subject matter experts help place these figures in context, however; for example, DARPA’s Gill Pratt noted that the cognition space is an area that is still in need of DoD basic research investment because, despite the large investment, R&D has yet to yield a generalizable “cognition capability” and it is far from achieving the broad capabilities that DoD systems will require in the future.

Robotics: Private sector investment in the robotics sector contains mixed levels of overlap with DoD concerns. Basic technologies that enhance maneuverability, strength, power consumption, and other factors will share applications in both the public and private sectors. Thus, private sector investment



should develop technology and drive down costs for DoD, especially considering the magnitude of investment and wide span of commercial applications for robotics. However, many DoD robotic systems will also require substantially greater performance.

One particular area of private sector work is in aerial systems, which attracted roughly \$500 million in private investment. In particular, many of these companies are working on smaller or micro UAVs for applications in

agriculture monitoring, disaster response, and search and rescue – which will provide DoD opportunities to purchase relatively low-cost systems if it can develop effective CONOPS to leverage them.

Additionally, in medical care, private sector developments are likely to have substantial overlap. Over 50 companies working in the medical robotics sector have attracted around \$700 million in private investment since 2010. Technology in this sector centers on instruments providing robotic assisted control during surgical procedures, but also touches areas such as powered prosthetics and hospital logistical challenges, such as dispensing medicines and material handling. Little of this investment is likely to produce the high-performance platforms required for more demanding missions, but there are many lower-complexity areas where DoD may find these technologies to be of great benefit.

Challenges: Analyzing the state of autonomy in the private sector presents a substantial challenge. Previous OTI Technical Assessments centered on a specific technology area that has defined boundaries, an established definition, and a distinct commercial sector. However, autonomy is not a specific technology; rather, autonomy represents a capability derived from the combination of numerous technology areas. Therefore, it is not feasible to determine the exact size, shape, or subsectors of the “autonomy industry,” as this is highly dependent on the technology one considers relevant and applications or capabilities of interest. Nonetheless, this analysis provides insights in to technology areas where private sector investment is likely to deliver valuable technologies to DoD and informs the R&D recommendations in the Autonomy Technical Assessment.

Conclusions: The private sector has and will continue to invest in technology areas that are crucial to the development of autonomous systems. In some cases there will be substantial overlap between DoD needs and private sector requirements. This overlap creates opportunities for DoD to leverage the private sector for its benefit. However, some DoD needs for autonomy are, and will remain, strictly DoD related challenges because of a lack of commercial applications. Thus, DoD should aim to leverage the private sector to solve lower complexity problems where significant near-term savings and advancements can be realized and conserve R&D investment for other, defense-unique areas.