

November 7, 2016



A Roadmap for US Robotics

From Internet to Robotics

2016 Edition

Organized By

University of California San Diego
Carnegie Mellon University
Clemson University
Cornell University
Georgia Institute of Technology
Northeastern University
Northwestern University
Oregon State University
SRI Inc.
Texas A&M University
The University of Utah
University of California Berkeley
University of Nevada - Reno
University of Southern California
University of Tennessee Knoxville
University of Washington
University of Wisconsin
Vanderbilt University
Yale University

Sponsored by:

National Science Foundation
University of California San Diego
Oregon State University
Georgia Institute of Technology

Table of contents

1. EXECUTIVE SUMMARY.....	5
1.1 INTRODUCTION.....	5
1.2 MAIN FINDINGS.....	6
1.3 THE ROADMAP DOCUMENT.....	7
1.4 FURTHER INFORMATION IS AVAILABLE FROM:.....	7
2. TRANSFORMING MANUFACTURING AND SUPPLY CHAIN.....	9
2.0 EXECUTIVE SUMMARY.....	9
2.1. INTRODUCTION.....	10
2.2. STRATEGIC IMPORTANCE OF ROBOTICS IN MANUFACTURING.....	10
2.2.1 <i>Economic Impetus</i>	10
2.2.2 <i>Growth Areas</i>	11
2.2.3 <i>“Consumerization” of Robotics</i>	12
2.2.4 <i>A Vision for Manufacturing</i>	12
2.3. RESEARCH ROADMAP.....	13
2.3.1 <i>The Process</i>	13
2.3.2 <i>Robotics and Manufacturing Vignettes</i>	14
2.3.3 <i>Critical Capabilities for Manufacturing</i>	15
2.4. REFERENCES.....	19
2.5. CONTRIBUTORS.....	20
3. NEXT GENERATION CONSUMER AND PROFESSIONAL SERVICES.....	23
3.1. INTRODUCTION.....	23
3.2. PRINCIPAL MARKETS AND DRIVERS.....	23
3.2.1 <i>Professional Service Robots</i>	24
3.2.2 <i>Consumer Service Robots</i>	28
3.2.3 <i>Transportation</i>	28
3.3 NEAR-TERM OPPORTUNITIES AND FACTORS AFFECTING COMMERCIALIZATION.....	29
3.4 REFERENCES.....	31
3.5 CONTRIBUTORS.....	31
4. HEALTH, INDEPENDENCE AND QUALITY OF LIFE.....	33
4.1. INTRODUCTION.....	33
4.1.1. <i>Definition of the Field/Domain</i>	33
4.1.2. <i>Societal Drivers</i>	34
4.2. AGING AND QUALITY OF LIFE IMPROVEMENT.....	36
4.2.1. <i>Motivating Scenario</i>	36
4.2.2. <i>State-of-the-art</i>	36
4.2.3. <i>Key Challenges and Unmet needs</i>	37
4.2.4. <i>Vision: 5-10-15 years</i>	38
4.3. SURGICAL AND INTERVENTIONAL ROBOTICS.....	38
4.3.1. <i>Motivating Scenario</i>	38
4.3.2. <i>State-of-the-art</i>	39
4.3.3. <i>Vision: 5-10-15 years</i>	39
4.4. REHABILITATION.....	40
4.4.1 <i>Motivating Scenario</i>	40
4.4.1. <i>Robotic Replacement of Diminished/Lost Function</i>	40
4.4.2. <i>Robot-Assisted Recovery and Rehabilitation</i>	41
4.4.3. <i>Behavioral Therapy</i>	42
4.4.4. <i>Key Challenges and Unmet needs</i>	43
4.5. CLINICAL WORKFORCE SUPPORT.....	43
4.5.1. <i>Motivating Scenario</i>	43

4.5.2. <i>State-of-the-art</i>	44
4.5.3. <i>Key Challenges and Unmet needs</i>	44
4.5.4. <i>Vision: 5-10-15 years</i>	45
4.6. DEPLOYMENT ISSUES.....	46
4.7. CONTRIBUTORS.....	47
5. ENHANCING PUBLIC SAFETY.....	49
5.1 STRATEGIC IMPORTANCE OF ROBOTICS IN ENHANCING PUBLIC SAFETY.....	49
5.1.1 <i>The Future Landscape</i>	49
5.1.2 <i>The Role of Unmanned Systems in Enhancing Public Safety</i>	50
5.1.3 <i>Motivating Vignettes</i>	55
5.2 SCIENTIFIC AND TECHNICAL CHALLENGES.....	58
5.2.2 <i>Force Application</i>	59
5.2.3 <i>Protection</i>	60
5.2.4 <i>Logistics</i>	60
5.3 TECHNOLOGY NEEDS AND GOALS.....	62
5.4 CONTRIBUTORS.....	66
6. EARTH AND BEYOND.....	68
6.1. OVERVIEW.....	68
6.1.1. <i>Robotic Handling of High-Consequence Materials</i>	68
6.1.2. <i>Robotic Reconnaissance for Human Exploration</i>	70
6.1.3. <i>Planetary Cave Exploration</i>	71
6.2. GAPS AND IMPACTS.....	71
6.2.1. <i>Agriculture</i>	71
6.2.2. <i>Environmental Monitoring</i>	72
6.2.3. <i>Space Robotics</i>	73
6.3. CONTRIBUTORS.....	77
7. RESEARCH ROADMAP.....	79
7.1. MECHANISMS AND ACTUATORS.....	79
7.2. MOBILITY AND MANIPULATION.....	80
7.2.1. <i>Mobility</i>	80
7.2.2. <i>Manipulation</i>	81
7.3. PERCEPTION.....	82
7.4. FORMAL METHODS.....	83
7.4.1. <i>Synthesis and verification for closed loop systems; incorporating uncertainty and dynamic environments</i>	84
7.4.2. <i>Safe behavior degradation</i>	84
7.4.3. <i>Formal methods and learning</i>	84
7.4.4. <i>Formal methods for human-robot interaction and collaboration</i>	84
7.5. LEARNING AND ADAPTATION.....	85
7.5.1 <i>Learning from Demonstration</i>	85
7.5.2 <i>Reinforcement Learning and Deep Learning</i>	85
7.6. CONTROL AND PLANNING.....	86
7.6.1. <i>Task and Motion Planning Under Uncertainty</i>	86
7.6.2. <i>From Specification to Deployment</i>	86
7.6.3. <i>Control and Planning in Constrained Environments</i>	87
7.6.4. <i>Manipulation</i>	87
7.6.5. <i>Dynamic environments</i>	87
7.6.6. <i>Coordination Among Multiple Agents</i>	87
7.7. HUMAN ROBOT INTERACTION.....	89
7.7.1. <i>Interface Design</i>	89
7.7.2. <i>Perceiving, modeling, and adapting to humans</i>	89

7.7.3. Sociability.....	89
7.7.4. Collaborative Systems.....	90
7.7.5. Robot-mediated communication.....	90
7.7.6. Shared Autonomy.....	90
7.7.7. Long-term interaction.....	91
7.7.8. Safety.....	91
7.7.9. Research Roadmap.....	91
7.8 MULTI-AGENT ROBOTICS.....	93
7.8.1. Distributed Control and Decision Making.....	93
7.8.2. Mixed Decentralized/Centralized Information Exchange Mechanisms.....	94
7.8.3. Human-Swarm Interactions.....	94
7.8.4. Heterogeneous Networks.....	94
7.8.5. Communication and Sensing in Multi-Robot Systems.....	95
7.8.6. Roadmap for multi-agent robotics.....	95
8. WORKFORCE DEVELOPMENT.....	97
8.1. INTRODUCTION.....	97
8.2. STRATEGIC FINDINGS.....	97
8.3. NEAR TERM OPPORTUNITIES AND FACTORS AFFECTING IMMEDIATE DEPLOYMENTS.....	98
8.4. CONTRIBUTORS.....	99
9. SHARED INFRASTRUCTURE IN ROBOTICS.....	101
9.1. FLEXIBLE RESEARCH PLATFORMS.....	101
9.2. COMMUNITY CONSENSUS VALIDATION BENCHMARK FRAMEWORKS.....	101
9.3. REFERENCE OPEN-ACCESS TESTBEDS.....	102
9.4. CONTRIBUTORS.....	103
10 LEGAL, ETHICAL, AND ECONOMIC CONTEXT.....	105
10.1. SAFETY.....	105
10.2. LIABILITY.....	106
10.3. IMPACT ON LABOR.....	107
10.4. SOCIAL INTERACTION.....	108
10.5. PRIVACY AND SECURITY.....	108
10.6. RECOMMENDATIONS.....	109
10.7. CONTRIBUTORS.....	109

1. Executive Summary

Robots for economic growth, improved quality of life and empowerment of people

1.1 Introduction

Recently the robotics industry celebrated its 50-year anniversary. We have used robots for more than five decades to empower people to do things that are typically dirty, dull and/or dangerous. The industry has progressed tremendously over the period from basic mechanical assist systems to fully autonomous cars, environmental monitoring and exploration of outer space. We have seen tremendous adoption of IT technology in our daily lives for a diverse set of support tasks. Through use of robots we will see a new revolution, as we not only will have IT support from tablets, phones, computers but also systems that can physically interact with the world and assist with basic daily tasks, work, and leisure activities.

The “old” robot systems were largely mechanical support systems. Through the gradual availability of inexpensive computing, user interfaces and sensors it is possible to build robot systems that were difficult to imagine before. The confluence of technologies is enabling a revolution in use and adoption of robot technologies for all aspects of daily life.

Nine years ago, the process to formulate an integrated roadmap was initiated at the Robotics Science and Systems (RSS) conference in Atlanta. Through support from the Computing Community Consortium (CCC) a roadmap was produced by a group of 120 people from industry and academia. The roadmap was presented to the congressional caucus and government agencies by May 2009. This in turn resulted in the creation of the National Robotics Initiative (NRI), which has been a joint effort between NSF, NASA, USDA and NIH. The NRI was launched 2011 and recently had its five-year anniversary. By 2013 a revision of the roadmap was generated with support from NSF and CCC.

Over the last few years we have seen tremendous progress on robot technology across manufacturing, healthcare applications autonomous cars and unmanned aerial vehicles, but also core technologies such as camera systems, communication systems, displays and basic computing. All this combined motivates an update of the roadmap. With the support of NSF two workshops took place 22-23 August in Portland, OR and 21-22 September in Atlanta, GA. In total the workshops involved 50 people mainly from academia and research institutes. The 2013 roadmap was reviewed and progress was assessed as a basis for formulation of updates to the roadmap.

The present document is a summary of the main societal opportunities identified, the associated challenges to deliver desired solutions and a presentation of efforts to be undertaken to ensure that US will continue to be a leader in robotics both in terms of research innovation, adoption of the latest technology and adoption of appropriate policy frameworks that ensure that the technology is utilized in a responsible fashion.

1.2 Main findings

Over the last 5-6 years we have seen introduction of 600,000 new jobs in manufacturing. During the same period, we have seen significant growth in adoption of robot systems in industry. The last 3 years have seen quarter by quarter records in sales of robot systems for use in manufacturing. The USA has never utilized as many robots as they do today.

A major new application domain has been in the adoption of collaborative robots that can operate side-by-side with humans. New standards (ISO 10218 and ISO/TS 15066) for utilization of collaborative systems have provided a framework for how to design systems that easily can be adopted for a diverse set of tasks in a cost-efficient manner.

Gradually sensors and computing power have become cheap enough to be easily adopted for robot applications, which is resulting in a revolution in control and flexibility of systems. However, the methods adopted for design, implementation and deployment are still relatively simple. There is a need for a continued effort to utilize modern methods in control and use of a richer set of sensors to deliver more flexible robot systems that have user interfaces that allow utilization with no or minimum training. There is a need for major new research on human robot interaction methods.

A major limitation in the adoption of robot manipulation systems is lack of access to flexible gripping mechanisms that allow not only pick up but also dexterous manipulation of everyday objects. There is a need for new research on materials, integrated sensors and planning / control methods to allow us to get closer to the dexterity of a young child.

Not only manufacturing but also logistics is seeing major growth. E-commerce is seeing annual growth rates in excess of 40% with new methods such as Amazon Express, Uber Food, ... these new commerce models all drive new adoption of technology. Most recently we have seen UPS experiment with use of Unmanned Vehicles for last mile package delivery. For handling of the millions of different everyday objects there is a need of have robust manipulation and grasping technologies but also flexible delivery mechanisms using mobility platforms that may drive as fast as 30 mph inside warehouses. For these applications there is a need for new R&D in multi-robot coordination, robust computer vision for recognition and modeling and system level optimization.

Other professional services such as cleaning in offices and shops is slowly picking up. The layout of stores is still very complex and difficult to handle for robots. Basic navigation methods are in place but it is a major challenge to build systems that have robust long-term autonomy with no or minimal human intervention. Most of these professional systems still have poor interfaces for use by non-expert operators.

For the home market the big sales item has been vacuum and floor cleaners. Only now are we starting to see the introduction of home companion robots. This includes basic tasks such as delivery services for people with reduced mobility to educational support for children. A major wave of companion robots are about to enter the market. Almost all these systems have a rather limited set of tasks they can perform. If we are to provide adequate support for children to get true education support or for elderly people to live independently in their home there is a need for a leap in performance in terms of situational awareness, robustness and types of services offered.

A new generation of autonomous systems are also emerging for driving, flying, underwater and space usage. For autonomous driving it is important to recognize that human drivers have a performance of 100 million miles driven between fatal accidents. It is far from trivial to design autonomous systems that have a similar performance. For aerial systems the integration into civilian airspace is far from trivial but it does offer a large number of opportunities to optimize airfreight, environmental monitoring, etc. For space exploration it is within reach to land on asteroids as they pass by earth or for sample retrieval from far away planets. For many of these tasks the core challenge is the flexible integration with human operators and collaborators.

The emergence of new industrial standards as for example seen with Industry 4.0 and the Industrial Internet facilitates access to cheap and pervasive communication mechanisms that allow for new architectures for distributed computing and intelligent systems. The Internet of Things movement will facilitate the introduction of increased intelligence and sensing into most robot systems and we will see a significant improvement in user experience. The design of these complex systems to be robust, scalable, and interoperable is far from trivial and there is a new for new methods for systems design and implementation from macroscopic to basic behavior.

As we see new systems introduced into our daily lives for domestic and professional use it is essential that we also consider the training of the workforce to ensure efficient utilization of these new technologies. The workforce training has to happen at all levels from K-12 over trade schools to our colleges. Such training cannot only be education at the college level. The training is not only for young people but must include the broader society. It is fundamental that these new technologies must be available to everyone.

Finally, there is a need to consider how we ensure that adequate policy frameworks are in place to allow US to be at the forefront of the design and deployment of these new technologies but it never be at the risk of safety for people in their homes and as part of their daily lives.

1.3 The roadmap document.

The roadmap document contains sections specific to different use-cases for robot technology across: transformation of manufacturing, next generation consumer and professional services, healthcare and well-being, ensuring public safety and exploring earth and beyond. Each of these areas are analyzed in detail in separate sections. Subsequently, a section provides a unified research roadmap across topical areas. Sections are devoted to workforce development and legal, ethical and economic context of utilization of these technologies. Finally a section discusses the value of access to major shared infrastructure to facilitate empirical research in robotics.

1.4 Further Information is available from:

Prof. Henrik I. Christensen
Director of Contextual Robotics
University of California San Diego
9500 Gilman Drive 0436
La Jolla, CA 92093-0436
hichristensen@ucsd.edu

2. Transforming Manufacturing and Supply Chain

2.0 Executive Summary

Restructuring of U.S. manufacturing is essential to the future of economic growth, the creation of new jobs and ensuring competitiveness. This in turn requires investment in basic research, development of new technologies, and integration of the results into manufacturing systems. Federal Investments in research in manufacturing can revitalize American manufacturing. Investing a small portion of our national resources into a science of cost-effective, resource-efficient manufacturing would benefit American consumers and support millions of workers in this vital sector of the U.S. economy. It would allow our economy to flourish even as the ratio of workers to pensioners continuously decreases. Such a research and development program would also benefit the health care, agriculture, and transportation industries, and strengthen our national resources in defense, energy, and security. The resulting flurry of research activity would greatly improve the quality of "Made in the U.S.A." and invigorate productivity of U.S. manufacturing for the next fifty years. This strategy has already been articulated in the administration's "Advanced Manufacturing Partnership" (AMP) and in the partial implementation of a "National Network for Manufacturing Innovation" (NNMI)

Robotics is a key transformative technology that can revolutionize manufacturing. American workers no longer aspire to low-level factory jobs and the cost of U.S. workers has been slowly rising due to insurance and healthcare costs. Even when workers are affordable, the next generation of miniaturized, complex products with short life-cycles requires assembly adaptability, precision, and reliability beyond the skills of human workers. Improved robotics and automation in manufacturing will: a) retain intellectual property and wealth that otherwise might go off-shore; b) save companies by making them more competitive; c) provide jobs for developing, producing, maintaining and training robots; d) allow factories to employ human-robot teams that leverage each others' skills and strengths (e.g., human intelligence and dexterity with robot precision, strength, and repeatability), e) improve working conditions and reduce expensive medical problems; and f) reduce manufacturing lead time for finished goods, allowing systems to be more responsive to changes in retail demand. Indeed effective use of robotics will *increase* U.S. jobs, improve the quality of these jobs, and enhance our global competitiveness. The advantages has already been recognized by companies such as NCR, Cisco, Apple, Lenovo and Tesla in their setup of new factories in the USA. Through utilization of robotics and automation the expectation is that such in-shoring will continue to flourish.

We summarize the strategic importance of robotics and automation technologies to manufacturing industries in the U.S. economy, describes applications where robotics and automation technologies will dramatically increase productivity, and outlines a visionary research and development roadmap with key research areas for immediate investment to reach these goals.

2.1. Introduction

This section summarizes the activities and results of a number of workshops on manufacturing and automation robotics that was carried out under the auspices of the Robotics-VO an organization sponsored by the National Science Foundation. An effort was organized to update the National Robotics Roadmap: From Internet to Robotics [NRR-13]. The objective was an update of the roadmap considering progress over the last 4-5 years. The research agenda proposed in this report will help strengthen the manufacturing sector of the U.S. economy, a well-trained, technologically-astute workforce, the creation of new jobs, and broad-based prosperity for Americans.

The terms “robotics” and “automation” have a precise technical meaning. According to the Robotics and Automation Society of the Institute of Electronics and Electrical Engineers, “*Robotics* focuses on systems incorporating sensors and actuators that operate autonomously or semi-autonomously in cooperation with humans. Robotics research emphasizes intelligence and adaptability to cope with unstructured environments. *Automation* research emphasizes efficiency, productivity, quality, and reliability, focusing on systems that operate autonomously, often in structured environments over extended periods, and on the explicit structuring of such environments.”

Our goal is two-fold: First, to determine the strategic importance of robotics and automation technologies in manufacturing industries in the U.S. economy (Section 2.2); second, to determine applications where robotics and automation technologies could increase productivity (Section 2.3); The presentations in this section feeds into the research roadmap presented in section 7.

2.2. Strategic Importance of Robotics in Manufacturing

2.2.1 Economic Impetus

The basis for the economic growth in the last century came from industrialization, the core of which was manufacturing. The manufacturing sector represents 12% of the U.S. GDP and about 9% of the total employment [WB-16, BEA-15]. Fully 70% of the net export of the U.S. is related to manufacturing [State-09], so the sector represents an area of extreme importance to the general economic health of the country. Within manufacturing, robotics represents a \$8B-industry in the U.S. that is growing steadily at 9% per year. This core robotics industry is supported by manufacturing industry that provides the instrumentation, auxiliary automation equipment, and the systems integration adding up to a \$30B industry.

The U.S. manufacturing economy has changed significantly over the last 30 years. Despite significant losses to Canada, China, Mexico and Japan over recent years, manufacturing still represents a major sector of the U.S. economy. Manufacturing, which includes the production of all goods from consumer electronics to industrial equipment, accounts for 12% of the U.S. GDP, and 9% of U.S. employment [WB-16]. U.S. manufacturing productivity exceeds that of its principal trading partners. We lead all countries in productivity, both per hour and per employee [FAM-11]. Our per capita productivity continues to increase with over a 100% increase over the last three decades. Indeed it is this rising productivity that keeps U.S. manufacturing competitive in the midst of recession and recovery and in the face of the amazing growth in China, India, and other emerging economies. Much of this productivity increase and efficiency can be attributed to innovations in technology and the use of technology in product design and manufacturing processes. Right now China is considered the manufacturing leader, but USA is expected to overtake China by 2020 in terms of manufacturing both in terms of value and productivity [CoC-16].

However, this dynamic is also changing. Ambitious foreign competitors are investing in fundamental research and education that will improve their manufacturing processes. On the other hand, the

fraction of the U.S. manufacturing output that is being invested in research and development has essentially remained constant over this period. The U.S. share of total research and development funding the world has dropped significantly to only 30%. Our foreign competitors are using the same innovations in technology with, in some cases, significantly lower labor costs to undercut U.S. dominance, so U.S. manufacturing industry is facing increasing pressure. Our balance of trade in manufactured goods is dropping at an alarming \$50 billion per decade. Additionally, with our aging population, the number of workers is also decreasing rapidly and optimistic projections point to two workers per pensioner in 2050 [UN-12]. Robotic workers must pick up the slack from human workers to sustain the increases in productivity that are needed with a decrease in the number of human workers [PCMR-11]. Finally, dramatic advances in robotics and automation technologies are even more critical with the next generation of high-value products that rely on embedded computers, advanced sensors and microelectronics requiring micro- and nano-scale assembly, for which labor-intensive manufacturing with human workers is no longer a viable option.

In contrast to the U.S., China, South Korea, Japan, and India are investing heavily in higher education and research [NAE-07]. India and China are systematically luring back their scientists and engineers after they are trained in the U.S. According to [NAE-07], they are "... in essence, sending students away to gain skills and providing jobs to draw them back." This contrast in investment is evident in the specific areas related to robotics and manufacturing. Korea has been investing \$100M per year for 10 years (2002-2012) into robotics research and education as part of their 21st Century Frontier Program. The European Commission invested more than \$600M into robotics and cognitive systems as part of the 7th Framework Programme. In the Horizon 2020 program that investment has been complemented by another \$900M for manufacturing and robotics. While smaller in comparison to the commitments of Korea and the European Commission, Japan is investing \$350M over the next 10 years in humanoid robotics, service robotics, and intelligent environments. Japan has also announced (2016) a major push to become leader in robotics with a 5 year investment of \$1B for industrial robotics. The non-defense U.S. federal investment in robotics and automation is small by most measures compared to these investments.

At the same time robotics is gaining significant importance for automation and logistics. In recognition of the importance of robotics Amazon during 2012 acquired the company KIVA Systems at a price of \$700M to have access to the best technology for warehouse automation. In addition, companies such as Apple [New York Times, Dec 8, 2012] and Lenovo are in-sourcing jobs as the cost of manufacturing in Asia no longer is so much cheaper that it pays off to outsource. Salary in China has grown by 340% over the last decade, whereas US salaries has seen very modest growth. In addition, during 2011 Tesla Motors in California opened up a factory for manufacturing of alternative fuel cars using heavy automation to enable all of the manufacturing to remain the United States. The recent (2015) Tesla battery factory in Nevada is another example of domestic manufacturing utilizing a high degree of automation.

2.2.2 Growth Areas

The Department of Commerce and the Council on Competitiveness [CoC-08, CoC-10, CoC-16, DoC-04] have analyzed a broad set of companies as to their consolidated annual growth rates. The data categorized for major industrial sectors is shown in the table below.

Sector	Average Growth/yr	CAGR
Robotics – manufacturing, service and medical	9%	2-22%
IP Companies	7%	6-11%
Entertainment/toys	6%	4-21%
Media / Games	8%	2-11%
Home appliances	1%	-5-6%
Capital equipment	6%	-1-13%
Automotive	3%	-1-6%
Logistics	12%	1-39%
Automation	6%	2-12%

Consolidated annual growth rates over a set of 280 U.S. companies for the period 2006-2016

Current growth areas for manufacturing include logistic including material handling, and robotics. Given the importance of manufacturing in general, it is essential to consider how technology such as robotics can be leveraged to strengthen U.S. manufacturing industry.

2.2.3 “Consumerization” of Robotics

Many advanced technologies have demonstrated that once they are introduced into the vast consumer market, the pace of innovation increases and the costs decrease. Notable examples include personal computers and mobile communications. Both of these technologies were initially developed and driven based on corporate needs and requirements. However, once they were introduced into the consumer market, the research dollars were amplified by corporate investments. This resulted in rapid technology development and dramatic cost reductions. This also spurred the creation of entirely new US companies and industries that currently make up a large percentage of the US GDP and dominate the NASDAQ.

Fostering a consumer market for robotics and robotics related technologies would have similar impacts. One simple example is the Microsoft Kinect interface. This interface which was developed for the home computer gaming market has advanced the use of voice and gesture interactions at a price point that makes it commercially viable for a number of commercial and business applications. An additional benefit of the “consumerization” of robotics would be the acceptance and familiarity by the target workforce. When people are accustomed to interacting with robots in their personal life, then they will have more acceptance of working with them in their professional life and will be less likely to view the robots as a threat. For example, two thirds of the owners of iRobot’s autonomous vacuum cleaner have named their Roomba and one third admits to taking their Roomba to visit friends.

2.2.4 A Vision for Manufacturing

U.S. manufacturing today is where database technology was in the early 1960's, a patchwork of ad hoc solutions that lacked the rigorous methodology that leads to scientific innovation. In 1970 when Ted Codd, an IBM mathematician, invented relational algebra, an elegant mathematical database model that galvanized federally funded research and education leading to today's \$15 billion database industry. Manufacturing would benefit enormously if analogous models could be developed. Just as the method to add two numbers together doesn't depend on what kind of pencil you use, manufacturing abstractions might be wholly independent of the product one is making or the assembly line systems used to assemble it.

Another precedent is the Turing Machine, an elegant abstract model invented by Alan Turing in the 1930s, which established the mathematical and scientific foundations for our now-successful high-tech industries. An analogy to the Turing Machine for design, automation and manufacturing, could produce tremendous payoffs. Recent developments in computing and information science now make it possible to model and reason about physical manufacturing processes, setting the stage for

researchers to "put the *Turing* into Manufac*Turing*". The result, as with databases and computers, would be higher quality, more reliable products, reduced costs, and faster delivery [GK-07, PCMR-11, FAM-11].

More effective use of robotics, through improved robotics technologies and a well-trained workforce, will *increase* U.S. jobs and global competitiveness. Traditional assembly-line workers are nearing retirement age. American workers are currently not well-trained to work with robotic technologies and the costs of insurance and healthcare continue to rise. Even when workers are affordable, the next generation of miniaturized, complex products with short life-cycles requires assembly adaptability, precision, and reliability beyond the skills of human workers. Widespread deployment of improved robotics and automation in manufacturing will: (a) retain intellectual property and wealth that would go off-shore without it, (b) save companies by making them more competitive, (c) provide jobs for maintaining and training robots, (d) allow factories to employ human-robot teams that safely leverage each other's strengths (e.g., human are better at dealing with unexpected events to keep production lines running, while robots have better precision and repeatability, and can lift heavy parts), (e) reduce expensive medical problems, e.g., carpal tunnel syndrome, back injuries, burns, and inhalation of noxious gases and vapors, and (f) reduce time in pipeline for finished goods, allowing systems to be more responsive to changes in retail demand.

Investments in research and education in manufacturing can revitalize American manufacturing. Investing a small portion of our national resources into a science of cost-effective, resource-efficient manufacturing would benefit American consumers and support millions of workers in this vital sector of the U.S. economy. Such investments would benefit health care, agriculture, and transportation, and strengthen our national resources in defense, energy, and security. The resulting flurry of research activity would invigorate the quality and productivity of "Made in the U.S.A." for the next fifty years. The DoD proposed NNMI Institute on Robotics could be an important mechanism to push forward robotics and automation in the US, but it must be complemented with basic research that can secure the full pipeline from basic research to economic growth.

2.3. Research Roadmap

2.3.1 The Process

The manufacturing technology roadmap describes a vision for the development of *critical capabilities* for manufacturing by developing a suite of basic *technologies* in robotics. Each critical capability stems from one or more important broad *application domains* within manufacturing. These point to the major technology areas for basic research and development (as shown in Figure 1). Integration of all the parts of this roadmap into a cohesive program is essential to create the desired revitalization of manufacturing in the U.S.

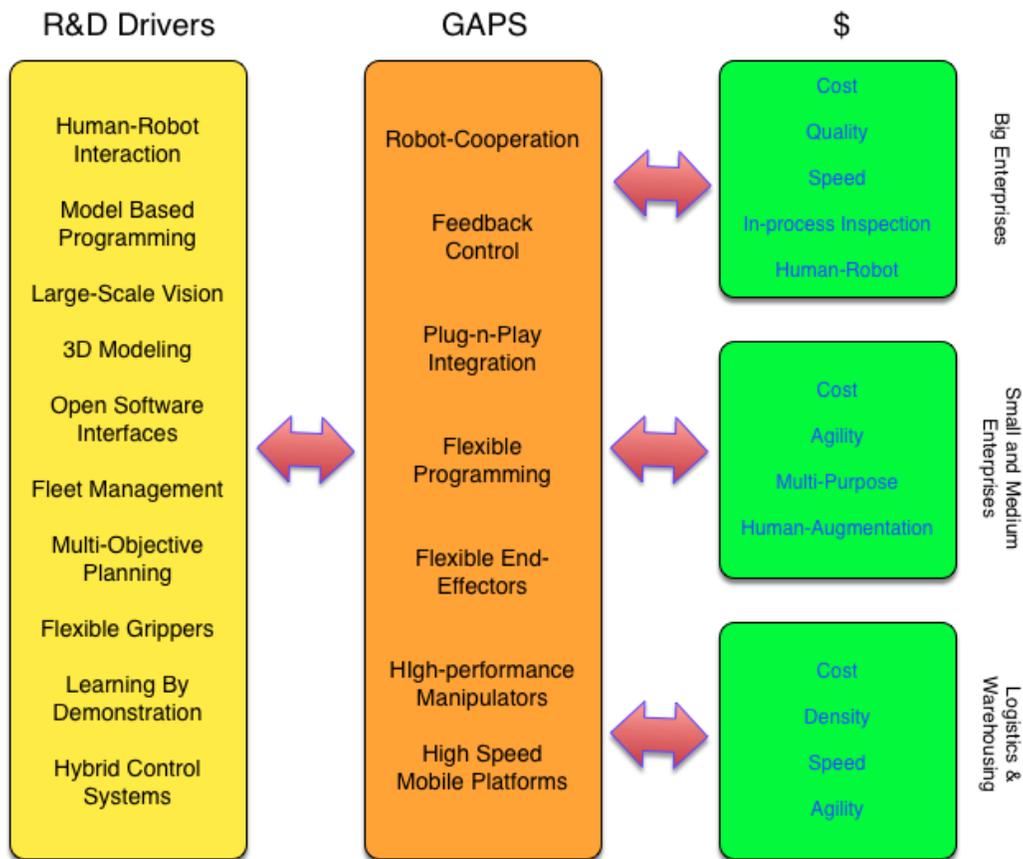


Figure 1: The roadmap process: Research and development is needed in technology areas that arise from the critical capabilities required to impact manufacturing application domains.

2.3.2 Robotics and Manufacturing Vignettes

We briefly discuss the motivating applications with vignettes and the critical capabilities required for a dramatic positive impact on the applications. The vignettes serve to illustrate paradigm changes in manufacturing and as examples of integration across capability and technology areas. The roadmap articulates five, ten and fifteen year milestones for the critical capabilities.

Vignette 1: Assembly Line Assistant Robots

An automotive manufacturer experiences a surge in orders for its new electric car design and needs to quickly merge its production capability with other earlier models already in production. Assembly tasks are rapidly reallocated to accommodate the new more efficient car model. A set of assembly line assistant robots are brought in and quickly configured to work alongside the retrained human workers on the new tasks. One practice-shift is arranged for the robot's sensor systems and robot learning algorithms to fine-tune parameters, and then the second shift is put into operation, doubling plant output in four days. Then, a change by a key supplier requires that the assembly sequence be modified to accommodate a new tolerance in the battery pack assembly. Engineers use computational tools to quickly modify the assembly sequence: then they print new instructions for workers and upload modified assembly programs to the assistant robots. This type of burst manufacturing is gradually entering our daily lives. As an example by August 2012 the company Rethink Robotics announced the robot Baxter that cost \$22k and can be programmed directly by demonstration with little or no training. All the major robotics

companies by now have collaborative robots, but few offer a solution in the range of \$25k. The cost reduction in setup and operation changes the business case for future use of automation.

Vignette 2: One-of-a-kind discrete-part manufacture and assembly

A small job shop with 5 employees primarily catering to orders from medical devices companies is approached by an occupational therapist one morning to create a customized head-controlled input device for a quadriplegic wheelchair user. Today the production of such one-of-a-kind devices would be prohibitively expensive because of the time and labor required for setting up machines and for assembly. The job shop owner reprograms a robot using voice commands and gestures, teaching the robot when it gets stuck. The robot is able to get the stock to mills and lathes, and runs the machines. While the machines are running, the robot sets up the necessary mechanical and electronic components asking for assistance when there is ambiguity in the instruction set. While moving from station to station, the robot is able to clean up a coolant spill and alert a human to safety concerns with a work cell. The robot responds to a request for a quick errand for the shop foreman in between jobs, but is able to say no to another request that would have resulted in a delay in its primary job. The robot assembles the components and the joystick is ready for pick-up by early afternoon. This happens with minimal interruption to the job shop's schedule.

Vignette 3: Rapid, integrated model-based design of the supply chain

The packaging for infant formula from a major supplier from a foreign country is found to suffer from serious quality control problems. The US-based lead engineer is able to use a comprehensive multi-scale, discrete and continuous model of the entire supply chain, introduce new vendors and suppliers, repurpose parts of the supply chain and effect a complete transformation of the chain of events: production, distribution, case packing, supply and distribution. An important aspect of the transformation is the introduction of 20 robots to rapidly manufacture the redesigned package

These vignettes may seem far-fetched today, but we have the technology base, the collective expertise, and the educational infrastructure to develop the broad capabilities to realize this vision in 15 years with appropriate investments in the critical technology areas.

2.3.3 Critical Capabilities for Manufacturing

In this section, we briefly discuss the critical capabilities and give examples of possible 5, 10, and 15 year milestones. After this, in Section 7 we describe promising research directions that would enable us to meet these milestones.

Adaptable and reconfigurable assembly

Today the time lag between the conceptual design of a new product and production on an assembly line in the U.S. is unacceptably high. For a new car, this lead-time can be as high as twenty-four months. Given a new product and a set of assembly line subsystems that can be used to make the product, we want to achieve the ability to adapt the subsystems, reconfigure them and set up work-cells to produce the product. Accordingly, the roadmap for adaptable and reconfigurable assembly includes the following goals over the next fifteen years.

- 5 years: Achieve ability to set up, configure and program basic assembly line operations for new products with a specified industrial robot arm, tooling and auxiliary material handling devices in under 24 hours.
- 10 years: Achieve ability to set up, configure and program basic assembly line operations for new products with a specified industrial robot arm, tooling and auxiliary material handling devices in one 8 hour shift.

- 15 years: Achieve ability to set up, configure and program basic assembly line operations for new products with a specified industrial robot arm, tooling and auxiliary material handling devices in one hour.

Autonomous navigation

Autonomous navigation is a basic capability that will impact the automation of mining and construction equipment, the efficient transportation of raw materials to processing plants and machines, automated guided vehicles for material handling in assembly lines and bringing completed products to inspection and testing stations, and logistics support operations like warehousing and distribution. Enabling safe autonomous navigation in unstructured environments with static obstacles, human-driven vehicles, pedestrians and animals will require significant investments in component technologies. The roadmap for autonomous navigation consists of the following milestones.

- 5 year: Autonomous vehicles will be capable of driving in any modern town or city with clearly lit and marked roads and demonstrate safe driving comparable to a human driver. Performance of autonomous vehicles will be superior to that exhibited by human drivers in such tasks as navigating through an industrial mining area or construction zone, backing into a loading dock, parallel parking, and emergency braking and stopping.
- 10 years: Autonomous vehicles will be capable of driving in any city and on unpaved roads, and exhibit limited capability for off-road environment that humans can drive in, and will be as safe as the average human driven car. Vehicles will be able to safely cope with unanticipated behaviors exhibited by other vehicles (e.g., break down or malfunction). Vehicles will also be able to tow other broken down vehicles. Vehicles will be able to reach a safe state in the event of sensor failures.
- 15 years: Autonomous vehicles will be capable of driving in any environment in which humans can drive. Their driving skill will be indistinguishable from humans except that robot drivers will be safer and more predictable than a human driver with less than one year's driving experience. Vehicles will be able learn on their own how to drive in previously unseen scenarios (e.g., extreme weather, sensor degradation).

Green manufacturing

As American architect William McDonough said, "pollution is a symbol of design [and manufacturing] failure." Our current approach to manufacturing in which components and then sub-systems are integrated to meet top-down specifications has to be completely rethought to enable green manufacturing. Today's solutions to reduce manufacturing waste mostly target process waste, utility waste and waste from shutdowns and maintenance. Our roadmap for green manufacturing emphasizes the recycling of all the components and subsystems used *throughout* the manufacturing process, starting from mining and processing of raw materials through production and distribution of finished products to recycling product materials. To create a step change new manufacturing techniques will need to be developed and products will have to be designed with this goal. For example, transitioning to additive manufacturing techniques would dramatically reduce waste for machined products/components. New logistics systems are also needed to enable widespread recycling; currently, it is often so difficult to recycle materials that companies either don't recycle or they don't universally recycle everything that they could. We are particularly concerned with re-use of the manufacturing infrastructure, recycling of raw materials, minimizing the energy and power requirements at each step and repurposing subsystems for production of new products.

- 5 years: The manufacturing process will recycle 10% of raw materials, reuse 50% of the equipment, and use only 90% of the energy used in 2010 for the same process.

- 10 years: The manufacturing process will recycle 25% of raw materials, reuse 75% of the equipment, and use only 50% of the energy used in 2010 for the same process.
- 15 years: The manufacturing process will recycle 75% of raw materials, reuse 90% of the equipment, and use only 10% of the energy used in 2010 for the same process.

Human-like dexterous manipulation

Robot arms and hands will eventually out-perform human hands. This is already true in terms of speed and strength. However, human hands still out-perform their robotic counterparts in tasks requiring dexterous manipulation. This is due to gaps in key technology areas, especially perception, robust high-fidelity sensing, and planning and control. The roadmap for human-like dexterous manipulation consists of the following milestones.

- 5 years: Low-complexity hands with small numbers of independent joints will be capable of robust whole-hand grasp acquisition.
- 10 years: Medium-complexity hands with tens of independent joints and novel mechanisms and actuators will be capable of whole-hand grasp acquisition and limited dexterous manipulation.
- 15 years: High-complexity hands with tactile array densities approaching that of humans and with superior dynamic performance will be capable of robust whole-hand grasp acquisition and dexterous manipulation of objects found in manufacturing environments used by human workers.

Model-based integration and design of supply chain

Recent developments in computing and information science have now made it possible to model and reason about physical manufacturing processes, setting the stage for researchers to "put the *Turing* into *ManufacTuring*". If achieved, as with databases and computers, would enable interoperability of components and subsystems and higher quality, more reliable products, reduced costs, and faster delivery. Accordingly our roadmap should include achievements that demonstrate the following milestones.

- 5 years: Safe, provably-correct designs for discrete part manufacturing and assembly so bugs are not created during the construction of the manufacturing facility.
- 10 years: Safe, provably-correct designs for the complete manufacturing supply chain across multiple time and length scales so bugs are not created during the design of the manufacturing supply chain.
- 15 years: Manufacturing for Next Generation Products: With advances in micro and nano-scale science and technology, and new processes for fabrication, we will be able to develop safe, provably-correct designs for any product line.

Nano-manufacturing

Classical CMOS-based integrated circuits and computing paradigms are being supplemented by new nano-fabricated computing substrates. We are seeing the growth of non-silicon micro-system technologies and novel approaches to fabrication of structures using synthetic techniques seen in nature. Advances in MEMS, low-power VLSI, and nano-technology are already enabling sub-mm self-powered robots. New parallel, and even stochastic, assembly technologies for low-cost production are likely to emerge. Many conventional paradigms for manufacturing will be replaced by new, yet-to-be-imagined approaches to nano-manufacturing. Accordingly the roadmap for nano-manufacturing and nano-robotics must emphasize basic research and development as follows.

- 5 years: Technologies for massively parallel assembly via self-assembly and harnessing biology to develop novel approaches for manufacturing with organic materials.
- 10 years: Manufacturing for the post-CMOS revolution enabling the next generation of molecular electronics and organic computers
- 15 years: Nano-manufacturing for nano-robots for drug delivery, therapeutics and diagnostics.

Perception for unstructured environments

Automation in manufacturing has proven to be simpler for mass production with fixed automation, and the promise of flexible automation and automation for mass customization has not been realized except for special cases. One of the main reasons is that fixed automation lends itself to very structured environments in which the challenges for creating “smart” manufacturing machines are greatly simplified. Automation for small lot sizes necessitate robots to be smarter, more flexible, and able to operate safely in less structured environments shared with human workers. In product flow layouts for example, robots and other machines go to various operation sites on the product (e.g., an airplane or a ship) to perform their tasks, whereas in a functional layout, the product travels to various machines. The challenges of one-of-a-kind manufacturing exacerbate these difficulties. The roadmap for perception includes the following milestones.

- 5 years: 3-D perception enabling automation even in unstructured typical of a job shop engaged in batch manufacturing operations
- 10 years: Perception in support of automation of small lot sizes, for example, specialized medical aids, frames for wheelchairs, and wearable aids.
- 15 years: Perception for truly one-of-a-kind manufacturing including customized assistive devices, personalized furniture, specialized surface and underwater vessels, and spacecrafts for planetary exploration and colonization.

Intrinsically safe robots working with humans: The democratization of robots

Much discussion has taken place around the topic of intrinsically safe robots, not the least of which is clarifying what the term actually means. Intrinsically safe equipment is defined as “equipment and wiring which is incapable of releasing sufficient electrical or thermal energy under normal or abnormal conditions to cause ignition of a specific hazardous atmospheric mixture in its most easily ignited concentration.” *ISA-RP12.6* In short, an intrinsically safe piece of equipment won't ignite flammable gases. This is certainly a requirement that must be addressed with robot systems, as with any equipment or systems designed for the manufacturing environment. However, it is clear that the term carries a heavier burden when applied to robots, perhaps related to the definition of “intrinsic” itself.

Intrinsic: belonging to the essential nature or constitution of a thing; originating and included *wholly within* an organ or part (Merriam-Webster online dictionary).

That is the crux of it: the expectation is that robots must be safe from the inside out, completely harmless to humans, no matter what the cost. It is part of the cultural fear that we might create something that turns on us...oh wait, we've already done that. In truth, there is no foolproof system.

To offer a comparison, consider the automobile: cars are dangerous. To be sure, the first horseless carriages were a menace to the other more traditional versions on the road, yet we have advanced to the point where people pass one another driving on the highway at speeds exceeding 70 mph. This is not because automobiles are intrinsically safe, but because we have learned to accept the risk. We created, over time, a transportation system that relies on human understanding of the capabilities, limitations, and risks inherent to operating a car on the highway. We *democratized* the automobile--that is, made it relate to, appeal to, and available to masses of people. Thus it became part of our society.

To democratize robots in the manufacturing arena, a similar model of risk/responsibility must be developed. Like driving, working in a manufacturing environment already presents a certain level of danger. The goal is not to *increase* that level when robots are added to the mix. An acceptable metric

for ascertaining whether that goal is met is the number of Lost Work Days. If it does not increase due to automation or robotics, then we are on the path to democratization. We must continue to develop and refine the existing safety standards, incorporating systems-engineered solutions for user-defined tasks.

Indeed, we must start with safety, but continue to encourage the development of collaborative solutions for user-communicated needs. This includes defining the capabilities, limitations, and risks inherent to each implementation. Acceptance of a risk/responsibility model for robots in the manufacturing environment will be driven by the diversity of demand for innovation. Social understanding of humans and robots in the workplace and culture at large will come with the democratization of robots. This can only happen over time as the consumer base of robot-users broadens. Natural language programming, control studies, and advances in materials technology are examples of potential pathways that can speed the process.

The roadmap for robots working with humans is as follows:

- 5 years: Broad implementation of easily programmed and adaptable safety-rated soft-axis guarding for fixed or mobile assembly robots on the factory floor.
- 10 years: Systems that automatically detect and respond appropriately to conforming/non-conforming human behaviors in the workspace while maintaining consistent performance.
- 15 years: Systems that can recognize, work with, and adapt to human or other robot behaviors in an unstructured environment, i.e. construction zones or newly configured manufacturing cells.

2.4. References

[BEA-16] Bureau of Economic Analysis, U.S. Department of Commerce Press Release, July 2016. <http://www.bea.gov/newsreleases/industry/gdpindustry/2016/pdf/gdpind116.pdf>.

[CoC-08,CoC-10] Council on Competitiveness, Competitiveness Agenda - New Challenges, New Answers, November 2008, 2010 (www.compete.org)

[CoC-16] 2016 Manufacturing Competitiveness Index, March 2016 (http://www.compete.org/storage/2016_GMCI_Study_Deloitte_and_Council_on_Competitiveness.pdf)

[Cloud-12] Cloud Robotics summary with links to resources, recent papers, presentations, etc: <http://goldberg.berkeley.edu/cloud-robotics/>

[DoC-04] U.S. Dept of Commerce, Manufacturing in America, Jan 2004 (ISBN 0-16-068028-X).

[E07] U.S. Fact Sheet, Economist, June 2007.

[EF-06] Fuchs, E. The Impact of Manufacturing Offshore on Technology Development Paths in

the Automotive and Optoelectronic Industries. Ph.D. Thesis. M.I.T. Cambridge, MA: 2006.

[GK-07] Goldberg, K., Kumar, V, "Made in the USA" can be Revitalized, San Jose Mercury News: Op-Ed, 24, October 2007

[MI-11] Milken Institute, Jobs for America, June 2011.

[NAE-07] Rising Above The Gathering Storm: Energizing and Employing America for a Brighter Economic Future, National Academy of Engineering, 2007.

[NRR- 13] Christensen, H. I., (ed.), National Robotics Roadmap, CCC, Washington, DC, March 2011.

[PCMR-11] PCAST, REPORT TO THE PRESIDENT ON ENSURING AMERICAN LEADERSHIP IN ADVANCED MANUFACTURING, June 2011

[PCMR-14] PCAST, Accelerating U.S. Advanced Manufacturing AMP2.0 Steering Committee Report, October 2014

[NNMI-12] PCAST, A preliminary design for a National Network for Manufacturing Innovation, 2012

[OECD] Organisation for Economic Co-operation and Development (OECD). 2013. *Main Science and Technology Indicators*. Vol. 2013/1. Paris

[UNESCO] United Nations Educational, Scientific and Cultural Organization (UNESCO), Institute for Statistics, Science and Technology. Table 25: Gross domestic expenditures on R&D, 2013. <http://www.uis.unesco.org/ScienceTechnology/Pages/default.asp>

[WB-06] Where is the Wealth of Nations? The International Bank for Reconstruction and Development, The World Bank, 2006.

[WB-11] World Bank, World Economic Indicators, 2011-report

[WB-16] "World Bank. 2016. The Little Data Book 2016. Washington, DC. World Bank. <https://openknowledge.worldbank.org/handle/10986/23968>

2.5. Contributors

This section is based on a revision of the 2013 section on manufacturing. The revision was performed by a small editorial group. The responsibility of the report lies entirely with the authors.

2016 Contributors were

Henrik I Christensen
UC San Diego

Howie Choset
Carnegie Mellon University

Venkat Krovi
Clemson University

Bill Smart
Oregon State University

The 2013 committee was composed of (affiliation has not been updated):

Jeff Baird
Adept Technologies

Gary Bradski
Industrial Perception

Michael Branicky
Case Western Reserve Univ

Rodney Brooks
Rethink Robotics

Jenny Burke
The Boeing Company

Henrik I Christensen
Georgia Tech

John Enright
KIVA Systems

Clay Flannigan
South Western Research Inst

Phil Freeman
The Boeing Company

Thomas Fullbrigge
ABB Research Corp

Joe Gemma
Staubli

SK Gupta
University of Maryland

Ginger Hildebrand
Schlumberger

Vijay Kumar
Univ. of Pennsylvania

Peter Luh
Univ. of Connecticut

Matt Mason
Carnegie Mellon University

Elena Messina
NIST

Erik Nieves
Motoman / Yaskawa

Curtis Richardson
Spirit Aerospace

Daniela Rus
MIT

Mark Spong
UT Dallas

Jason Tsai
FANUC Robotics

Stuart Shepherd
KUKA Robotics

3. Next Generation Consumer and Professional Services

3.1. Introduction

Service robotics is defined as those robotic systems that assist people in their daily lives at work, in their houses, for leisure, and as part of assistance in aging, and/or to help people with physical, cognitive, or sensory impairments. Industrial robots typically to automate tasks to achieve a homogenous quality of production or a high speed of execution; In contrast, service robots perform tasks in spaces occupied by humans and often in direct collaboration with people.

Service robotics is normally divided into professional and domestic consumer services. Generally speaking, professional service robotics is expected to serve as a workforce multiplier for increased economic growth, while domestic service robotics is expected to enable sustained personal autonomy. Professional service applications include inspection of power plants and infrastructure such as bridges, logistics applications such as delivery of meals and pharmaceuticals at hospitals, as well as commercial-scale lawn and cleaning technologies. The annual growth in professional service robots is 30%.

Personal service robots, on the other hand, are deployed for assistance to people in their daily lives in their homes, or as assistants to overcome mental and physical limitations. By far, the largest group of personal service robots consists of domestic vacuum cleaners; over 10 million iRobot Roombas alone have been sold worldwide, with continued growth in this market each year. In addition, a large number of robots have been deployed for entertainment applications such as artificial pets, personal assistants, etc. At 28% annual growth, with 4.7 million personal service robots sold in 2014 alone, this market is expected to remain one of the most promising in robotics for the coming years.

Autonomous aerial vehicles and self-driving cars are two additional technical areas that span a number of service applications and are poised to become a disruptive technology over the next 5-10 years. Below we discuss the range and impact of service robot technologies.

3.2. Principal Markets and Drivers

There is general agreement among those present at the meeting that we are still 10 to 15 years away from a wide variety of applications and solutions incorporating full-scale, general autonomous functionality. Some of the key technology issues that need to be addressed to reach that point are discussed in a later section of this report. There was further agreement among those present, however, that the technology has sufficiently progressed to enable an increasing number of limited scale and/or semi-autonomous solutions that are pragmatic, affordable, and provide real value. Commercial products and applications based on existing technology have already begun to emerge and more are expected as entrepreneurs and investors realize their potential.

One of the key factors contributing to the identified trends is our aging population. This impacts service robotics both in terms of the need to address a shrinking workforce as well as the opportunity to develop solutions that will meet their healthcare needs. As shown in the figure below, the United States is on the threshold of a 20-year trend that will see a near doubling of the number of retired workers as a percentage of the current workforce—from just over 2 retirees for every 10 workers today to just over 4 retirees for every 10 workers in 2030. In Japan, the situation is even worse and has fueled a major national initiative to develop the robotics technology needed to help care for their rapidly aging population. Generally speaking, professional service robotics is expected to serve as a workforce multiplier for increased economic growth, while domestic service robotics is expected to enable sustained personal autonomy.

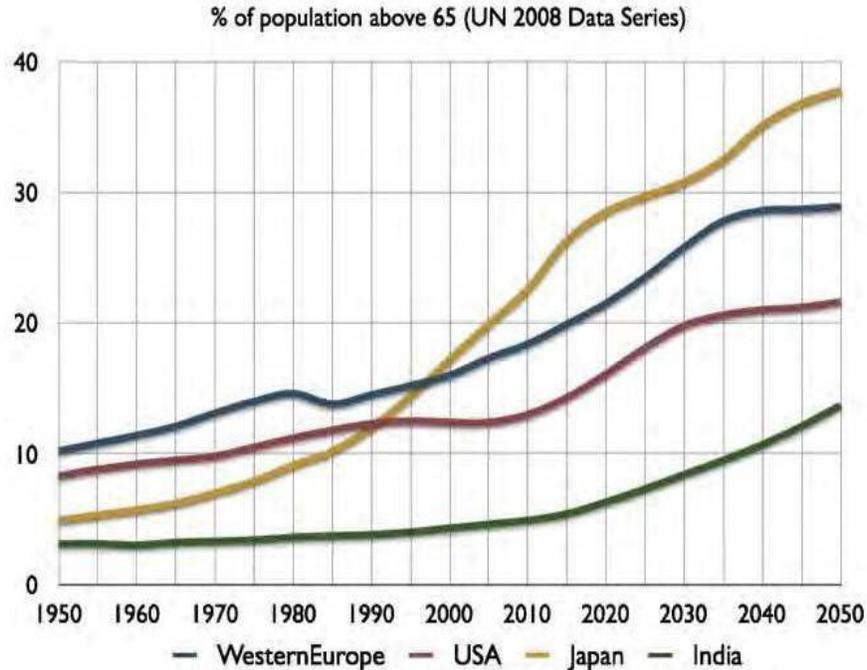


Figure 1 Changes in demographics in Western Europe, USA, Japan, and India

While increasing productivity and reducing costs are the common denominator of service robotics, each system is expected to uniquely provide a compelling solution to certain critical market specific issues or needs. Below we discuss the current state-of-the-art and near-term opportunities with respect to professional service robotics, consumer service robotics and self-driving vehicles.

3.2.1 Professional Service Robots

Professional service robots typically augment people for execution of tasks in the workplace. According to the IFR/VDMA World Robotics, more than 172,000 professional robots are in use today and the market is growing rapidly every year. Applications of professional service robots include logistics, professional cleaning, inspection, lawn care, and various other disciplines. In the following subsections we highlight significant areas of growth. Several closely related applications, including assistive healthcare, agriculture, and defense are covered in other sections of this report.

3.2.1.1 Logistics: automated delivery, movement of goods

Logistics and transport of components and manufactured goods comprises 10-15% of manufactured goods in the United States, and approximately 8% of the U.S. GDP, according to the “State of Logistics” report [CG13]. With a trend towards “just-in-time manufacturing” to reduce inventory overhead, logistics becomes an increasingly important component of supply-chain management. Driving down costs as well as improving responsiveness and flexibility of freight-in, freight-out, and internal logistics will significantly impact manufacturing in the United States. Amazon’s Kiva Systems robots, as well as increasingly automated FedEx or UPS distribution centers, are leading examples of robotics and automation applied in this manner. However, these examples are in controlled warehouse environments, with an entire space tailored to the purpose; Amazon’s shelves and warehouse infrastructure are designed around the robotic distribution of goods. To extend the automation of the logistics chain into the world, robots must have mobility that matches human mobility – robots must negotiate stairs, elevators, doorways, curbs, broken concrete,

cluttered environments, and go where people go. This type of advanced mobility is becoming realistic for robotic systems, legged and otherwise - and with such a solution, logistics will become fast, 24/7, on-demand, inexpensive, predictable, and well-tracked.



Figure 2 Logistics industry robots for store inventory from Bossa Nova, warehouse distribution from Fetch, hospital delivery from Aethon and hotel delivery from Savyoke.

Several existing and newly formed companies are beginning to address this market, including hotel delivery robots from Savyoke, hospital delivery robots from Aethon and Vecna, store inventory robots from Bossa Nova, aerial delivery drones from Amazon Prime Air and Google Project Wing, last-mile local delivery services from Starship Technologies, and people-friendly warehouse solutions from Fetch. Exponential growth is projected to continue for the logistics robotics market, making this one of the highest impact areas for investment over the coming 15 years.

3.2.1.2 Infrastructure

Infrastructure protection and inspection is another major application area for service robots. Robotics technology has tremendous potential to enhance the inspection and maintenance of our nation's bridges, highways, pipelines, and other infrastructure using ground, aerial and maritime robots. Unmanned maritime and aerial vehicles can inspect bridges and ports; unmanned ground vehicles can survey underground utilities such as buried tunnels of wires and pipes; and unmanned aerial vehicles can survey pipelines and the electrical grid. These robots will work in teams with one another and with human operators, using shared and full autonomy along with state-of-the-art sensors to extend the effectiveness of human judgment.

Already, the technology has been adapted to develop automated pipeline inspection systems that reduce maintenance and rehabilitation costs by providing accurate, detailed pipe condition information. Such systems, based on advanced multi-sensor and other robotics technology, are designed for underground structures and conditions that are otherwise difficult to inspect, including large diameter pipes, long-haul stretches, inverts, crowns, culverts, and manholes, as well as in-service inspections. Robotic platforms are able to navigate critical wastewater infrastructure to inspect sewer pipe unreachable by traditional means and produce very accurate 3D images of the pipe's inside surface.

A similar role of inspection, monitoring, and surveillance is performed by robots in the oil and gas industry, where the use of service robots has been projected to increase by 20% over the next 5 years. The inspection information, captured in digital form, serves as a baseline for future inspections and, as a result, can automatically calculate defect feature changes over time, leading to reduced human intervention, increased operational efficiency, reduced cost and improved safety.

With the aid of robotics technologies, inspection tasks can be done routinely with the transportation, energy, and communications infrastructure remaining in service, rather than the traditional "take it out of service and then inspect" scenario. The robots can be re-tasked to concentrate on certain areas based on terrorism alerts, to help prepare for a notice event such as Hurricane Sandy.



Figure 3 Infrastructure inspection robots: unmanned aerial inspection by ULC, and pipe inspection crawlers from EnviroSight and Honeybee Robotics.

Unmanned aerial vehicles (UAVs) have emerged as cost-effective platforms for performing infrastructure inspection tasks. UAV technology has become smaller, cheaper, and more reliable in the last 5 years, and several companies now sell platforms capable of autonomous navigation in the \$2000-\$3000 range (DJI, 3DRobotics and others). Cameras, laser scanners, infrared cameras, and other sensors can be mounted on these platforms, allowing for aerial inspection and monitoring.

Prior work has shown it's feasible to deploy more than 50 UAVs at once with existing logistical systems [Chung, 2016], and recent programs like the DARPA-funded fast lightweight autonomy program have pushed the frontiers of autonomous operation with current technology. Existing challenges in UAV systems include, but are not limited to, (1) safety, security, and privacy concerns with operation, (2) robustness to failure and failsafes when errors occur to allow for safe recovery, (3) in-the-loop adaptation to changing environments and conditions, (4) data processing for large datasets obtained during flight, (5) novel vehicle design, and (5) multi-vehicle coordination and collision avoidance in cluttered environments [Chung16].

FAA regulations have provided roadblocks for UAV operations in US airspace, which have previously required Certificates of Waiver of Authorization (COAs) only available to public institutions. The FAA has recently implemented new programs that allow for more flexible operation, including the Section 333 exemption for commercial operators and the Small UAS Rule (Part 107) for operating aircraft under 55 lbs. These new regulations have potential to open up UAV applications for inspection (bridges, buildings, powerlines, etc.), construction, photography (real estate, surveying), law enforcement, and agriculture. However, restrictions on line-of-sight operation and day operation still restrict use for applications like package delivery, urban search and rescue, firefighting, and medical supply and disaster supply delivery.

3.2.1.3 Telepresence and Tele-Labor

Robots designed to assist in and facilitate human communication and information sharing is another significant service robot application area, with telepresence robots holding the largest share of this market. Telepresence robots represent the next stage of evolution beyond stationary video conferencing, taking advantage of the existing telecommunications infrastructure to enable more effective collaboration and communication. In corporate settings, telepresence systems such as the Beam, VGo and Double enable remote or traveling employees to feel more physically connected to their place of work. The same platforms have the potential to make school accessible to children who otherwise cannot have a physical presence due to illness, injury or other physical challenges.

In healthcare, telemedicine robotic platforms, such as the InTouch Health Vita pictured below, feature cameras, microphones and speakers that allow physicians and patients to see and talk to each other. By facilitating remote interaction, while granting the physician the freedom of physical mobility and control over sensors that is unavailable with a standard computer, telemedicine robots expand community access to

medical specialists, especially in rural areas where there is a shortage of doctors. Telemedicine is currently being used in leading hospitals to diagnose patients suspected of suffering strokes, when every minute is crucial to prevent serious brain damage, as well as in intensive care units to provide access to specialists in areas such as neurology, cardiology, neonatology, pediatrics and mental health. Future remote robotic telepresence systems could have a major impact on acute and post-operative care, as well as on the long-term management of chronic conditions, allowing surgeons and therapists to visit patients and mentor/assist each other in pre-, intra-, post-operative, long-term recovery, and therapy scenarios. Researchers are already actively investigating the use of telepresence as an aging-in-place technology for use in the homes of older adults in order to enable the shrinking numbers of healthcare workers to check on their growing client list. Robot-mediated health communication has the potential to significantly lower healthcare costs and increase patients' access to the best care, allowing remote specialists to mentor local surgeons during a procedure, or therapists to conduct in-home assessments remotely.



Figure 4 Telepresence Beam platform from Suitable, telemedicine robot Vita from InTouch Health and the social robot Pepper from SoftBank Robotics.

3.2.1.4 Entertainment

Robotics technology is being incorporated into the entertainment industry throughout a wide range of applications, including automation for the movie industry such as Bot & Dolly (now at Google), intelligent and interactive toys like Anki Cozmo, programmable build-kit toys like Lego Mindstorms.



Figure 5: On the left: Cozmo robot, from Anki Robotics; incorporates sensors, onboard computing and "personality," along with a computer interface for user programming. On the right: The Iris cinematic motion control system by Bot & Dolly (company recently purchased by Google).

Games are becoming increasingly interactive, and including more and more robotics technology. For example, the game Pokemon GO is the first major commercial success for an augmented reality application. Augmented reality has long been envisioned as an important interface for control and interaction with robots, but with new commercial successes in gaming, experience and development of augmented reality should accelerate.

3.2.2 Consumer Service Robots

Consumer service robots are deployed for assistance to people in their daily lives in their homes, or as assistants to compensate for mental and physical limitations. So far, service robots for personal and domestic use are mainly in the areas of domestic (household) robots, which include vacuum and floor cleaning, lawn-mowing robots, and entertainment and leisure robots, including toy robots, hobby systems, hobby drones, and education platforms.

In 2014, it was estimated by IFR that 3.3 million robots were sold for domestic tasks, including vacuum cleaning, lawn-mowing, window cleaning and other types. The actual number might, however, be significantly higher, as the IFR survey is far from having full coverage in this domain. The value was approximately \$1.2 billion, representing an increase of 24% over the previous year. Approximately 1.3 million entertainment robots were sold in 2014 according to the IFR, 40% more than in 2013. Drones, similarly, are a fast-growing market, with an estimated 4.9 million units sold in 2014, with estimated growth of 30% per year in the near future. Applications range from first-person view (FPV) racing, to aerial photography, to recreational flying. More recent numbers are not yet available, but growth in this market is projected to continue increasing as the cost of technologies drops while capabilities increase.

Over the next 5-10 years, the consumer service robot market is expected to benefit from continued developments in the professional sector, with technologies such as telepresence, home inspection and improved cleaning robots mass-produced at reduced prices.



Figure 6 Popular consumer robots: Parrot Disco FPV Drone, Wonder Workshop Dash robot for education, iRobot Roomba vacuum and John Deere Tango lawn mower.

3.2.3 Transportation

Public transportation is on the verge of become increasingly automated. As robotics technology continues to improve and mature, unmanned transportation systems and solutions developed for limited-scale environments, such as airports, will be adapted for implementation in urban centers and other general purpose environments.

Robotics technology will significantly affect every aspect of how we transport people and goods in the coming decades, from personal transportation systems to intelligent highways to autonomous public transportation systems. Companies such as Segway and Toyota have introduced personal transportation robots that are ridden in standing position and controlled by internal sensors that constantly monitor the rider's position and automatically make the according adjustments.

Fully autonomous cars are also on the verge of coming to market. In 2016 Google passed the 2-million-mile mark for its self-driving vehicles, Tesla deployed a highway autopilot, Uber began transporting customers in autonomous vehicles, and the National Highway and Transportation Safety Administration (NHTSA) issued updated guidance to promote safe development of highly autonomous vehicles. At the same time, many carmakers are taking steps to make all vehicles "smarter" by providing shared autonomy features such as lane keeping, parking and braking assistance. Autonomy is also being explored for maritime shipping, rail, trucking and bus transport.

In parallel to the development of smarter cars, researchers are seeking to address transportation issues through the creation of “smart roads” by installing sensors, cameras, and automatic toll readers. A public-private national initiative called Vehicle Infrastructure Integration (VII) has been launched to merge smart cars and smart roads to create a virtual traffic information network and to bust up gridlock. Mass transportation systems are also expected to adopt robotics technology to provide operators with greater situational awareness and navigation assistance in crowded urban corridors thereby helping to control costs and increase safety.



Figure 8 Google self driving car and Otto self-driving truck.

3.3 Near-Term Opportunities and Factors Affecting Commercialization

Commercialization and economic impact of robotics technology is significantly affected not only by the technological progress, but also by the legal and policy framework, the extent of education and training for employees in new fields.

There are many parallels between the rise of computer science in the 20th century and its impact on the world’s economy with the current rise of Robotics technology. Like computing hardware, robot hardware is expensive to develop and to produce. The software component for Robotics technology is similar to operating systems in computing, in that common software platforms such as Robot Operating System (ROS) will enable faster progress.

The scale of manufacture as applications succeed will enable an accelerated pace of progress. Recent progress in shared autonomy for vehicles has led to economies of scale for important components of all robotics technologies. For example, laser scanners used on autonomous vehicles have been large (near a cubic foot), heavy, and expensive, costing many thousands of dollars. With a proven market to drive development, size is dropping to a few cubic inches, and cost is dropping to hundreds of dollars. With such improvements in these sensors, all other robots that operate in human spaces will benefit and improve.

With sustained research and development, we expect the following milestones to be reachable within the next 5, 10, and 15 years:

- **5 years:** Robots create semantic maps about their environment through exploration and physical interaction but also through instruction from humans. Robots exploit diverse mobility mechanisms in research laboratories to navigate safely and robustly in unstructured 2D environments and perform simple pick and place tasks. Relevant objects are either from a very limited set or possess specific properties. They are able to reason about tasks of moderate

complexity, such as removing obstructions, opening cabinets, etc. to obtain access to other objects.

- Increased use of factory warehouse logistics robots, to manage inventory and move materials.
- Autonomous vehicles will be capable of driving in any modern town or city with clearly lit and marked roads and demonstrate safe driving comparable to a human driver. Performance of autonomous vehicles will be superior to that exhibited by human drivers in such tasks as navigating through an industrial mining area or construction zone, backing into a loading dock, parallel parking, and emergency braking and stopping.

• **10 years:** Given an approximate and possibly incomplete model of the static part of the environment (possibly given a priori or obtained from data bases via the Internet, etc.), service robots are able to reliably plan and execute a task-directed motion in service of a mobility or manipulation task. The robot builds a deep understanding of the environment from perception, physical interaction, and instruction. The robot navigates multi-floor environments through stairways. The robot modifies its environment to increase the chances of achieving its task (e.g., remove obstructions, clear obstacles, turn on lights), and detects and recovers from some failures.

- Commercial applications come to market in package delivery, with application-specific use of UAVs, ground vehicles, and legged robots.
- Autonomous vehicles will be capable of driving in any city and on unpaved roads, and exhibit limited capability for off-road environment that humans can drive in, and will be as safe as the average human-driven car. Vehicles will be able to safely cope with unanticipated behaviors exhibited by other vehicles (e.g., break down or malfunction). Vehicles will also be able to tow other broken down vehicles. Vehicles will be able to reach a safe state in the event of sensor failures.

• **15 years:** Service robots including multiple mobility mechanisms such as legs, tracks, and wheels perform high-speed, collision-free, mobile manipulation in completely novel, unstructured, dynamic environments. They perceive their environment, translate their perceptions into appropriate, possibly task-specific local and global/short- and long-term environmental representations (semantic maps), and use them to continuously plan for the achievement of global task objectives. They respond robustly to dynamic changes in the environment (e.g., unexpected perturbation due to being pushed or jostled). They are able to interleave exploratory behavior when necessary with task-directed behavior. They interact with their environment and are able to modify it in intelligent ways so as to ensure and facilitate task completion. This includes reasoning about physical properties of interactions (sliding, pushing, throwing, etc.) between the robot, objects it comes into contact with, and the static parts of the environment.

- Increasing use of robots for all stages of logistics to achieve driverless-freight through autonomous trucks, autonomous planes, small robots delivering packages, warehouse robots moving heavy objects.
- Autonomous vehicles will be capable of driving in any environment in which humans can drive. Their driving skills will be indistinguishable from humans except that robot drivers will be safer and more predictable than a human driver with less than one year's driving experience. Vehicles will be able learn on their own how to drive in previously unseen scenarios (e.g., extreme weather, sensor degradation).

3.4 References

- [CG13] Third Party Logistics, Cap Gemini, 2013 https://www.capgemini.com/resource-file-access/resource/pdf/2013_Third-Party_Logistics_Study.pdf
- [Chung16] Chung, Timothy H., et al. "Live-fly, large-scale field experimentation for large numbers of fixed-wing UAVs." *Robotics and Automation (ICRA), 2016 IEEE International Conference on*. IEEE, 2016.

3.5 Contributors

The Service Robotics Section was discussed at the previously mentioned workshops and the observations were used for update of the roadmap by the following group of people:

Sonia Chernova Georgia Institute of Tech.	Jonathan Hurst Oregon State University	Hadas Kress-Gazit Cornell University
Anca Dragan UC Berkeley	Bill Smart Oregon State University	Venkat Krovi Clemson University
Dieter Fox University of Washington	Geoff Hollinger Oregon State University	

The update is based on the 2013 version of the roadmap which was prepared by (note affiliations have not been updated)

Odest Jenkins Brown University	Stefano Carpin UC Merced	Jana Kosecka George Mason University
Charlie Kemp Georgia Institute of Tech	Henrik I Christensen Georgia Institute of Tech.	Siddhartha Srinivasa Carnegie Mellon University
Joshua Smith University of Washington	Pieter Abbeel UC Berkeley	Gaurav Sukhatme University of Southern Calif.
Dieter Fox University of Washington	Ashutosh Saxena Cornell University	Stergios Roumeliotis Minnesota University
Ira Snyder Microsoft Research	Randy Trower Hoaloa Robotics	Kurt Konolige Industrial Perception
Radu Rusu Open Perception	James Kuffner Google	Trevor Blackwell Anybots
Parag Batavia Neva Systems	Bill Thomasmeyer Robotics Tech. Consortium	Larry Sweet Symbolic

4. Health, Independence and Quality of Life

4.1. Introduction

Healthcare robotics has the potential to increase access to care and to address critical workforce shortages in the healthcare system. Over 20% of the world's population has a motor, cognitive or sensory impairment, and with a rapidly aging population this number will only multiply.

Robotics technology has the potential to be a game-changer across multiple health sectors. Robots can be used to enable people with disabilities, support caregivers, and aid the clinical workforce. Robots have the potential to improve patient outcomes, reduce costs, and make people healthier and more independent and productive across the lifespan.

4.1.1. Definition of the Field/Domain

Robots have become routine in the world of manufacturing and other repetitive labor. While industrial robots were developed primarily to automate dirty, dull, and dangerous tasks, healthcare robots are designed for entirely different environments and tasks – those that involve direct and often unstructured and dynamically changing interaction with human users, in the surgical theater, the rehabilitation center, and the family home.

Robotics is already beginning to affect healthcare. Telerobotic systems such as the da Vinci Surgical System are being used to perform surgery that offer more intuitive and finer dexterity control of surgical instruments, potentially resulting in more reliable outcomes in common procedures such as hysterectomies and gall bladder removals. The use of robotics as part of a computer-integrated surgery system enables accurate, targeted medical interventions. It has been hypothesized that surgery and interventional radiology will be transformed through the integration of computers and robotics much in the way that manufacturing was revolutionized by automation several decades ago. Haptic devices, a form of robotics, are already used for simulations to train medical personnel. Robotic systems such as the MIT-Manus (commercially, InMotion), Lokomat (Hocoma) and Proficio (Barrett Medical) are also successfully delivering physical and occupational therapy. Rehabilitation robots enable a greater intensity of treatment that is continuously adaptable to a patient's needs. They hold the potential to amplify the impact of physical therapists through a greater number of hours spent in therapy, and in some scenarios have already proven more effective than conventional approaches, especially in assisting recovery after stroke, the leading cause of permanent disability in the US. The future potential for robots in convalescence and rehabilitation is even greater. Experiments have demonstrated that robotic systems can provide therapy oversight, coaching, and motivation that supplement human care with little or no supervision by human therapists, and can continue long-term therapy in the home both after hospitalization and for chronic conditions. Such systems have a therapeutic role not only for movement disorders (such as those resulting from stroke, traumatic brain injury, and other trauma) but also as intervention and therapeutic tools for social and behavioral disorders including autism spectrum disorder, ADHD, and other pervasive and growing disorders among children today. We also have seen emergence on the commercial market of human-operated wheelchair-mounted robotic arms with FDA-approval (e.g. the JACO from Kinova Robotics).

Robotic technology has also had a major impact on our quality of life. Home health care, mobility, wellness and well-being are being positively impacted by assistive robotics, human-robot interaction,

advanced prosthetics, and smart sensing, all areas that are central to the NRI. The emergence of “Smart Cities” and Internet of Things (IOT) initiatives led by private industry is supported by new sensing and robotic technologies coupled with advanced networked software, all components of NRI research.

Robotics technology also has a role in enhancing basic research into human health. The ability to create a robotic system that mimics biology is an important way to study and test how the human body and brain function. Furthermore, robots can be used to acquire data from biological systems with unprecedented accuracy, enabling us to gain quantitative insights into both physical and social behavior. Finally, socially interactive robots can be used to study human behavior as well as aid in diagnosis of behavioral disorders.

The spectrum of niches for robotic systems in medicine and health thus spans a wide range of environments (from the operating room to the family room), user populations (from the very young to the very old, from the infirm to the able bodied, from the typically developed to those with physical and/or cognitive deficits), and interaction modalities (from hands-on surgery to hands-off rehabilitation coaching). Technological advances in robotics have clear potential for stimulating the development of new treatments for a wide variety of diseases and disorders, for improving both the standard and accessibility of care, and for enhancing patient health outcomes.

4.1.2. Societal Drivers

Existing medical procedures can be improved to be less invasive and produce fewer side effects, resulting in faster recovery times and improved worker productivity. Revolutionary efforts aim to develop new medical procedures and devices, such as micro-scale interventions and smart prostheses, which would substantially improve risk-benefit and cost-benefit ratios. More effective methods of training of medical practitioners would lower the number of medical errors. Objective approaches for accountability and certification/assessment also contribute to this goal. Ideally, all these improvements would lower costs to society by lowering impact on families, caregivers, and employers. More directly, health care costs would be lowered due to improved quality, through fewer complications, shorter hospital stays, and increased efficiency of treatment.

Economic and population factors must be considered. In the United States, around 10% of the population is uninsured [CDC 2015]; many others are under-insured. The situation prevents individuals from receiving needed health care, sometimes resulting in loss of function or even life, and also prevents patients from seeking preventative or early treatment, resulting in worsening of subsequent health problems. Access to health care is most directly related to its affordability. Access to physically interactive therapy robots promise to reduce the cost of clinical rehabilitative care and are the focus of an ongoing Veteran’s Administration study of their cost-effectiveness. Socially assistive robotics efforts are working toward methods that could provide affordable in-home technologies for motivating and coaching exercise for both prevention and rehabilitation. It is also a promising domain for technologies for care taking for the elderly, toward promoting aging in place (i.e., at home), motivating cognitive and physical exercise toward delaying the onset of dementia, and providing companionship to mitigate isolation and depression.

Access to health care is also related to location. When disasters strike and result in human injury, distance and unstructured environments are obstacles to providing on-site care and removing the

injured from the scene. This has been repeatedly demonstrated in both natural disasters (such as earthquakes and hurricanes) and man-made disasters (such as terrorist attacks). Similar problems occur in the battlefield; point-of-injury care is needed to save the lives of many military personnel. Some environments, such as space, undersea, and underground (for mining) are inherently far from medical personnel. Finally, rural populations can live prohibitively far from medical centers that provide specialized health care. Telemedicine and assistive robotics can provide access to treatment for people outside populated areas and in disaster scenarios.

Population factors indicate a growing need for improved access and quality of health care. Demographic studies show that the US population will undergo a period of significant population aging over the next several decades. Specifically, the US will experience an approximately 40% increase in the number of older adults by 2030. Japan will see a doubling in the number of people over the age of 65, Europe will have a 50% increase, and the US will experience a ~40% increase in the number of older adults by 2030. The number of people with an age above 80 will increase by more than 100% across all continents. Advances in medicine have increased the life span and this, in combination with reduced birthrates, will result in an aging of society in general. This demographic trend will have a significant impact on industrial production, housing, continued education, and healthcare.

Associated with the aging population is increased prevalence of injuries, disorders, and diseases. Furthermore, across the age spectrum, health trends indicate significant increases in life-long conditions including diabetes, autism, obesity, and cancer. The American Cancer Society estimates that 1,685,210 new cancer cases (excluding the most common forms of skin cancer) will be identified and we will see 595,690 cancer deaths in the US during 2016. Furthermore, the probability of developing invasive cancers increases significantly with age [ACS Cancer Facts and Figures 2016].

These trends are producing a growing need for personalized health care. For example, the current rate of new strokes is 800,000 per year, and that number is expected to double in the next two decades. Furthermore, while stroke used to affect patients in their 60s and older, its instance is growing in the population in their 40s and up. Stroke patients must engage in intensive rehabilitation to attempt to regain function and minimize permanent disability. However, there is already a shortage of suitable physical therapists, and the changing demographics indicate a yawning gap in care in the near future. While stroke is the most common cause of movement impairments in adults, Cerebral Palsy (CP) is in children; both persist in life-long disabilities. About 10,000 infants and children are diagnosed with CP each year and there are over 764,000 persons in the US manifest symptoms of CP. Further, the number of neurodevelopmental and cognitive disorders is on the rise, including autism spectrum disorder, attention deficit and hyperactivity disorder, and others. Autism rates alone have quadrupled in the last quarter century, with one in 88 children diagnosed with the deficit today (up from 1 in 150 just a few years ago). Improved outcomes from early screening and diagnosis and transparent monitoring and continual health assessment will lead to greater cost savings, as can effective intervention and therapy. These factors will also offset the shrinking size of the healthcare workforce, while affordable and accessible technology will facilitate wellness, personalized, and home-based health care.

Increasing life-long independence thus becomes a key societal driver. It includes improving the ability to age in place (i.e., to enable the elderly to live at home longer, happier and healthier), improving mobility, reducing isolation and depression at all ages (which in turn impacts productivity, health costs, and family well-being). Robotics autonomy holds the potential to make existing

assistive machines, like powered wheelchairs and robotic arms, easier to operate and perhaps even accessible to populations whose severe motor impairments currently inhibit accessibility. The operation of such assistive robots can enhance and enable human independence. Improving care and empowering the care recipient also facilitates independence for caregivers, who have shifted from female stay-at-home relatives and spouses to employed family members of both genders, because the economics of in-home health care are unaffordable. Robotics technologies can improve safety and monitoring to avoid mis-medication, ensure consistency in taking medication, monitoring for falls, lack of activity, and other signs of decline.

All of the above features and properties of robotics technologies have the potential to prolong and improve productivity of the workforce and increase its size. With the decrease in available social security and retirement funding, people are working longer. Enabling people with disabilities, whose numbers are on the rise, to go into the workforce (and contribute to social security) would also offset the current reduction in available labor/workforce.

Finally, keeping technology leadership in the broad domain of health care is a key goal, given the size of the US population and its age demographics.

4.2. Aging and Quality of Life Improvement

4.2.1. Motivating Scenario

Due to the baby boom, the corresponding successive baby bust, and longer life expectancies, the average age of the national population is growing older. In order to maximize quality of life and to minimize the cost of care, a commonly agreed upon paradigm is to promote aging in place, where older adults can live at home provided that they receive some medical or service-related supplemental care. In this case, a person might need some help with light housework, decision support related to medical matters (medication management, nutrition, exercise regimen, etc.), or a conduit for improved social contact with the outside world. A robot can be utilized in these scenarios to perform service tasks, help a person maintain compliance with physician directives while also enjoying independent flexibility, and to either act as a social mediator or conduit for telepresence in order to increase the amount of human-human interaction that person receives in their regular routine.

4.2.2. State-of-the-art

Robots that provide social and cognitive support are beginning to appear in therapy, health, and wellness applications. The current state-of-the-art is able to achieve some of these critical tasks in laboratory or simulated settings. These socially assistive robots motivate their users to pursue healthy behaviors, engage them in a therapy program, and provide an easy-to-use natural interface. Such robots will recognize and display a wide range of human communicative cues such as speech, gesture, and gaze, and will create appropriate behavioral responses to offer effective and rich social interaction. They will be able to employ sophisticated models of embodied dialog that include verbal and nonverbal speech acts and handle imperfections that are typical in human communication.

4.2.3. Key Challenges and Unmet needs

Research challenges in achieving socially assistive robots for older adults include developing models of human behavior that accurately capture the nuanced and complex patterns of social interactions. Based on samples drawn from human experts, these models will allow robots to perform various roles such as consultant, therapist, buddy, and caregiver, and employ different strategies such as expressing authority or compassion or motivational competition, all to achieve desired behavior change in their users. Robots will also have to adapt to complex participation structures that are typical in day-to-day social interactions for situations such as group exercise and therapy, understanding the roles of all participants, following changes in speakership, appropriately acknowledging speakers, addresses, bystanders, and so on.

A research goal that is particularly pertinent to older adults is the ability to build and maintain relationships over long periods of time. Robots will need the capability to not only achieve short-term interactions, but also maintain these interactions over weeks and months, adapting their behavior to changes in the user's state of health, in responsiveness to different behavioral strategies, and in the relationship that has been established between the robot and its user. These changes will draw on health data as well as data on how human relationships change over time and employ learning strategies. Research on these capabilities will explore how much autonomy robots have in their interactions with their users; the robot might serve as an interface between a patient and a therapist or serve as a therapist itself.

The development of core capabilities for effective social human-robot interaction and interfaces must follow a human-centered design process and rigorous assessments with a range of stakeholders. User research in this process might involve the targeted health population in the early design process as well as formative evaluations of the design iterations that extend to patient, physician, family, therapist, and other members of their community. A key methodological limitation that research and development in this area will need to explore is identifying appropriate measures of success for natural interaction, validating these metrics across contexts and health applications, and developing methods to capture these measurements in real time as input to the robot for online assessment of the interaction and for learning.

Enabling cost-efficient and effective robot-mediated health or social communication requires the robotics research community address a number of challenges. Existing telepresence robots (e.g., InTouch, VGo) provide only visual and voice communication. Manipulation capabilities will enable the next level of physical interaction required to diagnose, treat, and even comfort patients or to engage extra-conversational interaction with a person remotely (i.e., playing a game of chess). Thus, any advances toward robots operating autonomously in human environments (e.g., navigation and dexterous manipulation) will also impact telepresence robots.

With technical advances in autonomous robots, telepresence robots could operate in semi-autonomous or autonomous modes. For example, these systems might allow the remote operator to issue only high-level commands for dexterous manipulation or navigation tasks or serve as personal health-monitoring devices when not actively in use as remote-presence portals, collecting information and providing medication reminders. It will be important to understand how a robot can effectively use a variety of levels of autonomy and still provide a seamless intuitive and acceptable experience for the end user. The design of adjustable autonomy interfaces for this type of scenario is an open challenge. The robot must communicate to the end users whether it is being

operated by a remote human or autonomously or some combination of the two. The interface available to the user might vary across scenarios and across different levels of autonomy.

4.2.4. Vision: 5-10-15 years

- In 5 years, robots will autonomously maintain one-time (e.g., a health interview) or short-term (e.g., a specific exercise) interactions, in specific, narrowly-defined domains, following appropriate norms of human social embodied communication, including social distance, gesture, expressions and other non-verbal cues as well as simple verbal content, instructions, and feedback.
- In 10 years, robots will autonomously maintain longer, repeated interactions in a broader set of domains, in controlled environments. They will offer a combination of human-led and robot-led interactions using open dialog including speech, gesture, and gaze behaviors in limited domains. They will be capable of providing prescribed intervention/therapy within precisely specified domains.
- In 15 years, robots will autonomously maintain multiple interactions over weeks and months in a broad set of domains. These robots will offer complex mixed-initiative interactions and fluently use multi-modal models of behavior that are generalizable across a broad set of social situations. They will adapt their behaviors to changes over time, including small fluctuations in mood, slow decline or improvement, and sudden unexpected changes, and shape the interaction to match the role and need of individual users.

4.3. Surgical and Interventional Robotics

4.3.1. Motivating Scenario

A pre-operative diagnostic test indicates that a patient may have cancer in an internal organ. The patient receives a Magnetic Resonance Imaging (MRI) scan, from which the existence and location of cancerous tissue is confirmed. The case is referred to a surgeon, who reviews digital models of the patient's anatomy based on the pre-operative images. An automated planning system uses these images as well as local and national surgical databases to guide the surgeon toward the most appropriate approach to the surgery. On the day before the procedure, the surgeon rehearses the surgery several times using a patient-specific simulation, visualizes the spatial extent of the cancer and develops an optimal surgical plan. On surgery day, a miniature robotic instrument is introduced into the patient's body through a very small incision. An imaging and navigation system guides the surgeon through the surgery and provides her with three-dimensional views of the anatomy, with cancerous tumors clearly highlighted. The system gives her the sense that she is inside of the patient's body and is able to see and feel the tissue while delicately removing all traces of the cancer. During the surgery, the navigation system tracks progress and automatically provides an optimal view of the anatomy as the surgeon works – acting as a digital assistant. The end result: the cancerous tissue is removed with very little impact on surrounding healthy tissue and the patient can leave the hospital on the same day with little pain and scarring, with the burden of cancer lifted from the patient's mind.

The development of surgical robots is motivated by the desire to:

- enhance the effectiveness of a procedure by coupling information to action in the operating room or interventional suite, and
- transcend human physical limitations in performing surgery and other interventional procedures, while still affording human control over the procedure.

Two decades after the first reported robotic surgical procedure, surgical robots for some procedures. Surgical robots are beginning to realize their potential in terms of improved accuracy and visualization, as well as enabling of new procedures.

4.3.2. State-of-the-art

Current robots used in surgery are under the direct control of a surgeon, often in a teleoperation scenario in which a human operator manipulates a master input device and a patient-side robot follows the input. In contrast to traditional minimally invasive surgery, robots allow the surgeon to maintain hand dexterity inside the body, scale down operator motions from normal human dimensions to very small distances, and provide a very intuitive connection between the operator and the instrument tips. The surgeon can cut, cauterize, and suture with a precision equal to or better than that previously available during only very invasive open surgery. A complete surgical workstation contains both robotic devices and real-time imaging devices to visualize the operative field during surgery. The next generation of surgical workstations will provide a wide variety of computer and physical enhancements, such as "no-fly" zones around delicate anatomical structures, seamless displays that can place vast amounts of relevant data in surgeon's field of view, and recognition of surgical motions and patient state to evaluate performance and predict health outcomes.

If the right information is available, many medical procedures can be planned ahead of time and executed in a reasonably predictable manner, with the human exercising mainly supervisory control over the robot. By analogy to industrial manufacturing systems, this model is often referred to as "Surgical CAD/CAM" (Computer-Aided Design and Computer-Aided Manufacturing). Examples include preparation of bone for joint reconstructions in orthopedic surgery and placement of needles into targets in interventional radiology. In these cases, the level of "automation" may vary, depending on the task and the relative advantage to be gained. For example, although a robot is easily able to insert a needle into a patient, it is currently more common for the robot to position a needle guide and for the interventional radiologist to push the needle through the guide. As imaging, tissue modeling and characterization, and needle steering/alignment technology improve, future systems are likely to become more highly integrated and actively place needles and therapy devices through paths that cannot be achieved by simply aiming a needle guide. In these cases, the human will identify the target, plan or approve the proposed path, and supervise the robot as it steers the needle to the target.

4.3.3. Vision: 5-10-15 years

A summary of current State-of-the-Art, key challenges and capabilities roadmap is listed below.

Surgical Scenario	Current State-of-the-Art	Key Challenges Unmet Needs	Road Map		
			5 Years	10 Years	15 Years
Surgery and Intervention	<ul style="list-style-type: none"> · Limited interfaces for physical interaction human and robot · Current Robots used in surgery are under the direct control of a surgeon 	<ul style="list-style-type: none"> · Modeling human behavior and dynamics · Sensing the human's physical behavior in a very large number of dimensions · Highly dexterous surgical robots that can move through body lumens and along tissue planes so as to minimize collateral tissue damage · Identify proximity and relative orientation between an end effector and deformable structure/tissues · Intuitive interfaces for physical interaction humans and robots · Tissue modeling and characterization · Endoscope steering and target alignment control · Maintain real-time registration in 3D space between instruments and deformable or non-deformable tissue 	<ul style="list-style-type: none"> · New devices and algorithms to enable more effective two-way exchange of information and energy between the human and the robot · Control interfaces and navigation systems that integrate real-time sensor and database information · Full suite of physical feedback to the surgeon as they control the robotic instruments as well as environmental compliance of the remote patient's tissue · Developing robot behaviors that will ensure appropriate interaction no matter what the human does · Uncertainty management 	<ul style="list-style-type: none"> · Intuitive and transparent human-robot interaction · Interface will estimate user's intent, rather than simply executing the user's commands that may be subject to human imperfections 	<ul style="list-style-type: none"> · Sensing a human's movement and inferring intent, · Algorithms development to provide context-appropriate forces to a human operator

4.4. Rehabilitation

4.4.1 Motivating Scenario

A patient is prescribed a physical therapy by a physician. This physical therapy may be used to treat a physical ailment, such as a muscular injury, recovery from a surgical procedure, or such a regimen typically involves both at-home and in-person exercises to complete. The in-person sessions are monitored by a physical therapist and shaped to ensure compliance with a given therapeutic regimen. Compliance with the therapy's instructions is highly correlated with positive outcomes and patient satisfaction, but is often not achieved to a therapeutic level. Compliance depends on two factors: motivation to do the often difficult exercises required for rehabilitation; and correct execution of the exercises required. Neither of these tasks explicitly requires physical interaction. Compliance can be monitored by in-person contact, but insurance typically does not reimburse for enough in-person contact to ensure that compliance is sufficient. Robotic therapeutic aids for rehabilitation could be applied to address the therapeutic gaps and to help ensure compliance.

4.4.1. Robotic Replacement of Diminished/Lost Function

Orthotic and prosthetic devices are worn to increase functionality or comfort by physically assisting a limb with limited movement or control, or by replacing a lost or amputated limb. Such devices are increasingly incorporating robotic features and neural integration. Orthoses protect, support, or improve the function of various parts of the body, usually the ankle, foot, knee and spine. Unlike

robotic devices, traditional orthoses are tuned by experts and cannot automatically modify the level or type of assistance as the patient grows and his or her capabilities change. Robotic orthoses are typically designed in the form of an exoskeleton, which envelopes the body part in question. They must allow free motion of limbs while providing the required support. Most existing robotic exoskeletons are research devices that focus on military applications (e.g., to allow soldiers to carry very heavy load on their backs while running) and rehabilitation in the clinic. These systems are not yet inexpensive and reliable enough for use as orthoses by patients.

A prosthesis is an artificial extension that replaces the functionality of a body part (typically lost by injury or congenital defect) by fusing mechanical devices with human muscle, skeleton, and nervous systems. Existing commercial prosthetic devices are very limited in capability (typically allowing only opening/closing of a gripper) because they are signaled to move purely mechanically or by electromyography (EMG), which is the recording of muscle electrical activity in an intact part of the body). Robotic prosthetic devices aim to more fully emulate the missing limb or other body part through replication of many joints and limb segments (such as the 22 degrees of freedom of the human hand) and seamless neural integration that provides intuitive control of the limb as well as touch feedback to the wearer. The last few years have seen great strides in fundamental technologies and neuroscience that will lead to these advanced prostheses. Further robotics research is needed to vastly improve the functionality and lower the costs of prostheses.

4.4.2. Robot-Assisted Recovery and Rehabilitation

Patients suffering from neuromuscular injuries or diseases, such as occur in the aftereffects of stroke, benefit from neurorehabilitation. This process exploits the use-dependent plasticity of the human neuromuscular system, in which use alters the properties of neurons and muscles, including the pattern of their connectivity, and thus their function. Sensory-motor therapy, in which a human therapist and/or robot physically assists (or resists) a patient during upper or lower extremity movements helps people re-learn how to move. This process is time-consuming and labor-intensive, but pays large dividends in terms of patient health care costs and return to productive labor. As an alternative to human-only therapy, a robot has several key advantages for intervention:

- after set up, the robot can provide consistent, lengthy, and personalized therapy without tiring;
- using sensors, the robot can acquire data to provide an objective quantification of recovery; and
- the robot can implement therapy exercises not possible by a human therapist.

There are already significant clinical results from the use of robots to retrain upper and lower-limb movement abilities for individuals who have had neurological injury, such as cerebral stroke. These rehabilitation robots provide many different forms of mechanical input, such as assisting, resisting, perturbing, and stretching, based on the patient's real-time response. For example, the commercially available MIT-Manus rehabilitation robot showed improved recovery of both acute and chronic stroke patients. Another exciting implication of sensory-motor therapy with robots is that they can help neuroscientists improve their general understanding brain function. Through knowledge of robot-based perturbations to the patient and quantification of the response of patients with damage to particular areas of the brain, robots can make unprecedented stimulus-response recordings. In order to optimize automated rehabilitation therapies, robots and experiments need to be developed to elucidate the relationship between external mechanical forces and neural plasticity. The understanding of these relationships will also give neuroscientists and neurologists insight into brain function, contributing to basic research in those fields.

In addition to providing mechanical/physical assistance in rehabilitation, robots can also provide personalized motivation and coaching. Socially assistive robotics (SAR) focuses on using sensory data from wearable sensors, cameras, or other means of perceiving the user's activity in order to provide the robot with information about the user that allows the machine to appropriately encourage, motivate and coach sustained recovery exercises. Early work has already demonstrated such socially assistive robots in the stroke rehabilitation domain, and they are being developed for other neuro-rehabilitation domains including traumatic brain injury frequently suffered by recent war veterans and those involved in serious traffic accidents. In addition to long-term rehabilitation, such systems also have the potential to impact health outcomes in short-term convalescence, where intensive regimens are often prescribed. For example, an early system was demonstrated in the cardiac ward, encouraging and coaching patients to perform spirometry exercises ten times per hour to prevent infection and speed healing. Such systems can serve both as force multipliers in health care delivery, providing more care to more patients, and also as a means of delivering personalized, customized care to all patients.

4.4.3. Behavioral Therapy

Convalescence, rehabilitation, and management of life-long cognitive, social, and physical disorders requires ongoing behavioral therapy, consisting of physical and/or cognitive exercises that must be sustained at the appropriate frequency and correctness. In all cases, the intensity of practice and self-efficacy have been shown to be the keys to recovery and minimization of disability. However, because of the fast-growing demographic trends of many of the affected populations (e.g., autism, ADHD, stroke, TBI, etc., as discussed in Section 1.2), the available health care needed to provide supervision and coaching for such behavior therapy is already lacking and on a recognized steady decline.

SAR is a comparatively new field of robotics that focuses on developing robots aimed at addressing precisely this growing need. SAR is developing systems capable of assisting users through social rather than physical interaction. The robot's physical embodiment is at the heart of SAR's assistive effectiveness, as it leverages the inherently human tendency to engage with lifelike (but not necessarily human-like or animal-like) social behavior. People readily ascribe intention, personality, and emotion to even the simplest robots, from LEGO toys to iRobot Roomba vacuum cleaners. SAR uses this engagement toward the development of socially interactive systems capable of monitoring, motivating, encouraging, and sustaining user activities and improving human performance. SAR thus has the potential to enhance the quality of life for large populations of users, including older adults, individuals with cognitive impairments, those rehabilitating from stroke and other neuromotor disabilities, and children with socio-developmental disorders such as autism. Robots, then, can help to improve the function of a wide variety of people, and can do so not just functionally but also socially, by embracing and augmenting the emotional connection between human and robot.

Human-Robot Interaction (HRI) for SAR is a growing research area at the intersection of engineering, health sciences, psychology, social science, and cognitive science. An effective socially assistive robot must understand and interact with its environment, exhibit social behavior, focus its attention and communication on the user, sustain engagement with the user, and achieve specific assistive goals. The robot can do all of this through social rather than physical interaction, and in a way that is safe, ethical and effective for the potentially vulnerable user. Socially assistive robots have

been shown to have promise as therapeutic tool for children, older adults, stroke patients, and other special-needs populations requiring personalized care.

4.4.4. Key Challenges and Unmet needs

Rehabilitation robots need to understand their user's state and behavior to respond appropriately. Because human state and behavior are complex and unpredictable, and because vision-based perception is an ongoing challenge in robotics (and a privacy concern as well), automated perception and understanding of human state and behavior requires the integration of data from a multitude of sensors, including those on the robot, in the environment, and worn by the user, and application of statistical methods for user modeling based on this multi-modal data. Fundamental mechanistic models of how robot interaction affects state and behavior are still in their infancy; further developing these models will enable more effective design of the control algorithms for medical and health care robots.

The ability to automatically recognize emotional states of users in support of appropriate, robot behavior is critical for making personalized robotics effective, especially for health-related applications that involve vulnerable users. Emotion understanding requires processing multi-channel data from the user, including voice, facial expression, body motion, and physiologic data and reconciling inconsistencies (e.g., between verbal and facial signals). The power of empathy is well recognized in health care: doctors who are perceived as empathetic are judged as most competent and have the fewest lawsuits. Further, creating empathy in synthetic systems is just one of the challenges of perceiving and expressing emotion. Early work in socially assistive robotics has already demonstrated that personality expression, related to emotion, is a powerful tool for coaching and promoting desired behavior from a user of a rehabilitation system.

Physiologic data sensors are typically wearable sensors and devices that provide real-time physiologic signals (e.g., heart rate, galvanic skin response, body temperature, etc.). Active research is addressing methods for extracting metrics, such as frustration and motivation, from physiologic data. The ability to capture physiologic data without encumbering a patient and to transmit those data to a computer, robot, or caregiver, has great potential for improving health assessment, diagnosis, treatment, and personalized medicine. It will enable intelligent assistance, appropriate motivation, and better performance and learning.

In rehabilitation, human-robot collaboration can ensure productive interactions between a client and a therapist, with increased efficiency and improved quality of care. The robot must be able to shift from assistant (providing support and information on client performance) to director in driving the therapeutic exercises when alone with the client. Implementing such guidance requires that the robot understands the task the therapist is trying to accomplish and the current state of both client and therapist, and that it has the physical and/or social means for providing assistance.

4.5. Clinical Workforce Support

4.5.1. Motivating Scenario

There are far more people needing healthcare than there are clinicians to provide it. Clinicians across all aspects of care are overworked and overloaded, constantly, particularly in primary and acute care. Healthcare workers have the most hazardous industrial job in America with regards to nonfatal

occupational injuries and illness, according to the National Institute for Occupational Health and Safety (NIOSH). They are also at the highest risk of suicide and substance abuse compared to any other profession, according to NIOSH.

A key way robots can aid the healthcare workforce is through reducing both cognitive and physical workload for clinicians. This support can range from robots that deliver supplies from one part of a hospital to another, to intelligent systems that can provide logistical support, to robots that help clinicians lift patients. Autonomous vehicles might provide patient or clinical transportation, or mobile manipulators might help clean up highly infectious waste. There are many places where robotics technology can be hugely impactful to the clinical workforce.

Robots are also used extensively in medical education. Over 180,000 doctors, nurses, EMTs, and other first responders train annually on high fidelity robotic patient simulators, which are lifesized mannequins that can breathe, bleed, respond to medication, and interact with learners. These simulators are used extensively to help clinicians practice procedural and communication skills before treating actual patients, thereby reducing risking lives. Furthermore, surgeons of all disciplines utilize task trainers to hone their motor skills and learn new procedures, which are lifelike models of anatomical regions of interest, some of which contain actuated elements.

Finally, telerobots may be used to help clinicians reach patients in rural areas, bring remote expertise to consult on a case, or protect clinicians from harm when treating highly infectious diseases like Ebola. This technology has proven effective in tele-dermatology and tele-psychiatry, though the evidence is still lacking for tele-surgery and other manipulative tasks.

4.5.2. State-of-the-art

In terms of reducing physical workload for clinicians, straightforward object transportation tasks - supply delivery, waste removal, etc. - are solvable with existing technology, and are being used in many hospitals. Robots are also being tested to assist clinicians with patient transfer (e.g., bed to chair).

In terms of cognitive support, there have been recent advancements in providing embodied cognitive aids to clinicians in hospitals, to help with scheduling and logistics, which are promising. This is another huge sunk cost for clinicians, but a tractable problem in the planning space which could make a huge difference.

Finally, telepresence robots have advanced in terms of their on-board sensing capability, and are starting to also have low-cost manipulators on them which will give them a greater degree of flexibility in terms of telemanipulation. Design advancements have helped enable remote users of these robots achieve a greater level of situational awareness in the telehealth application space, which can help address a range of important interaction challenges.

4.5.3. Key Challenges and Unmet needs

The biggest challenge in this domain is that robots in healthcare settings are an inherently disruptive technology, which will alter the clinical workflow in unforeseen ways. This does not mean the disruption will necessarily be problematic, but it will change things. Furthermore, care is very

localized - every care setting is completely different. This further exacerbates both the social and technological challenges to deploying robots within health systems. Extensive research in this space is necessary to deeply, and longitudinally explore this context.

Another challenge is the workforce itself. Clinicians tend to have fairly low technology literacy levels. Poorly designed technology coupled with poor socio-technical integration is a huge challenge. This was seen with the advent of electronic health records (EHRs), where an untold number of lives and dollars were lost with their too-rapid deployment. It is critical in healthcare robotics to proceed cautiously.

Thus, it is important technologists and those marketing healthcare robotics technology to the clinical workforce are fully transparent about the true capabilities and limitations of a system, and provide extensive inter-professional training to ensure workers are able to safely work with robots. There have been several instances of companies grossly exaggerating a robot's capabilities, understating the required amount of training time clinicians need to become proficient at using a robot, and substantially overcharging rural hospitals for technology that does not actually benefit healthcare workers or patients. Some of these deceptions have led to patient harm.

Another challenge in building robots to provide cognitive support to the clinical workforce is that they have very busy, chaotic, stressful jobs, which are in highly dynamic environments. It is still difficult to know when and how to intervene, particularly in critical care settings. This brings up a lot of unforeseen technical challenges in robotics which we are only just beginning to realize as a discipline.

4.5.4 Vision: 5-10-15 years

- In 5 years, robots will autonomously support clinicians by completing non-value added, well-defined tasks, such as: delivering supplies, removing waste, and administering medication. Telepresence robots will be used for tele-psychiatry, tele-dermatology, and tele-wellness promotion across a range of care settings including rural health, home health, and the Indian health service.
- In 10 years, robots will support clinicians in dangerous manipulation tasks, including: patient transfer (bed to chair), patient mobility, and highly infectious disease care (e.g., ebola). Expressive, interactive "smart" robotic patient simulators will provide high-fidelity experiences to learners for both new clinician training as well as retraining.
- In 15 years, robots will be as seamlessly integrated into the clinical workflow as mobile phones. They will be easy for providers with low technology literacy levels to use. They will provide cognitive support to clinicians with administrative care tasks, such as scheduling exams and bringing patients to them or arranging patient visit schedules for clinicians. They will learn new interactive paradigms on the fly to be unobtrusive to clinical work.

4.6. Deployment Issues

Deployment of complete health robotics systems requires practical issues of safe, reliable, and continuous operation in human environments. The systems must be private and secure, and interoperable with other systems in the home. To move from incremental progress to system-level implications, the field of medical and health robotics, needs new principled measurement tools and methods for efficient demonstration, evaluation, and certification.

The challenge of system evaluation is compounded by the nature of the problem: evaluating human function and behavior as part of the system itself. Quantitative characterization of pathology is an existing problem in medicine; robotics has the potential to contribute to solving this problem by enabling methods for the collection and analysis of quantitative data about human function and behavior. At the same time, some health care delivery is inherently qualitative in nature, having to do with therapy, motivation, and social interaction; while such methods are standard in the social sciences, they are not recognized or accepted by the medical community. Because medical and health robotics must work with both trained specialists and lay users, it is necessary to gain acceptance from both communities. This necessitates reproducibility of experiments, standards, code re-use, hardware platform re-use/sharing, clinical trials, sufficient data for claims of efficacy, and moving robots from lab to real world. As systems become increasingly intelligent and autonomous, it is necessary to develop methods for measuring and evaluating adaptive technologies that change along with the interaction with the user.

Affordability of robotic technology must be addressed at several different levels. The hospital pays a significant cost in terms of capital investment to acquire a robot, the maintenance costs are high, and the cost of developing robots is immense, given their complexity and stringent performance requirements for medical applications. Policies are needed to address regulatory barriers, the issue of licensure and state-by-state certification, rules for proctoring and teaching with robots, and reimbursement via insurance companies. Finally, we need to consider the culture of both surgeons and patients; both groups must have faith robotic technology for widespread acceptance.

The ultimate goal of robotics for health and wellness is for a consumer to be able to go to a store and purchase an appropriate system, much like one buys a computer today, and then integrate that system into the home without requiring retrofitting. The technology must be shown to be effective, affordable, and accepted.

To create a robust and vibrant healthcare robotics industry to meet this goal, first resources must be directed toward funding collaborative ventures that bring together the necessary expertise in engineering, health, and business. Funding is specifically needed in the areas of incubating and producing complete systems and evaluating those on patient populations in trials that are a year long or longer. Currently no funding agency exists for such incubation: the research is too technological for NIH, too medical for NSF, and too far removed from an immediate market to be funded by business or venture capital. As a result, there is a lack of critical mass of new, tested and deployed technological innovations, products and businesses to create an industry.

A thriving industry requires training in research, implementation, evaluation, and deployment of healthcare robotics. Universities are already taking the first step to facilitate this by developing interdisciplinary programs that bridge medical and engineering training at the undergraduate and

graduate levels. There is also increased attention to K-12 outreach, using the already popular and appealing topic of robotics. Health-related robotics in particular effectively recruits girls into engineering, addressing another important workforce trend, since women play a key role in both healthcare and informal caregiving. As the use of assistive robots becomes more widespread in the society, their adoptability will significantly rely on the personalization of the control and interaction technologies in an automated and continuous manner, as every disability is unique and one solution does not fit all. User-centered and adaptable designs of personalization at three levels, interfaces for interaction, controllers for action, and feedback in interaction, will result in greater acceptability of the assistive technologies such as prosthetics, exoskeletons, and smart wheelchairs

4.7. Contributors

The 2016 update of the roadmap was edited by:

Brenna Argall Northwestern University	Maja Mataric Univ. of Southern California	Michael Yip UC San Diego
Nahid Sidki SRI Inc.	John Hollerbach Univ. of Utah	Taskin Padir Northeastern University
Jake Abbott Univ. of Utah	Jaydev Desai Georgia Tech	Laurel Riek UC San Diego
David Feil-Seifer University of Nevada, Reno	Ayanna Howard Georgia Tech	

The 2013 roadmap upon which the document is based was prepared by the following group of researchers (the affiliations have not been updated):

Jake Abbott Univ. of Utah	How Choset Carnegie Mellon University	Henrik I Christensen Georgia Institute of Tech
Mihir Desai Univ. of Southern California	Simon DeMaio Intuitive Surgical Inc	Pierre Dupont Harvard Children’s Hospital
Rod Grupen Univ. of Mass.	Gregory Hager Johns Hopkins University	Blake Hannaford University of Washington
John Hollerbach University of Utah	Charlie Kemp Georgia Institute of Tech.	Venkat Krovi SUNY-Buffalo
Katherine Kuchenbecker Univ. of Pennsylvania	Art Kuo University of Michigan	Gerald Loeb Univ. of Southern California
Maja Mataric Univ. of Southern California	Bilge Mutlu Univ. of Wisconsin – Madison	Marcia O’Malley Rice University

Allison Okamura
Stanford University

David Rinkensmeier
UC Irvine

Brian Scassellati
Yale University

Nabil Simaan
Vanderbilt University

Reid Simmons
Carnegie Mellon University

Andrea Thomaz
Georgia Institute of Tech.

Kate Tsui
Univ. of Mass – Lowell

Michael Zinn
Univ. of Wisconsin - Madison

5. Enhancing Public Safety

5.1 Strategic Importance of Robotics in Enhancing Public Safety

Protecting America encompasses those who seek to protect the United States and its citizens, including federal government military, law and safety services as well as state and local public safety services, including law enforcement and first response. These organizations and individuals are responsible for protecting and deterring incidents, mitigating injuries and loss of life to themselves and civilians, and minimizing property damage. Robotics can help military and emergency service personnel achieve these goals.

The Department of Defense (DoD) and the defense industry generally refer to robotic systems as **unmanned systems**, a term that applies to all forms of military, border patrol, homeland security and emergency response robots that keep humans out of harm's way. Unmanned systems offer tremendous versatility, persistent functionality, the capacity to reduce the risk to human life, and an ability to provide contributing functionality across all key warfighting areas. These systems provide the U.S. military, federal agencies and local emergency services officials with an increasingly valuable means for conducting a wide range of operations in the chaos and uncertainty of de-engineered environments. Military operations, such as precision targeting and precision strike are conducive to using unmanned systems, as are missions and applications shared with public safety services, such as chemical, biological, radiological, and nuclear (CBRN) detection, counter-improvised explosive device (C-IED) actions, and humanitarian assistance. While the largest number of deployed unmanned systems appear in the military, providing a great deal of firsthand evidence of how effective unmanned systems can be, they are also highly prevalent in civilian bomb squads and are beginning to be utilized in other emergency services environments. The collective experience, together with the recognition that the capabilities provided by unmanned systems will continue to expand, serves to raise expectations for the growing role of unmanned systems in many unstructured situations. The recent use of robots during the Fukushima Daiichi incident (3-11) underscored both the successes and challenges for civilian use of robotics in disaster scenarios.

5.1.1 The Future Landscape

The strategic environment and the resulting national security challenges facing the United States for the next 25 years are diverse, encompassing outside our national borders (OCONUS) threats from the saber rattling of North Korea and Kim Jong Un to the unpredictability of militant organizations, such as the Islamic State of Iraq and the Levant (ISIL) and the Taliban, as well as threats within our national borders (CONUS), spanning the recent budget sequestration deal to homegrown militant groups and terrorism. As a result, the United States faces a complex and uncertain security landscape that challenges both finances and capabilities. The rise of new OCONUS powers, the growing influence of non-state actors, the spread of weapons of mass destruction and other irregular threats, and continuing socioeconomic unrest will continue to pose profound challenges to international order. While these OCONUS threats continue to develop, fears of terroristic migration to the homeland are very present. Our Emergency Services are challenged with the concerns for day to day operations, while being ever mindful of nefarious actions and intentful manmade disaster. Our first responder's have been introduced to active shooter events, homegrown IED's and explosives, intentionally set structure fires, ambushes, etc. More so than

ever, our public safety personnel must remain on high alert, be vigilant, and prepared for worst case scenarios.

Over the next two decades, U.S. military forces will operate in a geostrategic environment of considerable uncertainty with traditional categories of conflict becoming increasingly blurred. This era will be characterized by protracted confrontation among state, non-state, and individual actors using violent and nonviolent means to achieve their political and ideological goals. Likewise, homeland security agencies face considerable uncertainty as climate change and natural fluctuations in weather patterns, coupled with more Americans living in harm's way of natural and manmade disasters and attacks, put an unprecedented number of civilians at risk of experiencing increasingly violent environments. Changes in workforce demographics also challenge the technical competence of federal agencies, demanding innovative responses and careful compromises. Clearly, there is much that robotics and autonomy can offer in these stressful times.

The DoD has responded with a call for greater autonomy and a concerted, agency-wide discussion of how autonomy research can and must respond.

Imagine if...

- *We could covertly deploy networks of smart mines and UUVs to blockade and deny the sea surface, differentiating between fishing vessels and fighting ships...*
...and not put U.S. Service personnel or high-value assets at risk.
- *We had an autonomous system to control rapid-fire exchange of cyber weapons and defenses, including the real-time discovery and exploitation of never-seen-before zero day exploits...*
...enabling us to operate inside the "turning radius" of our adversaries.
- *We had large numbers of small autonomous systems that could covertly enter and persist in denied areas to collect information or disrupt enemy operations...*
...a "sleeper presence" on call.
- *We had large numbers of low-cost autonomous unmanned aircraft capable of adaptively jamming and disrupting enemy PNT capabilities...*
...destroying their ability to coordinate operations.
- *We had autonomous high performance computing engines capable of not only searching "big data" for indicators of WMD proliferation, but of deciding what databases to search...*
...to provide early warning and enable action

And imagine if we are unprepared to counter such capabilities in the hands of our adversaries.

- *Defense Science Board Summer Study on Autonomy, June 2016.*

Civilian homeland security agencies have been less proactive in defining and preparing for these coming threats, but many in the emergency response community are planning for leaner, more agile times ahead. Our nation's Fire Service and Law Enforcement communities must both become supported and educated in the use of unmanned systems. This technology can help fill gaps in first response, while expanding operational presence and extending capabilities that address threats in an efficient timely manner.

5.1.2 The Role of Unmanned Systems in Enhancing Public Safety

Unmanned systems enhance public safety through military uses of force projection primarily outside our national borders, while emergency services can utilize unmanned systems for unforeseen events within our borders. Military operational functional areas include engagement, sustainment, mobility,

and survivability/force protection, while emergency services operational areas include search and rescue, real-time data gathering to inform decision making, mapping, damage assessments, asset tracking, and during large scale events, creation of ad-hoc communication networks, and enhancement of deliverable payloads. Unmanned systems can also help reduce the load on personnel and mitigate the risks to the personnel responsible for these areas by providing early warning and information as well as increase stand-off from hazardous areas. The intent is to exploit the inherent advantages of unmanned systems, including their persistence, size, speed, maneuverability, and better sensing capabilities. As the technology continues to advance, the DoD and public safety services envision robots seamlessly operating with manned systems to aid in human decision making, while reducing the degree of required human control.

The nation's military and emergency services communities understand the effect that innovation and technology in unmanned systems can have on the future of our collective response to emerging threats in an ever-changing global and national environment. The DoD is committed to harnessing the potential of unmanned systems in its efforts to strengthen the nation's warfighting capability, while husbanding resources and maintaining fiscal responsibility. At the same time, public safety services use of unmanned systems against homeland threats is on the rise. Both communities believe unmanned systems must:

- Provide capabilities more efficiently through modularity, commonality and interoperability.
- Be more effective through greater autonomy, better performance, as well as more intuitive, natural and flexible command and control capabilities.
- Be more survivable with improved and resilient communications, development for anti-permissive environments, and more security from tampering.
- Be trustworthy, constructive and reliable partners that contribute to the overall human-robot team's success and performance.
- Take the "human" out of unmanned. Unmanned systems must strive to reduce the number of personnel required to operate and maintain the systems.
- Reduce the likelihood of casualties and harm to civilians and property.

The DoD is working to advance operational concepts that leverage the synergies between manned, (semi-)autonomous, and remote-controlled systems to achieve the capabilities and desired effects on missions and operations worldwide, while optimizing commonality and interoperability across space, air, ground, and maritime domains. Pursuing this approach with unmanned systems will help the DoD sustain its position as the dominant global military power and better enable national decision makers to adapt to an ever changing global and national environment. Similar efforts to develop operational concepts in the emergency services domain are in their infancy. For example, a working group is being stood up to address operational concepts for urban search and rescue, and the National Fire Protection Association and the National Institute of Standards and Technology are developing unmanned system equipment and training standards.

Unmanned systems technology developed to satisfy military needs is oftentimes well-suited for "dual use" in first response applications. In fact, many of the robotics applications and uses described in other sections of this road-map are based on technologies resulting from Government funded research projects intended to advance the state of the art and practice to meet defense requirements. The considerable investment being made by DoD in developing and maturing such technologies can also contribute to the development of unmanned systems products and applications for the emergency services sector through what is referred to as safety, security, and rescue robotics.

The emergency services are a great place to establish operational relevance of unmanned systems. While at first glance it can appear easy to integrate DoD and commercial off-the-shelf unmanned systems, the process can quickly become complicated and is not without its challenges. Today public safety must first establish favorable public perception, create concise operational policies and procedures, obtain funding streams adequate for both program development and equipment purchases, educate Incident Commanders to the benefits and limitations of unmanned systems, and become users of technology appropriate for emergency response. There are also considerations of staffing, creation of appropriate dispatch protocols, and while operating in the National Airspace, being able to comply with FAA regulations that were never created with the first responder in mind. That being said, federal, state, and local emergency service agencies already incorporate ground robots into bomb squads and use aerial vehicles for border security. Further, unmanned systems have already contributed to worldwide humanitarian response. Examples include, but are not limited to the UAS mapping in Nepal after the devastating earthquake, unmanned surface vehicles assisting in the water rescues of fleeing Middle Eastern refugees, and countless applications of UAS being deployed to assist with wildfires, flooding and other natural disasters throughout the United States. These applications are the tip of the iceberg, as law enforcement, fire rescue, disaster management, bridge inspection, and port officials begin to adopt these technologies.

Unmanned Aerial Systems

Recent technological advances and the new small UAS operation and certification rules (Part 107, RIN 2120–AJ60) issued by the Department of Transportation and the Federal Aviation Administration have made flying small (< 55 lbs) UAS popular with the general public and viable tools for the emergency services sector. The proliferation of inexpensive, small UAS has been very visible to the general public, even though the DoD fully embraced UAS capability in tactical, operational and strategic roles for the last 15 years. As the DoD has drawn down the number of U.S. troops deployed in support of Operation Enduring Freedom and Operation Iraqi Freedom and as targeted military interventions have increased, the use of UAS has expanded to support joint forces on the ground by providing surveillance and close in air support. The air domain has consumed a very large share of the overall DoD investment in unmanned systems, which has resulted in fielding many UAS capable of executing a wide range of missions. The original DoD UAS missions focused primarily on tactical reconnaissance; however, this scope has expanded to include a much wider range of capabilities. DoD UAS, for example, are playing a greater role in engagement missions, both as a designator and as the platform launching a munition. The ability of select UAS to conduct multiple strike missions-and time-critical targeting is well documented. As of March 2014, the Army had logged 2,000,000 flight hours on vehicles ranging from the persistent large Gray Eagle to the mini tactical Raven UAS. The current strategy maximizes the use of unmanned assets and minimizes troop deployments. Further, threats to the homeland will use these systems to monitor and protect the nation's civilians; thus, the number of fielded systems and applications will continue to expand, the number of flight hours is expected to dramatically increase.

The DoD's many successful OCONUS UAS deployment programs have left many in the emergency services hopeful that the technology will "trickle down" to their respective agencies. Until very recently, the primary use of this technology within CONUS has been the Predator UAS by the U.S. Customs and Border Patrol, while the Coast Guard intends to deploy the ScanEagle in 2017. While these military based systems have successes within the federal public safety sector, they have been less successful in the emergency services sector, particularly at the state and local levels. A number of factors, including the cost, required staffing, favorable public perception, regular training requirements, fragile technology that limits system reliability, complex command and control

systems, and the FAA's flight regulations have hindered UAS deployments. Small UAS (e.g., micro and mini) have the potential to transform the emergency services by providing real time intelligence gathering during and after kinetic events. An example of such a deployment is the use of Chinese based DJI Phantom 3 UAS to capture images and map the destruction after the massive earthquakes in Nepal. While these commercially based systems are small and inexpensive compared to their DoD counterparts, they are unable to operate in extreme environments and have very basic autonomy, imprecise GPS navigation, very limited sensor and actuator payload capacity, and restricted power supplies that limit deployment durations. Small UAS technology has great potential to enhance our nation's emergency response; however, two situations must occur. First, we must design the technology specifically for the first response domain, rather than the only options being off-the-shelf equipment or transitioned DoD equipment. Second, federal regulations must be tailored and designed to facilitate deploying UAS by the state and local service communities.

Unmanned Ground Systems

Unmanned Ground Vehicles (UGVs) have most frequently been used by military and public safety services to mitigate explosive devices. UGV systems, such as the Andros family of robots and the Packbot have been used to investigate and mitigate domestic explosive and hazardous materials threats as well as improvised explosive devices in the military theater. Hundreds of local bomb squads own these UGVs, which are sometimes used to support and protect emergency services personnel. While these UGVs have been used to breach and enter hostile buildings in order to provide intelligence and surveillance information to law enforcement, the use of UGVs by emergency services personnel has been limited, at best to these types of situations due to the expensive technology, cumbersome command and control, lack of autonomy, etc.

DoD UGVs support a diverse range of operations, including maneuver, maneuver support, and sustainment. Maneuver operations include closing with and destroying the enemy using movement and fires. Maneuver support missions include facilitating movement by mitigating natural and artificial obstacles and hazards. Sustainment missions include maintaining equipment, supplying the force with logistics and providing medical service and support. Since 2001 the DoD has acquired and deployed thousands of UGVs. Approximately 8,000 systems of various types saw action in Operation Enduring Freedom and Operation Iraqi Freedom. As of September 2010, UGVs have been used in over 125,000 missions, including suspected object identification and route clearance, to locate and defuse improvised explosive devices (IEDs). During these counter-IED missions, Army, Navy, and USMC explosive ordnance teams detected and defeated over 11,000 IEDs using UGVs.

The rapid fielding and proliferation of military UGVs have helped with many missions, but resulted in many challenges, not the least of which are configuration, sustainment and maintenance costs. UGVs will continue to provide tremendous benefit to the military commanders and have the potential to do the same for emergency service commanders. For example, fire department personnel are particularly interested using UGVs for fire based response, including structural collapse, confined space, and hazardous environments. In order to meet the challenges anticipated for using UGVs in the future, both the DoD and emergency services will require improvements in the UGV command and control interfaces, human-machine teaming, navigation and manipulation actuation, autonomy, reliability, endurance and survivability. These improvements need to keep pace with advances in 360° sensing, recording fidelity, and CBRNE detection and decontamination.

Unmanned Maritime Systems

Maritime domain awareness both abroad, as well as along and within our national borders is becoming increasingly critical to our national security. Counties along the national shoreline account for 39% of the total population, which increased by almost 40% between 1970 and 2010 and is projected to increase another 8% (10 million people) by 2020. Further, four of the five largest US metropolitan areas serve as port cities that handle a significant portion of the nation's imports and exports. Over 90% of information, people, goods and services flow across the world's oceans. Protecting the country's residents and economic prosperity is dependent on the ability to persistently monitor ocean surface and sub-surface activities, in order to identify, classify and mitigate emerging threats. Unmanned Maritime Systems will play a critical role in expanding the nation's undersea superiority and addressing growing challenges in piracy, natural resource disputes, drug trafficking and weapons proliferation.

The use of Unmanned Underwater Vehicles (UUVs) and Unmanned Surface Vehicles (USVs) has increased, largely because the enabling technology has reached an inflection point. Building upon decades of academic research in maritime vehicle autonomy, command and control, sensors and power systems, industry has successfully fielded a portfolio of UMS options that include:

- *Small UUVs* are 3" to 10" in diameter and are deployed from a variety of platforms (e.g., a larger UUV, submarine, surface craft, man-portable) for short duration (~24 hours) missions, such as mine countermeasures, survey and bottom mapping or inspection of underwater infrastructure (e.g., oil and gas). The military has employed small UUVs to perform hull inspections, assess in-port security conditions and aid in intelligence preparation of the operational environment missions. Public safety services are showing growing interest in using small UUVs to inspect dams, evaluate water quality and respond to emergencies.
- *Medium UUVs* are deployable from a variety of platforms (although special handling equipment is often necessary) and have a diameter between 10" and 21". The larger power sources permit slightly longer duration missions, up to a few days, such as search and rescue/recovery and fixed infrastructure inspection and service. The oil and gas community leverages UUVs to assess the integrity of offshore drilling infrastructure, which serves as an exemplar of how autonomous underwater systems can reduce risk to humans.
- *Buoyancy Gliders* convert changes in buoyancy to forward motion, allowing for long duration (months) surveillance missions primarily aimed at data sampling. At the surface, a glider takes on additional water weight and adjusts its mass to maneuver into a nose-down pitch, gliding forward as it descends. At depth, the glider reverses the process, expelling water and pointing its nose up to float back to the surface. The Naval Oceanographic Office leverages over 100 gliders, with roughly 1/3 of the gliders deployed globally at a given time to gain situation awareness and understanding of the world's ocean environment.
- *Wave Gliders* are powered by wave and solar energy. The system consists of a float equipped with solar panels used to recharge batteries and an attached glider operating several meters below the water's surface. With water speeds of 1-3 knots, these vehicles can gather data for up to a year, operating alone, or in fleets.
- *Remotely operated vehicles (ROV)* allow DoD and emergency service personnel to operate the vehicles from either shore or a surface vessel via the vehicle's tether cable. Micro-ROVs weigh a few pounds and can enter restricted areas, such as pipes, while the largest ROVs can dig trenches and lay cables at depths to 6000m. ROVs facilitate gathering information

pertinent to situation awareness and when equipped with actuators can be used for maintenance and construction of infrastructure in a broad range of environments.

- *Surface Vehicles (USV)* range in size from the man-portable MARTAC Mantas to the 12m Bonefish. DoD applications of this technology includes persistent surveillance and mine countermeasures. USVs allow emergency service providers to remotely operate and deliver life saving equipment, including in place devices (helmets and life jackets), shore based lanyards, and communications equipment. USVs enable remote operation in hazardous environments, such as severe flooding or extreme surf, which reduces personnel's exposure to extreme conditions. Modular payload systems will support search and rescue, search and recovery, damage assessment, and evidence collection missions. Payload examples include side scan sonar, forward looking infrared, and surface mounted radar systems.

Building on the successes of the first-generation UMS and advances in foundational autonomy technology for marine, air and ground systems, a shift is underway from UMS's merely serving as an extension of the sensor systems of manned ships and submarines to an integrated force multiplier capable of longer duration, more complex missions. However, it is noted that in order to field Large UUVs (diameter between 21" and 84") and Extra Large (XL) UUVs (diameter >84"), significant investment in infrastructure for launch from ship or shore, and in technology development for long duration missions will be required. Various national laboratories, aerospace contractors and academic research institutions are running Large and XL UUV demonstrations to experiment with innovations in propulsion, undersea communications, power systems and machine intelligence. Various DoD organizations have proven capabilities in collaborative multi-vehicle autonomy, USV navigation based on COLREGS, and detection and tracking of submarines in deep ocean regions. The resulting technology can serve as the foundation for developing future capabilities for the DoD and the emergency services.

5.1.3 Motivating Vignettes

The following vignettes offer examples of the increased capability and flexibility inherent in unmanned systems as the DoD and emergency services communities continue to field unmanned technologies and integrate resulting systems into existing structures.

Nuclear Contamination Threat

A seismic disturbance is detected 150 miles southeast of Anchorage, Alaska followed several minutes later by a more significant event in the same location. An interagency DoD/homeland defense reconnaissance UAS detects a radiation plume emanating near Montague Island at the mouth of Prince William Sound. The UAS maps the plume as it begins spreading over the sound, and a U.S. Coast Guard offshore patrol cutter deployed from Kodiak employs its embarked unmanned helicopter to drop buoys with chemical, biological, radiological, and nuclear (CBRN) sensors in the Sound and within narrow passes to measure fallout levels. The plume begins to spread over the sound and threatens the city of Valdez. All vessel traffic, mainly oil tankers transiting in and out of the Sound, is stopped, and operations at the oil terminal are suspended. Oil storage facilities at the terminal are quickly filled to capacity, and the flow from Prudhoe Bay is shut down.

Due to the growing contamination of the local environment, disaster response officials decide to request the support of the military because of their experience, both with operations in CBRN zones and with unmanned systems. An EQ-25, a very high altitude, extreme endurance UAS capable of operating at 75,000 feet for two months on station without refueling is dispatched over the Sound to ensure long-term, high-volume communication capability in the high-latitude, mountainous region.

A U.S. Navy amphibious transport dock ship anchors near an entrance to Prince William Sound and begins operations with its USV and UAS detachments. The radiation plume has now encompassed the evacuated town of Valdez, and UAS fly repeated sorties to the town, dock, and terminal areas to deploy UGVs with sensors and collect samples for analysis. The UAS and recovered UGVs are met and serviced back at the base by UGVs equipped to decontaminate the returning UMS after each sortie.

A USV proceeds to the focus of contamination and lowers a tethered remotely operated vehicle (ROV) to conduct an underwater search for the source. The USV's sonar quickly locates a large object in very shallow water and, on closer inspection by the ROV images the severely damaged hull of what appears to be a 50 year old, former Soviet-era, nuclear attack submarine. The hull is open to the sea, and the ROV places temperature gradient sensors on the hull and inserts gamma sensors into the exposed submarine compartments. The Joint Task Force quickly determines that the reactor fuel core is exposed to the sea due to not being shut down and is still critical.

With conditions deteriorating, two unmanned Homeland Defense CBRN barges fitted with cranes, containers, and remote controls arrive from Seattle. USVs are stationed in the narrow straits leading into the Sound with hydrophones to broadcast killer whale sounds to frighten fish outside the Sound away from the contaminated area. Over the next two weeks, with the assistance of U.S. and coalition ROVs equipped with cutting torches, grappling fixtures, and operating from USVs, one remotely operated submersible barge is able to work around the clock with impunity against exposure levels to recover the exposed fuel sources and to isolate them in specially designed containers. A second barge similarly retrieves sections of the crippled submarine. Both barges operate with a high degree of autonomy, providing critical information regarding the contamination to the Joint Task Force, while limiting exposure of personnel to the radioactive contamination.

The UGVs continue monitoring contamination levels and collecting samples, but also begin decontaminating the oil terminal control station and the local power and water facilities. Highly contaminated soil is placed into steel drums, and larger UGVs dig pits and bury contaminated building and pipeline materials. Advanced sensor technology and control logic allows the UGVs to operate around the clock during basic decontamination procedures, checking in, with human operators only when higher-level decisions are needed or for basic monitoring. UUVs crisscross the seafloor of the Sound to locate and tag remnants of the submarine for later collection. UAS fly continuously through the National Airspace System (NAS) at low altitude to monitor and map the declining radiation contours, at medium altitude to map cleanup operations, and at high altitude to relay control commands and data from the nearly one hundred unmanned vehicles at work. It is the largest coordinated use of international air, ground, and maritime unmanned systems ever conducted.

Littoral Pipeline Threat

An unmanned aircraft system (UAS) and an unmanned underwater vehicle (UUV), deployed from a U.S. Navy combat ship, are on patrol off the west coast of Africa to monitor the littoral oil infrastructure of a developing nation-state allied militarily and economically with the United States and friendly European governments. The UUV in its assigned patrol area detects an anomaly: a remote pipeline welder controlled by an unknown force. The underwater remote welder is positioning itself to intersect a major underwater oil pipeline. Using its organic "smart software" processing capability, the UUV evaluates the anomaly as a possible threat, takes a compressed data "snapshot" using its onboard video/acoustic sensor, and releases a communications buoy to transmit an alert signal and the data snapshot. The communications buoy's low probability of

intercept (LPI) data are relayed via the UAS to other units in the area and to the Joint Maritime Operations Center (JMOC) ashore. The commander onboard the Navy ship directs the UUV and the UAS to provide persistent intelligence, surveillance, and reconnaissance (ISR) and command and control (C2) relay support. As a result, the UAS, thanks to a recently fielded, advanced technology propulsion upgrade that enables it to stay on station for 24 hours before being relieved, detects a suspect vessel near the UUV anomaly and transmits corroborating ISR data.

A JMOC analysts recognize the pipeline welder in the UUV data snapshot as one recently stolen and acquired by rebel anti-government forces. The JMOC dispatches an Allied quick reaction force (QRF) from a nearby airfield and re-tasks a special warfare combatant-craft crewman (SWCC) Mk V to investigate and neutralize the potential hostile surface vessel controlling the stolen pipeline welder. The SWCC Mk V launches its own small UAS to provide a low-level ISR view ahead of its navigation track, while providing an LPI secure communications path among the special forces QRF team.

The JMOC receives a Signals Intelligence (SIGINT) alert that the suspect hostile surface vessel, having detected the LCS or visually sighted the UAS launched by the SWCC Mk V, is launching a Russian Tipchak, a medium-altitude, long-endurance UAS, capable of launching short-range air-to-air missiles (AAMs) or air-to-surface missiles (ASMs). Realizing the hostile UAS can pose a risk or even jeopardize the QRF, the JMOC commander launches a USAF UAS optimized for air interdiction and ground strike. The USAF UAS, empowered by rules of engagement allowing autonomous operation, immediately conducts an air-to-air engagement and neutralizes the Tipchak UAS. The SWCC Mk V's special forces team conducts a visit, board, search, and seizure (VBSS) on the suspected hostile vessel supporting the UUV pipeline interdictor. Since the threat is neutralized, the unmanned systems update their patrol status, cancel the alert status, and recover or resume their assigned patrol sectors.

Homeland Critical Infrastructure Protection and Inspection

The Port Authority of New Jersey, New York, and the Port of Miami receive notice that a terrorist event may occur within the next two weeks along either the Manhattan and Miami waterfront. They must prevent the unspecified event and prepare to respond should it happen, yet not restrict commerce or transportation. Both port authorities immediately re-task their USVs, which have been performing continuous routine inspection of piers, pilings, and seawalls 24/7 in fog, shadows, night, and varying temperatures. The updated underwater maps of the coastal urban infrastructure allow both agencies to prioritize continuous monitoring of targets with high value to terrorists. At the same time, an artificial intelligence planning algorithm identifies the shipping ports as potential economic terrorism targets and the cruise ship terminals as potential public terror targets. The algorithm also uncovers two other high consequence targets: a major fuel pipeline for Manhattan and an electrical and communications conduit in Miami.

UUVs are tasked to monitor the area at depths below the drafts of ships using advances in GPS-denied localization and computer vision. The long endurance UUVs will immediately surface and alert authorities if an anomaly is detected, otherwise will only routinely send a "heartbeat" indication that it is still on task. Small fixed wing UAS festooned with ribbons begin to fly over the port areas, circling near approaching smaller boats as if the UAS were public entertainment, but providing authorities with close up views of the boats without alerting terrorists or causing concern among the general public. The UAS surveillance continues in the evening relying on infrared and laser illuminations to monitor and inspect. Port authorities stop the larger boats and use ROVs to inspect

the hulls for possible hidden devices. The UAS monitoring produces a global map that alerts whenever a small vehicle appears to be approaching a vulnerable area. On land, ground robots are used to increase port security and to constantly circulate throughout the port to apply sensitive radiation and chemical detectors. Small UAS perch in trees or structures near where pipelines and conduits come ashore, ready to fly to provide over-watch if people approach by land or by boat. Within days of the initial alert, a high altitude, long endurance UAS spots a small commercial fishing boat with what appears to be suspicious cargo on the deck headed for the littoral area where the electrical and communications lines cross into Manhattan. After reviewing the threat data, the port authority deploys a swarm of brightly colored small UAS to do low altitude acrobatics around a nearby cruise liner, and to fly over the fishing boat on the way back to the landing zone to provide authorities with better intelligence without alerting the boat's pilot. The imagery confirms that the deck holds suspicious cargo. The Coast Guard acts to cut off access to the conduits and possible secondary targets, and deploys several fast USVs to circle the suspicious vessel and broadcast for it to stop engines and stand-by. One of the USV's detects the crew of the fishing vessel disposing of the suspicious cargo by dropping it overboard and uses its sonar to follow it as it sinks to the sea floor.

The other USV's continue to contain and monitor the suspect fishing vessel and one identifies what may be explosives strapped to the hull, possibly for a suicide mission, if captured. The Coast Guard vessels, while staying at a safe distance, deploy a USV equipped with a ROV used for hull and deep sea oil rig inspection to defuse the explosives attached to the hull of the fishing vessel. The Coast Guard moves in to board the fishing vessel and capture the crew, while the USV with the ROV moves on to locate and inspect the disposed cargo in collaboration with the USV that initially detected and tracked it. The ROV identifies the suspicious cargo as an underwater mine, deactivates it, and removes it from the sea floor. To help prevent a possible two-pronged attack, surveillance is heightened in all sensitive areas, with the deployed unmanned systems remaining on alert until authorities deem the threat has passed and all systems can return to their routine inspection tasks.

5.2 Scientific and Technical Challenges

Current and future advancements in technology will enable both single unmanned systems and teams of unmanned systems to work jointly with people, leveraging human flexibility and domain-knowledge in order to perform greater numbers of missions across multiple capability areas. In order to correlate similar needs, leverage effective solutions, and synchronize related activities, the DoD uses a Joint Capability Areas (JCAs) framework around which capabilities and capability gaps can be aligned across the DoD and across portfolios. These same capability areas generally apply to emergency service agencies. The nine Tier One JCAs each represent a collection of related missions and tasks that are typically conducted to bring about the desired effects associated with that capability. Mapping current and projected unmanned systems against the JCAs provides a sense of the portfolio of unmanned systems and how such systems currently, and in the future, can contribute to both the DoD and emergency services missions. Unmanned systems are key contributors in the Intelligence, Surveillance and Reconnaissance, Force Application, Protection, Logistics and Command and Control JCAs.

5.2.1 Intelligence, Surveillance and Reconnaissance

The Intelligence, Surveillance and Reconnaissance capability aligns with the Battlespace Awareness JCA in which all unmanned systems have the ability to contribute significantly into the future. This capability area applies across virtually all other JCAs and lends itself to tasks and missions conducted in all DoD and emergency services domains. Applications of this capability range from tasks, such as ground or aerial urban reconnaissance, which today are performed by UAS, such as Predators, Reapers and Global Hawks and UGVs, such as PackBots and Talons, to tasks, such as

Expeditionary Runway Evaluation, Nuclear Forensics, and Special Reconnaissance. UMS examples include using Slocum gliders for mapping ocean environments and the Remus 100 for identifying potential mines during mine countermeasure missions. Intelligence, surveillance and reconnaissance capabilities are also required to provide information for improved situational awareness and informed decision making by the public safety services. Similar to the DoD mission commander, emergency services personnel need to know where to position critical assets and how to identify hazardous threats. Future human-machine teams require a number of technological advances to provide military and public safety service personnel with the necessary intelligence and awareness, including: unmanned systems technologies hardened to withstand deployment in extreme environments (e.g., high temperature fires and radiation) and intelligent on-board data processing, perception and analysis in order to communicate more effectively with human and unmanned teammates, while also making accurate predictions and decisions of what actions to take in order to improve the team's intelligence, surveillance and reconnaissance mission.

5.2.2 Force Application

Force Application is a JCA that involves conducting maneuvers in order to enable gaining control and generating the intended effect in an environment. Military based force application includes target detection and identification, ballistic or non-ballistic firing solutions, selection of the firing platform, and battle damage assessment (BDA) of the results of the engagement. Emergency services force application includes extending the operational reach through real-time data collection, comprehensive damage assessments, communication enhancement, and various response activities. The Force Application JCA has seen a growing number of contributions from the proliferation of unmanned systems. Today, Predator, Reaper, and Gray Eagle UAS are weaponized to conduct high value target prosecution. The projected mission areas for UAS include air-ground, air-to-air combat, and suppression and defeat of enemy air defense, while emergency services UAS can be used to gather vital scene information through various payloads and provide the emergency services commanders with the necessary situational awareness to improve emergency response and enhance overall operations. Today small commercial grade UAS, such as the Draganflyer Commander, Leptron RDASS, and a very large selection of the DJI multirotor systems are being employed by the first responder community to gather intelligence in areas such as: large commercial fires, flood response, search and rescue, wildland fires, CBRNE response, damage assessments, and mapping of disaster areas. There is additional benefit in using small UAS for less critical activities, such as pre- and post-event planning, 3D mapping of accident scenes, and investigations. On the ground, UGVs are projected to conduct missions, such as dismounted offensive and defensive operations, and to some degree, mounted operations, such as armed reconnaissance. UGVs for the emergency services have been used to support SWAT team operations and in the future will be deployed for crowd control and suspicious device detection and suppression. In the maritime domain, UUVs and USVs are projected to be particularly suited for mine laying and mine neutralization missions, while public safety is presently using the Hydronalix Emergency Integrated Lanyard (EMILY), a USV, designed to make rapid access to victims caught in extreme flood and surf conditions. As the technology advances there will be a growing reliance on UMV's to protect and inspect national infrastructure. Current unmanned systems payload capabilities, sensors or munitions are driven by a number of factors; however, as unmanned systems become larger or payloads shrink in size and weight, unmanned system capabilities will grow.

DoD personnel must comply with the law of war, whether the weapon system is manned or unmanned. For example, Paragraph 4.1 of DoD Directive 2311.01E, DoD Law of War Program, May 9, 2006, requires that: "[m]embers of the DoD Components comply with the law of war during

all armed conflicts, however such conflicts are characterized, and in all other military operations." Current armed unmanned systems deploy lethal force only in a fully human-approved and initiated context for engagement decisions. The United States for decades have operated defensive systems for manned ships and installations that have human-supervised autonomous modes. For the foreseeable future, decisions over the use of force and the choice of which individual targets to engage with lethal force will be retained under human control in unmanned systems.

5.2.3 Protection

The protection JCA focuses on preventing and mitigating the susceptibility and vulnerability of military personnel, public safety personnel and the nation's residents and assets. Protection includes measures taken to harden positions, systems and personnel, while maintaining the self-awareness, early warning capabilities to preclude detection or surprise. Unmanned systems are ideally suited for many protection tasks that are deemed dull, dangerous or dirty. Military protection activities include early warning associated with flank security of ground, air or maritime forces. The protection JCA includes the wide array of threats - whether the threat is forces, systems or chemical agents. As the future enables greater autonomy with respect to perception, navigation and manipulation, unmanned systems will be able to perform, such tasks as firefighting, decontamination, infrastructure inspection, forward operating base deployment, installation and security of sensitive or cordoned areas, obstacle construction and deconstruction (breaching), inspections of both vehicles and personnel, explosive ordnance duties that include, detection, neutralization, and disposal, both casualty extraction and evacuation, and maritime interdiction. Human-machine teaming, perhaps as ad-hoc teams will require collaboration within and across domains in order to achieve the protection JCA objectives. The Dallas emergency services remotely controlled an Andros F5 UGV to mitigate an active shooter incident on July 7, 2016. While the exact implementation of this JCA in Dallas may be of concern to many, the resulting action served to protect personnel and the civilian population. Future systems will need to provide distributed advanced perception and intelligent decision making as well as actuators that can maneuver and manipulate objects and structures throughout human-built environments. It is important that these future unmanned systems are capable of both perceiving and interpreting the status of their human teammates, while completing the communication loop and sharing their own operational status.

5.2.4 Logistics

The Logistics JCA is ideally suited for employing unmanned systems in all domains to deploy, distribute, and supply military forces and public safety personnel. Transportation of supplies and personnel or civilians is an application particularly suited for unmanned systems in widely varying terrain and domains. Maintenance related tasks, such as inspection, decontamination, and refueling can be performed by unmanned systems. Munitions and material handling are tasks that can be assigned to unmanned systems to enhance safety as well as increase efficiency. Additionally, casualty evacuation and care, human remains evacuation, and urban search and rescue can also be assisted by unmanned systems. Unmanned systems will perform Logistics tasks on home station as well as forward deployed. The TerraMax UGV and K-Max UAV have been tested to support military logistic missions, but unmanned systems to support logistics are not yet common. Further, corporations such as DHL and Amazon are investigating the use of UAVs to support commercial logistics and the results are expected to influence the available technologies to support logistics missions related to public safety. Future systems will need to support both teaming with military and emergency services personnel, but also understanding of and interaction with the nation's civilian population. Further, these systems will require system hardening to withstand extreme environments along with reliable perception and intelligence capabilities and novel actuators to

transport and manipulate supplies, personnel and civilian casualties to support potentially long-duration complex missions.

5.2.5 Command and Control

Command and Control is a JCA focused on how to synchronize and communicate information regarding the mission and environment in order to facilitate military and public safety personnel's comprehension and decision making by providing predictions of potential mission planning and action outcomes and the allocation of resources necessary to achieve mission objectives. Current unmanned system command and control systems range from line-of-sight radio-based single operator direct control of a single UAV, such as the eBee, to more advanced flight planning and monitoring systems, such as the eBee's flight manager system, to common operator control systems that permit single or multiple operators to deploy heterogeneous unmanned systems for various missions simultaneously, such as the Navy's common control system, to the more complex ground control stations used for deploying the Predator UAV. The breadth of unmanned systems, including size, capability, longevity, payload capacity, autonomous capabilities, etc. provide clear reasons why varied command and control interfaces are needed. It is important to understand that machines are good at processing large amounts of data, often in a repeatable manner, but have limitations in perception, intelligence, and communication; in contrast, human personnel are well-equipped to make higher-level decisions. The majority of current unmanned systems still rely on direct commands from the human (teleoperation), have limited autonomous capabilities (e.g., follow a specified path, avoid obstacles), and limited intelligent perception, prediction, and decision making capabilities. Future unmanned system will require capabilities that allow the systems to perceive and understand the environment, identify the most relevant information to share with their human and unmanned teammates, to cue actions to their human and unmanned teammates, predict potential action outcomes, and execute those actions based on the mission, domain and environment. For example, a military UAS tasked with locating a target must perceive and determine that the correct target has been located and begin tracking that target, while simultaneously notify personnel at the forward operating base. Similarly, a public safety UAS may need to determine that an individual is acting suspiciously, begin tracking the individual's movements and intelligently communicate the decision triggers and description of the individual, behaviors, and location to the responsible personnel. As unmanned system autonomy improves in its capabilities and reliability, these systems will be partnered directly with humans. The resulting human-machine teams will be required to understand the overall mission, the team's goals and each member's responsibilities, as well as the intentions and current state of one another. This more complex command and control interaction will require advanced interaction capabilities that can be reliably, efficiently and effectively used in extreme environments that require cumbersome personal protective equipment (e.g., heavy bullet proof vests and Level A protection suits) or manipulating other equipment. Examples of advanced interaction capabilities include physiological monitoring of human cognitive and physical state, learning models of typical mission and personnel profiles (e.g., general roles and actions as well as individual differences), and naturalistic interaction capabilities (e.g., gesture and natural language interpretation). Regardless of the command and control structure, future systems must facilitate the ability of the military and public safety personnel, as well as the nation's civilian population to develop trust in the unmanned systems. The development of trust by the military and public safety personnel will require transparent systems that provide relevant, reliable and useful information to support mission situational understanding, decision-making and resource allocation. Gaining the trust of civilians, such as mobile wounded victims who are instructed to egress by following a UGV to a decontamination station, will require unmanned systems that embody human social behaviors.

5.3 Technology Needs and Goals

Unmanned systems necessary for achieving the described JCA missions require developing and fielding systems that include the following technology needs and capabilities:

- Data Representation: translate vast quantities of existing intelligence and sensor data into a shared and relevant environmental understanding.
- Intelligent Perception: enable greater onboard processing that facilitates improved change detection, aided (AiTR) and automatic (ATR) target recognition, image and sensor data analysis in complex environments (e.g., fires smoke, collapsed structures), and enhanced tracking of mission operations and personnel (e.g., structural firefighting, flood response, and wildfire activities).
- Longevity: enable mission endurance to extend from minutes for small unmanned systems to hours, days, weeks and months so that unmanned systems can conduct long endurance and persistent missions.
- Robust Systems: enable hardened unmanned systems, both actual hardware and software to breach and operate in environments that currently disable the technology (e.g., extreme hot and cold temperatures, wind, water/rain, complex radio/electrical signal interference, and CBRNE environments).
- Agile Actuators: enable unmanned systems to maneuver freely in complex environments (e.g., rubble piles, high winds and currents, human-made structures) and manipulate objects across broad spectrum of domains (e.g., clear debris from sand and dirt to destroyed vehicles, open doors, move soft and compliant human limbs and bodies).
- Distributed Perception: provide the systems with their own organic perception from on-board sensors so that they can autonomously contribute to missions.
- Intelligence: enable greater cognitive functions and collaborative awareness individually and among unmanned systems in a cohort.
- Individual Autonomy: enable actions that result from cognitive functions (e.g., move to physical proximity of the target or victim and deploy a sensor to further inspect an area of interest or assess the victim's state, respectively).
- Shared Autonomy: enable seamless transfer of and shared decision making and resource allocation to actions with team members, either human or unmanned in order to improve operational efficiency and flexibility.
- Control: advance capabilities to provide efficient, intuitive and natural interactions for tasking, and system configuration to support interaction and coordination between humans and unmanned systems.
- Information Sharing: provide personnel with intuitive and transparent information related to the environment, team and individual unmanned system state that supports rapidly and accurately building situational understanding and making decisions.
- Human State Assessment: model predicted and perception of current human cognitive and physical state, communications and intent so that the unmanned system can intelligently respond, assist or adapt system interactions and actions appropriately in order to improve team performance.
- Teaming: support human-machine teaming by developing and modeling information transfer and human-only team behaviors to improve communication, individual member and overall team performance, while building trust.

The technology challenges can translate to notional five, ten, and fifteen year goals. These envisioned goals are presented by domain for each of the JCAs. Unless otherwise noted, the goals apply to both the DoD and emergency services (ES) domains.

	5 Year	10 Year	15 Year
Unmanned Air Systems			
Intelligence, Surveillance and Reconnaissance	Maintain geospatial relationships with unmanned and manned aircraft. Respond to orders to conduct and investigate points and areas of interest. Greater ATR/AiTR. Improved intelligence assessment. Perception of location and flight path intention of friendly air systems. Ability to search for specific threats, includes both image and non-image based sensors. Improved target recognition and human presence detection. Small UAS platforms and advanced sensor capabilities designed for emergency services (ES only).	Ability to respond to standard ATC procedures in both military and civilian operations. On-board automated target recognition. Ability to detect and respond to threat UAS. Advanced intelligence assessment. Hardened-small UAS and sensors that withstand extreme environments.	Coordinated multiple unmanned vehicles (UAS, UGS, and UMS) to collect intelligence or search for threats. Hardened-miniaturized advanced sensors that withstand extreme environments.
Force Application	Limited air-air counter-UAS capability. Includes both the detection suite and for DoD, support munition. Extend operational reach to facilitate event planning. Standardized operational concepts, credentialing, and training for emergency services (ES only).	Intelligent damage assessment. Extend operational reach to extreme environments to facilitate event planning. Advanced air-air counter-UAS capability (DoD only).	Advanced comprehensive damage assessment. Intelligently extend operational reach to extreme environments and provide event planning. Multiple target counter-UAS capability. Cooperative engagement capability in a counter-UAS engagement (DoD only).
Protection	Integrated human-machine teaming for the purpose of early warning. Ability to identify threats to downed-personnel and systems. Tracking capabilities of emergency personnel (ES only).	Automated inspection of infrastructure. Detection of CBRNE hazards. Location communication, including in GPS-denied environments of downed personnel, casualties/victims and systems. Armed systems contributing flank security to a dismounted force (DoD only).	Communicate route to, or autonomously guide responders to downed personnel, casualties/victims and systems in GPS-denied environments. Collaborative engagement of UASs against a standoff threat (DoD only).
Logistics	Unmanned logistic resupply.	Unmanned medical evacuation with on-board human assistance and intervention.	Automated re-stocking, take-off and landing of autonomous UAS cargo planes, operating 24/7.
Command and Control	Common operator control interfaces within UAS classes. Beyond line-of-sight, autonomous single UAS missions. Predictive planning systems that integrate basic situational information. Interaction modalities to support dismounted human operators and teammates.	Common operator control interfaces across UAV classes. Line-of-sight, autonomous multiple unmanned vehicles (UAS, UGS, and UMS) missions. Suggested cueing of other unmanned vehicles (UAS, UGS, and UMS) and personnel. Predictive planning systems that integrate advanced environmental conditions and unmanned system capabilities and states. Natural interaction modalities that accommodate stationary control stations and mobile command by humans wearing restrictive PPEs.	Beyond line-of-sight, autonomous multiple unmanned vehicle missions. Automatic cueing of other unmanned vehicles (UAS, UGS, and UMS) and personnel. Real-time predictive planning systems that integrate extreme environmental conditions, as well as human and unmanned systems capabilities and states. Natural interaction modalities that automatically assess and accommodate human intentions and humans' predicted and current state.

	5 Year	10 Year	15 Year
Unmanned Ground Vehicles			
Intelligence, Surveillance and Reconnaissance	Self-positioning for optimal sector coverage.	Self-emplacement and self-recoverable early-warning devices. Self-awareness of role within team of UGSs responsible for optimal sector coverage. Hardened systems to withstand extreme environments. Advanced intelligence assessment. Suggested cueing of other unmanned vehicles (UAS, UGS, and UMS) and personnel.	Self-awareness of role within team of humans and UGSs responsible for optimal sector coverage. Automatic cueing of other unmanned vehicles (UAS, UGS, and UMS) and personnel.
Force Application	Intelligent damage assessment. Extend operational reach to facilitate event planning. Mounting infantry heavy weapons (mortar-50 cal - missiles) capability on UGS capable of following dismounted operation (DoD only). Laser designation capability (DoD only).	Cooperative autonomous A-G, G-G, G-A engagements. Advanced comprehensive damage assessment. Extend operational reach to confined and complex spaces. Basic, real-time semi-autonomous manipulation of hazards in extreme environments.	Cooperative engagement and ability to provide suppressive fires and maneuver against a fixed position. Advanced, autonomous manipulation of hazards in extreme environments.
Protection	Firefighting systems. Communicating location, including in GPS-denied environments, of downed personnel, casualties/victims and systems. Autonomous guidance of human teammates or civilians through GPS-available environments.	Armed systems contributing to flank security to a dismounted force. Communicating route to, or autonomously guide responders to downed personnel, casualties/victims and systems in GPS-denied environments. Autonomous guidance of human teammates or civilians through complex and dangerous environments. Basic semi-autonomous infrastructure inspection.	Collaborative engagement of UGSs against a standoff threat. Autonomous guidance of human teammates or civilians through GPS-denied and extreme environments. Advanced infrastructure inspection.
Logistics	Integrated manned-unmanned convoys in which multiple, large, optionally manned vehicles autonomously traverse defined secondary routes as either leader or follower vehicle under the supervision of a nearby operator. Unmanned medical evacuation in accessible environments.	Material handling UGSs identify, unload, load and secure containerized or palletized cargo fully autonomously at distribution centers under all environmental conditions.	Fully automated logistics management system that tracks in-country inventory and loads and routes UGSs with needed supplies via ground lines of communication to units for just-in-time restocking with the option of no human input. Unmanned medical evacuation in extreme environments.
Command and Control	UGS to human communication mechanisms that improve situational awareness. Shared autonomy that permits simple decisions based on strategic objectives (e.g., transport route). Joint-autonomy systems that intelligently request assistance from, and make suggestions to human operators. Interaction modalities that accommodate stationary control stations and mobile command by humans wearing restrictive PPEs. Integrated human-machine teaming in accessible environments.	Shared autonomy that permits autonomous complex decisions based on strategic objectives. Suggested cueing of other unmanned vehicles (UAS, UGS, and UMS) and personnel. Real-time predictive planning systems that integrate extreme environmental conditions and unmanned system capabilities and states. Naturalistic interaction modalities that automatically accommodate human teammates in extreme environments. Integrated human-machine teaming in extreme environments.	Shared autonomy that permits autonomous complex decisions based on flexible objectives. Automatic cueing of other unmanned vehicles (UAS, UGS, and UMS) and personnel. Real-time predictive planning systems that integrate extreme environmental conditions as well as human and unmanned system capabilities and states. Naturalist interaction modalities with multiple UGSs for extreme environments that automatically assess and adapt for human intentions as well as humans' predicted and current state. Advanced human-machine teaming in extreme environments.

	5 Year	10 Year	15 Year
Unmanned Maritime Systems			
Intelligence, Surveillance and Reconnaissance	Automated following COLREGs. Remote sensor deployment. Metal and plastic mine detection by UMSs (DoD only).	Persistent automated surface and subsurface monitoring (user-in-the-loop). Human detection by UMSs. Semi-autonomous facility and infrastructure inspection and anomaly detection. Collaborative operations to operate as part of a wide area detection.	Persistent automated surface and subsurface monitoring (user-off-the-loop) with worldwide reach in all weather. Autonomous facility and infrastructure inspection and anomaly detection. Relocatable detection zones. Detection avoidance.
Force Application	Placement of hull-borne devices on static targets.	Remote maritime threat interdiction response. Counter submarine capability (DoD only).	Automated maritime threat interdiction response. Human team delivery. Coordinated multiple unmanned vehicles (UAS and UMS) to collect intelligence or detect and track threats.
Protection	Automated following COLREGs.	Automated interdiction of manned and unmanned threats. Armed UMSs contributing to flank security (DoD only).	Collaborative engagement of UMSs against a standoff threat. Fully automated ship and shore installation security from maritime threats (DoD only).
Logistics UMV	Automated health monitoring. Continuous ship and shore installation inspection. Automated hull cleaning (DoD only).	Automated UMS prognostics. Semi-automated underwater refueling. Automated resurfacing and painting (DoD only). Automated preventive ship and shore installation maintenance (DoD only). Condition based UMS maintenance (DoD only).	Fully automated ship and shore based operations (no hands-on operator interaction). Automated underwater refueling.
Command and Control	Near-term predictive planning systems that integrate environmental conditions and unmanned system capabilities and states. Automatic and accurate data processing and information extraction to simplify the intelligence gathering task.	Suggested cueing of other unmanned vehicles (UAS and UMS) and manned vessels. Advanced predictive planning systems for extended missions that integrate environmental conditions and unmanned system capabilities and states. Communication displays and tools to quickly and accurately bring human operators “back in the loop” after moderate periods of no communications with UMSs. Automatic and accurate data processing and information extraction from multiple sensor sources to simplify the intelligence gathering task.	Automatic cueing of other unmanned vehicles (UAS and UMS) and manned vessels. Advanced predictive planning systems for long-duration missions that integrate environmental conditions, as well as manned vessels and unmanned systems (UMS and UAS) capabilities and states. Communication displays and tools to accurately and quickly bring human operators “back in the loop” after extended periods of no communications with UMSs.

5.4 Contributors

The prior version of this chapter was based on the content of the Department of Defense “*Unmanned Systems Integrated Roadmap, FY2011 – 2036*”, input and feedback from a defense robotics workshop hosted by the Robotics Technology Consortium on December 4th, 2012 in Washington, DC, and contributions from the individuals listed below.

The current chapter is based on the content of the 2016 National Academies of Sciences, Engineering and Medicine “*Mainstreaming Unmanned Undersea Vehicles into Future U.S. Naval Operations*” and the Department of Defense “*Unmanned Systems Integrated Roadmap, FY2013 – 2038*”, input from the robotics roadmap workshops held on August 23, 2016 in Portland, OR and September 21-22, 2016 in Atlanta, GA and contributions from the individuals listed below.

2016 Contributors

Julie A. Adams
Vanderbilt University

Jim Overholt
AFRL

Coitt Kessler
Public Safety Unmanned Systems SME

Heidi Perry
Charles Stark Draper Laboratory

Heather Knight
Stanford University

Richard Voyles
Purdue University

2013 Contributors

Grant Bagley
*Concepts to Capabilities
Consulting*

Brian Julian
MIT Lincoln Laboratory

Mike Passaretti
Honeybee Robotics

Robert Bechtel
Soar Technology, Inc.

Vijay Kumar
UPENN

Robin Pope
SAIC

Sara Blackmer
Pratt & Miller Engineering

Alberto Lacaze
Robotic Research, LLC

Mike Ryan
BRTRC

Henrik Christensen
Georgia Tech

Chris Mailey
AUVSI

Adam Schuman
Chatten Associates, Inc.

Edward Chupein
USAF

Mario Mairena
AUVSI

Jim Shinn
ARA

Andrew Culhane
TORC Robotics, LLC

Jay McConville
Chandler/May Inc.

Bill Thomasmeyer
Robotics Technology Consortium

Robert Finkelstein
Robotic Technology, Inc.

Adrian Michalick
Lockheed Martin

Eduardo Torres-Jara
Worcester Polytechnic Institute

Matthew Fordham
Applied Research Associates

Gordon Fraser
ICS NETT

Satyandra K. Gupta
USC

Glen Henshaw
NRL

Rich Ivey
BAE

Wesley Mitchell
Lockheed Martin

Robin Murphy
Texas A&M University

Don Nimblet
Lockheed Martin

John Northrop
John H. Northrop & Associates, Inc.

James Overholt
TARDEC

Tom Van Doren
HDT Robotics, Inc.

Rocco Wall
*Integrated Microwave Technologies,
LLC*

Peter Wells
QinetiQ North America

Joe Zinecker
Lockheed Martin

Noah Zych
Oshkosh Corporation

6. Earth and Beyond

6.1. Overview

Earth is an enormous sphere, some 12,700 km in diameter, with roughly 150 million square kilometers of land surface and 360 million square kilometers of water surface. The deepest ocean is nearly 11,000 meters below the water surface, while the tallest mountain is nearly 9,000 meters above sea level. Earth's arable land accounts for just under 14 million square kilometers with the United States, India, Russia, China and the European Union accounting for nearly half of that total. Another 1.5 million square kilometers is used for "permanent crops" (such as orchards and vineyards) and nearly 34 million square kilometers are used for pasture. The atmosphere is less than 500 kilometers thick and beyond that, a single moon orbits the Earth. Within the larger solar system are four terrestrial planets, of which the Earth is classified, and four Jovian planets ("gas giants"). The scale of humans' efforts to provide food and water to maintain life, to transport products to maintain commerce, and to extract raw materials to supply production suggests myriad ways robotics can contribute to the safe, efficient, and sustainable support of human society. Beyond simply maintaining life, robotics can actually help to enhance the quality of life, and support one of the unique characteristics of humanity: the curiosity to explore and explain the world and universe around us.

Population and food production by 2050. Climate change (windmill manufacturing like Boeing aircraft mfg) Sensing and monitoring the environment for intelligent management and clean-up. Robotic ocean monitoring & exploration. UAS Air traffic management.

6.1.1. Robotic Handling of High-Consequence Materials

The handling of "high-consequence materials" is becoming an increasing concern within several government agencies, as well as the public advocacy groups that oversee their operations. An example of a "high-consequence material" is ebola virus samples that researchers use in high-level "biosafety level 4" (BSL-4) facilities to investigate the virus and to explore cures and vaccines. The recent outbreak of ebola brought to light important deficiencies in not only the breadth of our medical research and our preparedness to respond to medical epidemics, but to the risks that human scientists and technicians routinely expose themselves to – and, potentially to the public, at large, should an accident occur – in order to develop cures and safeguards for the public health. Similarly, space scientists will be increasingly dealing with samples from extra-planetary sources, such as comets, asteroids and, eventually, distant worlds, that will require careful isolation and handling.

Perhaps no government agencies are more concerned with high-consequence materials than those that deal with nuclear materials and wastes. The Department of Energy Office of Environmental Management (DOE EM) is the primary steward of the land, water and waste storage facilities involved in the development of nuclear weaponry and nuclear power generation that began during World War II. Many of the relics of early uranium and plutonium processing facilities remain in place, today, waiting for a remediation. These sites which are located throughout the US involve wide ranges and concentrations of high consequence chemical and radioactive materials. In the past two decades, considerable progress has been made in removing and stabilizing the hazards associated these facilities. In many cases, "green field" conditions have been established assuring future public safety and health. Many decades of remaining remediation work remain to properly address the remaining facilities, many of which are the most complex and involve the highest

residual nuclear radiation levels that exist. On a continuing basis, these facilities must be monitored and their status formally verified. Future remediation projects will require the extensive use of various levels of remote operations involving complex remote manipulation and handling in uncertain and unstructured environments. A unique aspect of nuclear remediation is radiation effects on materials and components. It is common for nuclear radiation levels to be significant with respect to material and component degradation. Therefore, the accumulated radiation exposure of robotic systems must be carefully addressed through radiation hardening, exposure monitoring, periodic replacement of components. In many scenarios, work will be done by humans with direct contact and exposure to hazardous conditions. The highest priority is to assure worker and public safety throughout the decades of future remediation projects. The genesis of robotic manipulation goes back to the very early days of nuclear research in the 1940's. There is a certain degree of irony in the fact that some of the greatest challenges and opportunities for modern robotics technology exist in this same domain today.

Nuclear remote operations are notoriously expensive and slow due to the complexities and hazards of the remote environments, which impose stringent requirements on operational preparedness and supervision. Modern day robotics technologies offer fresh opportunities to increase remote work efficiencies, to increase the range of admissible tasks that can be performed reliably, and to free remote human operators for repetitive and fatiguing tasks. As examples, emerging cognitive robotics and machine learning concepts offer promise of a next generation of remote telerobotics that provides remote human operators with reliable remote systems that can truly share tasks with remote intelligent agents (remote mobile sensor/manipulator systems) intuitively and fluidly, similar to the way human work teams execute and collaborate during direct contact operations. The net result will be overall remediation systems that complete projects faster and cheaper than possible today. In addition to remote operations, there is an expanse of remediation project work that will be done by humans directly, in some cases using established personal protective equipment (PPE) such as protective suits and respirators. DOE EM is interested in enhancing worker protection and safety through robotics technology as well. There is particular interest in assistive robotics such as effective exoskeletons which improve worker capacity and stamina and to reduce job-related injuries.

In summary, the nuclear remediation challenges that DOE EM faces represent classes of problems that provide grand and challenging opportunities for reliable and effective modern-day robotic solutions. The ages and technical complexities of the facilities results in highly uncertain and unstructured remote work environments that are similar to, and on the same level of, future space and undersea challenges. It is believed that the motivations and requirements implicit in practical nuclear remediation applications will provide essential research drivers and success metrics that “pull” future research and developments to success in many other application domains.

i. Clean-up of the Portsmouth Uranium Enrichment Facilities

Gaseous diffusion was the chemical process that was developed to concentrate natural Uranium 235 for use in nuclear weapons. The nature of the process resulted in massive equipment and facility physical sizes, tens of thousands of acres of roofed facilities. The original gaseous diffusion plant was constructed at the Oak Ridge complex and was fully remediated in the 1990's and early 2000's using more or less conventional demolition methods and contact operations. The uranium enrichment process involves comparatively low level radioactive contamination, a wide range of toxic chemicals and nuclear criticality concerns regarding the residual hold-up of fissionable uranium inventory in

the process piping and equipment. As a result, the remediation at Oak Ridge was a lengthy and expensive endeavor.

Plans are now underway for the remediation of the other gaseous diffusion plant located in Portsmouth, Ohio. The remediation of the Portsmouth Gaseous Diffusion Plant represents a prime opportunity for the integration and use of robotics. Gaseous diffusion is multi-stage sequential refinement process on the scale of kilometers in length and many process components in the ton range. The two main process buildings at Portsmouth are each over a kilometer long, stretching nearly 2.4 kilometers in total length. This physical scale and multiple instances of the same process components leads to many opportunities for robots to aid in the remediation.

A wide range of robotic innovations are possible, and the following are examples of interesting possibilities. Pipe climbing/crawling robots that can move about the expanse of piping and process equipment to search for and locate residual uranium build ups would be of great value. Large scale demolition robots that could remove, segment and package large process components, many of which involve classified materials, could substantially reduce worker hazards and increase productivity. Payload amplifying exoskeletons could play a major role in reducing worker injuries and hazard exposures as well. Modern sensor-based mobile robots could dramatically improve the resolution and speed of pre and post operational radioactive and toxic materials surveys.

6.1.2. Robotic Reconnaissance for Human Exploration

Robotic reconnaissance prior to human activity has the potential to significantly increase scientific and technical return from planetary exploration missions. Robotic reconnaissance involves operating a planetary rover underground control, or IVA astronaut control, to scout planned sorties prior to human EVA. Scouting can be: (1) traverse-based (observations along a route); (2) site-based (observations within an area); (3) survey-based (systematic collection of data on transects); or (4) pure reconnaissance. Scouting can be done far in advance to help develop overall traverse plans. Scouting can also be done just prior to an EVA to refine an existing traverse plan, i.e., used to adjust priorities and modify timelines.

Although orbital missions can produce a wide variety of high-quality maps, they are limited by remote sensing constraints. Instruments carried by planetary rovers can provide complementary observations of the surface and subsurface geology at resolutions and from viewpoints not achievable from orbit. This surface-level data can then be used to improve planning for subsequent human sorties and missions, especially by reducing uncertainty in targeting and routing. Moreover, surface-level data can be used to improve crew training and to facilitate situation awareness during operations. As a practical example of how robotic reconnaissance would be extremely useful for future human planetary exploration, consider what took place during the last human mission to the Moon. During Apollo 17's second EVA, the crew drove from the lander site to the South Massif, then worked their way back. At Station 4 (Shorty Crater), Harrison Schmitt discovered numerous deposits of orange, volcanic glass – perhaps the most important discovery of the mission. However, time at the site was severely limited due to the remaining amount of consumables (e.g., oxygen) carried by the astronauts. Had the presence of this pyroclastic material been identified in advance through robotic scouting, the EVA timeline could have been modified to allow more time at Shorty Crater. Alternatively, the traverse route could have been changed to visit Shorty Crater first.

6.1.3. Planetary Cave Exploration

Planetary caving has been envisioned for a century, but has been beyond reach because no ways were known to gain cave entrance. Compelling motivations for cave exploration include studying the origin, geology, life signs and suitability for human haven that are not possible from the surface. The impossibility of access has recently been transformed by discovery of hundreds of Skylights on the Moon and Mars, and intimation of others in the solar system. Skylights are planetary holes with steep cylindrical or conical walls of rock. Some expose entrances to the prized underground caves. There is great scientific ambition to explore life signs, morphology and origins in these recently-discovered and completely unexplored areas. Surface robot technologies and missions have succeeded for half a century, but those capabilities fall short for descending walls, bouldering over unweathered floors, and getting around in caves. The caves deny light and line-of-sight communications and call for new forms of interface and autonomy. New but achievable robotic technologies are required for exploring these holes and tunnels.

Kilometer-scale landing precision is sufficient for many missions, but meter-scale accuracy, if achieved, could guide a lander to bisect and view a Skylight hole from a once-only, close-up, down-looking birds-eye view. After planning a landing spot near the rim, the robot touches down, disconnects, and proceeds to explore the floor and cave. A rover could approach, view, circumnavigate, model, and study the Skylight's apron, rim, and portion of walls that are visible from a safe standoff. Descent might occur by rappel, or by lowering like a spider from a line spanning the Skylight. Either requires unprecedented robotic rigging and anchoring. Sensing at high resolution across and down great distances and in extreme lighting variances poses new challenges for perception and onboard modeling. Repeated descents and ascents are desirable for whole-Skylight coverage, and for close-ups of the walls and floors. After thorough exploration, the rover re-attaches to its dangling tether and ascends like a spider to its highline, then out of the hole and on to the next Skylight.

6.2. Gaps and Impacts

6.2.1. Agriculture

According to the USDA¹, the world's population will reach 9 billion people by 2050 and total worldwide food production must double. In the U.S. there are nearly 2.2 million farms and the total arable land is reaching its limit. Yet, through productivity gains, the average U.S. farmer went from feeding 19 people in 1940 to feeding 155 people today. Yet, in the process of ever-increasing production and productivity, agriculture now consumes roughly 50% of the land in the U.S., 10% of the total energy budget, and 70-80% of available fresh water, while depositing 20 million tons of fertilizer and pesticides. And by focusing productivity gains on engineered plant types that are more resilient to transport and handling, many argue the nutritional value of our agricultural product has diminished.

To continue that pace of productivity gains, robotics will need to play an ever increasing role in

¹ Sonny Ramaswamy, "Big Data and The Future of Agriculture," Intl. Conf on Soil, Big Data, and Future of Agriculture, June, 2015 in Canberra, Australia.

Dan Schmoldt, "Strategic Vision for CPS and Robotics," 2015, USA.

Precision Agriculture - Field Crops

Precision field crop agriculture is seen as one of a handful of key, early-adoption, commercial opportunities for drones and drone-based technologies because of the demonstrated ability to survey large areas from an elevated vantage point to measure and map crop needs. Mechanized equipment for planting, watering, fertilizing, and harvesting crops already is largely tied to GPS for adaptability based on location, so robotic aerial sensors to feed the equipment is natural. In addition, inspection of high value crops, such as fruits and nuts, has attracted interest in ground-based robotics. Various systems for determining disease, blight, readiness for harvest, and labor assistance with transport have been prototyped.

Precision Animal Agriculture - Livestock and Poultry

Exploration of precision animal agriculture has been less robust than precision field crops, but many opportunities exist for robotic assistance. Like plants, animals are sensitive to the proper amounts of food and water intake, are subject to disease and injury, and go through reproductive cycles that can disturb normal activities and production. However, the time scales to respond for precision animal agriculture are much shorter than for precision crop agriculture. Most dairy cows are fed at least twice daily, for example, and levels of sustenance are targeted to the median of the herd. Metabolic disease is an issue of great national and societal import as the general public has become increasingly focused on general animal well-being and the possible overuse of antibiotics administered in a proactive manner. For example, metabolic diseases like ketosis and acidosis are extremely prevalent on U.S. dairy farms and these conditions can impact milk production and product quality.

6.2.2. Environmental Monitoring

Air

The air is comprised of gases, which makes survival on the planet possible. As per Newton's third law, every action force is opposed with an equal and opposite reaction force. As a result of energy consumption from rapid economic development and industrialization there is an equal and opposite reaction in the form of air pollutants emission. Air pollution has an adverse effect on all living organisms which share the planet earth. As a country, America represents 4.4 percent of the world's population but contributes around 40 percent of the world's emission of carbon dioxide. Major cities like Washington D.C., New York, Chicago, San Francisco, etc., where population is very high, are some of the cities with high particle pollution and people living in these cities report more health issues than normal. In this respect of knowledge, it is of high importance to continuously monitor the air quality and pollutants in the atmosphere.

Concept of Unmanned Aerial Vehicle (UAV) to monitor concentration of gases in the air could be used. UAVs could be simply equipped with sensors like Electrochemical Cell sensor, Metal Oxide Semiconductor sensor etc., to detect various pollutants in the air. Based on the requirements, various types of UAVs like fixed wing, a vertical take-off and landing (VTOL), blimp etc., can be used. Recent advances like Consensus Based Bundle Algorithm (CBBA) and decentralized network architecture in the field of multiple UAVs make the air monitoring task even more enticing. Wireless sensor networks (WSN) is already a well-established process for air monitoring. Integrating WSNs with multiple UAVs could be helpful in mitigating issues faced by static sensor networks and could open up new ways for air monitoring systems. Using UAVs has its own advantages and

limitations; some advantages are that they can monitor large areas in user defined intervals and also UAVs can navigate to remote and inaccessible places. Major challenges in using UAVs are Federal Aviation Administration (FAA) restrictions and privacy concerns in major areas. Integrating these monitoring technologies with cloud computing might be of benefit as a real time relay of information to the recipients can be achieved and a large amount of sensor data can be stored and processed to make useful conclusions using cloud infrastructure. Monitoring air continuously with cloud Robotics will help in providing important data about air quality to the policy makers, scientists, and everyone concerned and enable them to act accordingly.

Water

All life depends on water, and we are all citizens of a watershed. Improper human activities typically lead to water contamination, which can disrupt and disorganize a community and even become a threat to people's lives. The US Environmental Protection Agency (EPA) reported that 55% of rivers and streams in the US are in poor biological condition. Moreover, nutrient pollution has been threatening the lakes and coastal water tremendously, while the harmful algae boom is killing massive marine fish and shellfish in the oceans. Late reaction to water crisis not only results in massive financial support and more time invested to recover the situation but also leads to social and political affairs, such as the Flint Water Crisis. Due to this, a water quality monitoring system, which features real-time monitoring and early warning, is of great significance.

Currently, there are successful prototypes nationwide such as the unmanned aerial vehicle (UAV), unmanned underwater vehicle (UUV) and unmanned surface vehicle (USV) water sampler, human-portable underwater robot, etc., for water and sediment sampling. They link aquatic ecologists with robotic researchers to enable cost-effective water sampling and monitoring system. Due to its strong mobility and recharging capability, UAV has realized water sampling in-situ for unknown situation. UUVs used in inspecting unreachable sewer pipe can produce 3D images of the gratitude of water pollutant. Similarly, UUVs exhibits strong capability in sediment sampling in order to monitor and predict the degradation of contaminant. Being freed from wireless signal shielding, USV is the ideal candidate in building water sensing network. Already, USV combined with active localization, extreme environment adaptability, and environmental energy harvesting has been applied in ocean and even iced lake monitoring.

Nowadays, Cloud Robotics renders more possibilities for robotic water monitoring system. Since the limitation in the communication as well as autonomous navigation of UUV are conquered by connecting the USV via wires, the combination of UUV and USV, carried with multiple accurate sensors, shall be able to provide a high dimensional 24-hour real time offshore sampling results. A database collecting those large magnitude of monitoring results will lay a solid foundation to a decision making mechanism based on machine learning. The decision generated from this intelligent robotic monitoring system will become the essential and regular reference for government in making public policies and guidelines.

6.2.3. Space Robotics

Driven by our natural curiosity, mankind through the ages has demonstrated a relentless desire to explore the unknown. In addition to opening up new worlds, this yearning for exploration has historically proven to generate economic growth and further a nation's resources, knowledge, and

power. The inventions inspired by the needs of exploration, including the discovery of new goods and materials, have served to generate enormous returns to a nation's economy. Since its inception in 1958, NASA has repeatedly demonstrated the veracity of this axiom by having accomplished many great scientific and technological feats in fulfilling its mission as our nation's agent for exploration beyond the bounds of our planet.

Much of what we know about the Solar System (and beyond), we owe to robotic probes, orbiters, landers, and rovers. These robot explorers have traveled on behalf of mankind through dark and deep space in order to observe, measure, and visit distant worlds. Equipped with sensors for guidance and observation, onboard avionics for control and data processing, actuation for locomotion and positioning, these robots have performed critical science and engineering tasks in-orbit and on planetary surfaces. Research in robotics, telerobotics, and autonomous systems has provided necessary technology to accomplishing these missions.

Looking forward, robotics, telerobotics, and autonomous systems figure heavily in NASA's strategy and are prominently mentioned in the US Space Policy released June 28, 2010. The policy states as one of its goals to "Pursue human and robotic initiatives" to develop innovative robotic technology and directs NASA to "Maintain a sustained robotic presence" in the solar system to conduct science and prepare for future human missions. The policy also indicates the need for immediate and sustained development and maturation of autonomous system technologies for numerous purposes, including the effective management of space power systems that will enable and significantly enhance space exploration and operational capabilities.

Robots and autonomous systems are already at work in all of NASA's Mission Directorates. Ongoing human missions to the International Space Station (ISS) have an integrated mix of crew working with both Intra Vehicular Activity (IVA) and Extra Vehicular Activity (EVA) robots and supporting autonomous systems on-board spacecraft and in mission control. Future exploration missions will further expand these human-robot "Co-Explorer" partnerships. While unmanned science missions are exclusively robotic in flight, they are integrated with Earth-based science and operations teams connected around the globe. In the future, NASA will see even more pervasive use of robotic co-explorer systems. Accordingly, NASA has developed a separate roadmap (synthesized herein) for robotics and autonomous systems technology expected to be integrated for dozens of planned flight missions of the four NASA Mission Directorates over the next 25 years.

The benefits to NASA of robotics and autonomous systems technology include: extending exploration reach beyond human spaceflight limitations; reducing risks and cost in human spaceflight; increasing science, exploration and operation mission performance; improving capabilities for robotic missions; providing robots and autonomy as a force multiplier (e.g., multiple robots per human operator); and enhancing autonomy and safety for surface landing and flying UAV's.

The benefits of this technology outside of NASA are potentially even more significant and include: bringing manufacturing back to America; developing new electric vehicles, more effective wind turbine control, better smart grids, and other green technology; enabling strategic asset inspection, repair and upgrade; increasing the extent and performance of automated mining and agriculture; creating more capable prosthetics, rehabilitation, surgery, telesurgery, and assistive robots; extending the reach of undersea robotics for exploration and servicing; infusing robots in education to stimulate Science, Technology, Engineering and Mathematics; enhancing the capabilities of personal

service, emergency response, hazardous material handling, and bomb disposal robots; and increasing the use of automated transportation via land, air, and sea.

These external benefits are consistent with NASA's strong record of developing and transferring innovative technology to the private sector. NASA technology can be found in virtually every civilian and military aircraft, in sensors for air quality, in breakthroughs to help the medical community better treat illnesses, and in new materials that keep our law enforcement and first responder personnel safe. NASA spin-off technologies have saved thousands of lives, have helped create tens of thousands of jobs, and have resulted in over \$6.2 billion in cost savings to companies and their customers. By one estimate, the total return on investment to the United States' economy, resulting from technology that NASA more or less freely shares with the public or US companies, is on the order of 700% return for every dollar invested in space exploration.

Achieving human-like performance for piloting vehicles

Machine systems have the potential to outperform humans in endurance, response time and the number of machines that can be controlled simultaneously. Humans have safety limits on flight or drive-time that do not exist in machines. Human response time, coupled with human machine interfaces, results in significant delays when faced with emergency conditions. Humans are poor at parallel processing the data and command cycles of more than a single system. But machine systems continue to lag behind humans in handling extremely rare cases, improvising solutions to new conditions never anticipated, and learning new skills on the fly. Achieving human-like (or better) performance leverages machine proficiency at controlling complex systems and requires: (1) getting the human out of the control loop and (2) engaging the human at the appropriate level (i.e. strategic direction, intent, etc.).

Access to extreme terrain in zero, micro and reduced gravity

Current crew rovers cannot access extreme Lunar or Martian terrain, requiring humans to park and travel on foot in suits. In micro gravity, locomotion techniques on or near asteroids and comets are undeveloped and untested. Access to complex space structures like the ISS is limited to climbing or positioning with the SSRMS. Challenges include developing robots to travel into these otherwise denied areas, or building crew mobility systems to move humans into these challenging locations. In addition to improved mechanisms and power, access to extreme terrain requires significant advances in robotic perception (sensors and algorithms) and vehicle control (servo, tactical, and strategic) capabilities. Perception is particularly important for detecting and assessing environmental obstacles, hazards, and constraints (e.g., locations to drive over, to grip, etc.).

Grappling and anchoring to asteroids and non-cooperating objects

Grappling an object in space requires a manipulator or docking mechanisms that form a bi directional 6 axis grasp. Grappling an asteroid and then anchoring to it is an all-new technology. Grappling approaches attempted on man- made objects may not apply to asteroids, since these techniques count on specific features such as engine bells that will not be available on a natural object. Similarly, grappling an object that is tumbling has not been attempted.

Exceeding human-like dexterous manipulation

The human hand is generally capable. A robotic equivalent, or superior grasping ability, would avoid the added complexity of robot interfaces on objects, and provide a sensate tool change-out capability for specialized tasks. Dexterity can be measured by range of grasp types, scale, strength and reliability. Challenges include fundamental 1st principles of physics in the development of actuation and sensing. Other challenges include 2 point discrimination, contact localization, extrinsic and intrinsic actuation, back-drivability vs. compliance, speed/strength/power, hand/glove coverings that do not attenuate sensors/motion but are rugged when handling rough and sharp objects.

Full immersion, telepresence with haptic and multi modal sensor feedback

Telepresence is the condition of a human feeling they are physically at a remote site where a robot is working. Technologies that can contribute to this condition include fully immersive displays, sound, touch and even smell. Challenges include 1st principles of physics in the development of systems that can apply forces to human fingers, displays that can be endured for long periods of telepresence immersion, and systems that can be used by people while walking or working with equipment concurrently with the telepresence tasks.

Understanding and expressing intent between humans and robots

Autonomous robots have complex logical states, control modes, and conditions. These states are not easily understood or anticipated by humans working with the machines. Lights and sounds are helpful in giving cues as to state, but need to be augmented with socially acceptable behaviors that do not require advanced training to interpret. Likewise, robots have difficulty in understanding human intent through gesture, gaze direction or other expressions of the human's planned behavior. In order to improve the quality, efficiency, and performance of human-robot interaction for space applications, a key challenge is to enable humans and robots to effectively express (communicate) their state, intent, and problems. This is true regardless of whether humans and robots are in proximity, or separated by great distance.

Rendezvous, proximity operations and docking in extreme conditions

Rendezvous missions include flybys of destinations without landing or docking. Proximity operations require loiters at destinations with zero relative velocity. Docking drives latching mechanisms and electrical/fluid couplings into a mated condition. Major challenges include the ability to rendezvous and dock in all ranges of lighting, work across near to far range, and achieve a docked state in all cases.

6.3. Contributors

The 2016 roadmap update was prepared by

Bill Hamel
UT Knoxville

Richard Voyles
Purdue University

Aaron Dollar
Yale University

Taskin Padir
North Eastern University

Magnus Egerstedt
Georgia Institute of Tech.

The 2013 version of the chapter was based on the content of NASA’s “Robotics, Tele-Robotics, and Autonomous Systems Roadmap” published in April, 2012, input and feedback from a tele-workshop held on November 26th, 2012, and contributions from the individuals listed below.

Rob Ambrose
NASA

Roland Menassa
General Motors Corp

Richard Voyles
NSF

Sarah Bergbreiter
University of Maryland

Rob Platt
SUNY Buffalo

Ian Walker
Clemson University

Terry Fong
NASA

Luis Sentis
UT Austin

David Wettergren
Carnegie Mellon University

Jeremy Frank
NASA

Mark Schwabacher
NASA

Red Whittaker
Carnegie Mellon University

Vijay Kumar
Univ of Pennsylvania

Vytas SunSpiral
SGT, Inc.

Brian Wilcox
NASA

William Thomasmeyer
RTC

7. Research Roadmap

In this section the main research challenges and opportunities across all the applications areas are outlined and provide an outline of progress over a 5, 10 and 15-year period.

7.1. Mechanisms and Actuators

How robots are built is changing. Robots of the past consisted of rigid links with actuators mounted to joints. Although many robotic tasks can still be accomplished with this design approach, new applications and new economic impact will be enabled by new manufacturing techniques, new materials and construction paradigms, and merging the design of actuator, mechanism, and control as a holistic process to generate compact systems that are both highly capable and energy efficient.

New Manufacturing Techniques: Additive manufacturing (3D printing) techniques are democratizing robot design, allowing complex shapes and structures by anyone with a printer. This democratization is enabling experimentation with new types of materials and sensor/actuator integration within robotic structural elements. 2D planar manufacturing processes, such as laser-cutting, are being used to create complex 3D geometries using origami-inspired methods. MEMS-based fabrication techniques make it possible to fabricate truly microscale robotic elements. Additive manufacturing can be used not only to produce useful components, but also as a part of the manufacturing process to generate molds for other materials, or forms for composite structures.

New materials and construction paradigms: 3D printed parts and softer polymers formed in 3D-printed molds, sometimes formed with other materials in a composite structure, have the potential to create a new paradigm of robot design that is more similar to soft biological machines and less similar to hard metal machines. While this field is in its early stages, it is clear that soft materials are far more effective than hard materials for gripping, manipulation, traction, and many physical interaction tasks. The strength and the challenge of soft materials is the complex dynamics of the materials; while compliance in a robot finger may be useful for gripping, it is also challenging to model, sense, and actuate. Continued development will yield new sensor paradigms, new actuators and transmissions (such as hydraulic bladders), and greater integration of the dynamics afforded by soft materials with the control methods for robot motion.

Merging the design of actuator, mechanism and control: There is an intricate interplay between the dynamics of mechanical devices, actuators, and the algorithmic complexity required to control them. Some algorithmic problems can be solved, or their solution can be greatly facilitated, by intelligent mechanical design; but only a subset of dynamic control problems can, or should, be solved in software. To illustrate a simple case, a spring-like behavior can be implemented either as a physical spring or through an actuator. But the choice has implications: If implemented through direct actuation, the performance may be poor due to actuator torque limits, high actuator inertia, speed limits, or other inherent actuator dynamics; and efficiency may be very low due to transmission losses or other actuator losses. However, while implementing spring-like behavior with a physical spring may address some concerns, it also constrains the device to *only* spring-like behavior, with no option to change that behavior. The behavior cannot be changed easily either for design revision during controller development, or to enable the robot to achieve multiple different tasks, some of which may not require spring-like behavior. This example illustrates that integrating the design of algorithm, actuator, and mechanism will enable new robotic tasks and outperform

traditional machines. New technologies for actuator, manufacturing, and construction paradigms will synergistically enable progress, as the line between control algorithm, hardware, and actuation blur.

Goals:	5 Years	10 Years	15 Years
Human-safe robot arms for manufacturing	Increasing adoption of compliant human-safe arms for simple manufacturing	Machines can partner with humans, including hand-off and other physical cooperation	Robots are increasingly common partners with humans in manufacturing, though still with limited capabilities
Human-like walking and running	Experimental robots walking outdoors and indoors, isolated demonstrations	Improved understanding of legged locomotion science enables efficient, agile locomotion demonstrations	Initial applications in commercial and military use, including logistics and telepresence

7.2. Mobility and Manipulation

7.2.1. Mobility

Mobility in the real world is enabled by perception, planning, and new mobility implementations (quadrotors, legged machines, swimming robots). Although significant progress has been made in recent years in all aspects of mobility, the problem is deep, and continued progress will yield important applications across all areas of our economy. Robots will become commonplace and as useful as cars and smartphones, amplified as such intelligence and information can be mobilized and applied to physical interaction with the world.

New Mobility Implementations: Robots should be able to go both where humans can go and to place where we cannot send humans. Both the DARPA Robotics Challenge (DRC), and the Fukushima Daiichi power plant disaster that inspired it, illustrated the limits of the state-of-the-art in their own ways; In the case of Fukushima Daiichi, no robots were capable of going into the power plant to assess the situation, even when such information was desperately needed. The 2015 DARPA Robotics Challenge demonstrated that state-of-the-art humanoid robots are slower than humans by an order of magnitude in performing tasks such as turning valves, using hand drills and flipping electric switches. Furthermore, the DRC robots relied heavily on pre-scripted motions and hence lacked autonomy to carry out the simulated tasks relevant to disaster response. We also learned that reliability in completing tasks is prohibitively low to make robots practical even for an overly simplified set of tasks. As a result, it is clear that much more work remains for mobile robots to go where people can go in an agile, reliable, and efficient manner.

Research into legged locomotion applies not only to mobile robots, but to human locomotion as well. Powered exoskeletons and prosthetic limbs are progressing rapidly, but again, have significant

room for improvement. Robot exoskeletons and prosthetics that are as agile and efficient as human legs will have a significant impact on quality of life for millions of people.

For nearly all aspects of locomotion and physical interaction, including swimming, flying, and walking, animals are inspiring examples of what is possible. As such, continued research into bio-inspired methods remain important. Bio-inspiration is distinct from bio-mimetics - animals should not necessarily be copied, but principles gained from understanding how animals work can then be re-interpreted for engineered systems, using different methods and tools to potentially outperform animal mobility, perception, and planning.

Perception and Planning: New capabilities in autonomous mobility are exemplified by museum tour guides and autonomously driving cars. Nevertheless, a number of important open problems remain. Participants identified **3D navigation** as one of the most important challenges in the area of mobility. Currently, most mapping, localization, and navigation systems rely on two-dimensional representations of the world, such as street maps or floor plans. As robotic applications increase in complexity and are deployed in every-day populated environments that are more unstructured and less controlled, these 2D representations will not be sufficient to capture all aspects of the world necessary for common tasks. It will be important for robots to perceive three-dimensional world models in support of navigation and manipulation. These 3D representations should not only contain the geometry layout of the world; instead, maps must contain task-relevant **semantic** information about objects and features of the environment. Current robots are good at understanding where things are in the world, but they have little or no understanding of *what* things are. When mobility is performed in service to manipulation, environmental representations should also include **object affordances**, i.e. knowledge of what the robot can use an object for achieving **semantic 3D navigation** will require novel methods for sensing, perception, mapping, localization, object recognition, affordance recognition, and planning. Some of these requirements are discussed in more detail later in this section. Participants also identified robust **navigation in crowds** as an important mobility challenge.

7.2.2. Manipulation

Substantial progress in manipulation is needed for almost all of the service robotics applications identified in the previous section. These applications require a robot to interact physically with its environment by opening doors, picking up objects, operating machines and devices, etc. Currently, autonomous manipulation systems function well in carefully engineered and highly controlled environments, such as factory floors and assembly cells, but cannot handle the environmental variability and uncertainty associated with open, dynamic, and unstructured environments. As a result, participants from all three break-out groups identified **autonomous manipulation** as a critical area of scientific investigation. While no specific directions for progress were identified, the discussions revealed that the basic assumptions of most existing manipulation algorithms would not be satisfied in the application areas targeted by this effort. **Grasping and manipulation** suitable for applications in open, dynamic, and unstructured environments should leverage prior knowledge and models of the environment whenever possible, but should not fail catastrophically when such prior knowledge is not available. As a corollary, truly autonomous manipulation will depend on the robot's ability to **acquire adequate, task-relevant environmental models** when they are not available. This implies that—in contrast to most existing methods which emphasize planning and control—**perception** becomes an important component of the research agenda towards autonomous manipulation.

Participants identified novel **robotic hands** (discussed in the subsection on Hardware), **tactile sensing** (see Sensing and Perception), and highly-accurate, physically realistic simulators as important enablers for autonomous manipulation.

7.3. Perception

Sensing and perception are of central importance to all aspects of robotics, including mobility, manipulation, and human-robot interaction. Participants were convinced that innovation in sensing and perception will have profound impact on the rate of progress in robotics.

Participants believed that new sensing modalities as well as more advanced, higher-resolution, lower-cost versions of existing modalities would be areas of important progress. For example, participants expect important advances in manipulation and mobility alike from dense 3D range sensing, including LIDAR and RGB-D sensing. Robustness and accuracy across wide range of environments is critical for further advancement. Advances in dexterous manipulation are likely to require skin-like tactile sensors for robotic hands and more specialized depth and appearance sensors for short range sensing. Additional sensors, for example acoustic sensors and specialized sensors for safety, were discussed by the participants. These sensors could take various forms, such as range or heat sensing to detect the presence of humans, or could be implemented by special torque sensors as part of the actuation mechanism, capable of detecting unexpected contact between the robot and its environment. Skin-like sensors for the entire robotic mechanism would also fall into this category.

The data delivered by sensor modalities must be processed and analyzed by near real time algorithms for perception in complex and highly dynamic environments under varying conditions, including differences between day and night and obscuration like fog, haze, bright sunlight, and the like. Approaches to perception capable of long term adaptation (weeks, years) will need to be developed. Participants identified the need for progress in high-level object modeling, detection, and recognition, in improved scene understanding, and in the improved ability to detect human activities and intent.

Integrative algorithms that use multiple modalities, such as sound, 3D range data, RGB image, tactile, are important to be considered. Participants believe that task-specific algorithms that integrate well with planning algorithms and consider dynamic physical constraints are needed. For example, novel algorithms for affordance recognition are important for areas such as dexterous manipulation for performing tasks in human environments. Creating contextual models that are situation-aware is important to be considered in robotics perception algorithms.

Robots are slowly beginning to operate in unconstrained environments and as such there is a need to provide robust perceptual functionality to cope with the environmental variation. Perception is critical to navigation and interaction with the environment and for interaction with users and objects in the proximity of the system.

Today perception is focused on recovering geometry, object recognition, and semantic scene understanding. We need to develop algorithms that go beyond recognition and geometry to task-relevant characteristics of entities such as objects (rigid and deformable), piles, environments, or people. Such characteristics include material properties, object affordances, human activities, interaction between people and objects, physical constraints of the environments. These are all necessary precursors for the development of advanced robot capabilities.

Computational models capable of handling uncertainty and scalability of basic perceptual capabilities along with frameworks for integrating them in a task-dependent manner need to be investigated.

In 5, 10, and 15 years the following goals are possible with sustained research and development:

- *5 years*: Sensing and perception algorithms should integrate information over time for robust operation in large scale settings. The robot will be able to perceive task-relevant characteristics of a wide-variety of environments and objects and will be able to recognize and locate and search for thousands of objects in cluttered environments.
- *10 years*: Basic capabilities of operating in static environments will be extended to dynamic environments. This will enable demonstration of robot system that can perceive dynamic events and human activities, so as to learn from and cooperate with humans. It is necessary to develop robotics-specific perception algorithms for domains such as dextrous manipulation, mobility, human-robot interaction, and other tasks. Development of large-scale learning and adaptive approaches that improve the perception over time will be necessary for deployment of systems capable for operating over extended periods of time.
- *15 years*: Demonstration of a robot that integrates multiple sensory modalities such as sound, range, vision, GPS, and inertial to acquire models of the environment and use the models for navigation, search and interaction with novel objects and humans. The focus will be on operation over long periods of time in cluttered, dynamic environments along with the adaptation of perceptual capabilities through exploration and/or interaction with humans.

7.4. Formal Methods

As autonomous systems, such as driverless cars and delivery drones, become a reality, a major challenge for robotics will be to develop methods and tools for safe and verifiable autonomy. This will require approaches for specifying safe robot behaviors as well as explicit assumptions for safe operation. Formal methods are mathematical approaches for reasoning about systems and their requirements. These approaches capture desired system behavior in a formal specification, usually using mathematical logic, and enable verification, synthesis and validation; *verification* is the process of ensuring a given model of a system satisfies its specification, *synthesis* is the process of generating a correct-by-construction system from the specification and *validation* is the process of ensuring the physical system (not only its model) satisfy its specification.

These approaches are crucial for the deployment of robots in the world; they will form the basis for certifying robotic systems, for ensuring the safety, security and predictability of robots and for enabling robotic “introspection” - the ability of a robot to report to people about the success and possible failures of its own behavior.

Over the past several years these approaches have been used to verify and synthesize properties of a vast collection of complex robots such as the Atlas humanoid robot. Collision avoidance algorithms have been formally proven, controllers for high-speed flying in a forest have been automatically synthesized and human-robot interaction has been formally defined. Future directions in research include the following:

7.4.1. Synthesis and verification for closed loop systems; incorporating uncertainty and dynamic environments

Robots operate in environments that are dynamic, and make decisions based on noisy information they gather using sensors. Given extreme circumstances, every robot will fail in its task; for example, a self-driving car might leave its lane in a whiteout. The next challenge is to develop formal methods that holistically reason about perception-actuation loops and take into account uncertainty about the state of the robot and the world. This will enable people to understand the limits of robot operation - in which environments the robots will be successful, under what conditions they will fail and when they should be trusted.

7.4.2. Safe behavior degradation

Typically, when robots fail, they fail spectacularly; self-driving cars collide with obstacles, humanoid robots fall and flying robots crash to the ground. Formal methods will be used to monitor the robot's behavior, reason about imminent failure and deploy strategies for graceful degradation where the robot fails safely.

7.4.3. Formal methods and learning

With the recent and continuing impact of machine learning on robotics, it is becoming imperative to develop techniques that guide, quantify and verify the performance of such algorithms in physical environments. Formal methods will enable reasoning about the robustness of such algorithms, monitor the robot's behavior to detect deviations from the expected and provide additional inputs that can be then used for training. Furthermore, formal methods may enforce constraints on the learning process thereby enabling "safe learning" where the system is guaranteed to satisfy certain properties while learning.

Formal methods in other domains such as software verification have shown a steady progress in scaling up tools to deal with realistic problems; however, formal methods in the context of robotics are not yet at the full system scale. Advances in machine learning can be leveraged to scale up verification and synthesis. Two distinct directions for research emerge: How to learn while satisfying safety, whereas the other is developing learning approaches for safety.

7.4.4. Formal methods for human-robot interaction and collaboration

As robots are moving away from industrial cages and into homes, workplaces and healthcare, they increasingly interact and collaborate with people. The next challenges are to formally model these interactions and collaborations, define specifications that include the person's role, and verify or synthesize the complete system - the robots behavior in the context of the interaction. Such methods will enable answering questions such as in a shared autonomy situation, when should the robot assume more control and when should it ask for more human intervention? In an assistive device, how can a robot help the person only when needed? How should the behavior of the person in the interaction be reasoned about to detect situations in which the human is impaired and the robot should take over?

7.5. Learning and Adaptation

Perception and planning/control enable robots to estimate the state of the world and decide how to act to achieve the task they are supposed to, like package goods or fill and bring a glass of water to an older adult. However, they only work when certain assumptions are met: the robot needs a dynamics model of how its actions are affecting state, a sensor model to make sense of the observations it gets, a clearly specified objective or reward function to optimize, and problem domain small or structured enough for planning and perception algorithms to produce solutions in acceptable computation times.

These assumptions will not always be met in our applications. A surgical robot might need to use a new tool, and not know how to model the interaction between it and the tissues it needs to contact. The uncertainty induced by human environments might render a service robot incapable of computing a good policy in a reasonable amount of time. A personal robot might clean up a room in a way that does not match at all what its owner wanted, and might be put away in a closet and never used again. And a factory robot might need to perform a new task as the factory rolls a new product out, its objective function changing from the task it was programmed to perform.

Machine learning has the potential to help robots adapt to these situations, enabling them to learn from their own experiences and improve over time, as well as to learn from humans.

7.5.1 Learning from Demonstration

Instead of an expert re-programming the robot to perform a new task, learning from demonstration empowers end-users (factory workers, service workers, consumers) to teach the robot what to do. At the moment, it is possible to demonstrate a task to a robot “kinesthetically”, i.e. by physically guiding the robot, and have it replay that motion and adapt it locally, e.g. to reach a new target or to avoid an obstacle. Remaining challenges include extracting task structure (identifying subgoals and objectives of the task), handling more than just kinematics (learning about the forces that the robot needs to apply), and being able to learn by observing people directly (as opposed to relying on kinesthetic demonstrations). More research is needed in making learning from demonstration algorithms easy to use by non-experts. Furthermore, learning and adapting to people’s preferences about the task will also be important: it is not always about achieving the task successfully, especially in collaborative and service robots adapting to what the end-user wants will be a crucial part of acceptance.

7.5.2 Reinforcement Learning and Deep Learning

One research avenue that has demonstrated a lot of successes in recent years is learning policies directly from experience. The robot practices a task, learning directly how to map states to the actions it should take. This can bypass planning, perception, and modelling of the world altogether, or learning can be harnessed for modelling and for speeding up planning. Policies learned through deep learning have led to impressive results, enabling robots (or AI systems) to play Atari games at human-level performance, or to win against Lee Sedol at Go. Other areas such as computer vision, speech recognition, and natural language processing have benefitted tremendously from progress in deep learning.

However, the physical world is far less structured and involves a continuous and high-dimensional space of states and actions. It also requires physically acting in the world to acquire data instead of using simulation only to improve. Therefore, while these new learning paradigms have the potential to greatly improve robot performance, making them work on physical robots performing complex

tasks opens up various research challenges. Among these challenges are: generating good estimates of the learned system's uncertainty, generalizing to new task domains, learning in regimes where data is expensive and sparse, and combining model-based reasoning and deep learning

7.6. Control and Planning

Robots of the future will need more advanced control and planning algorithms capable of dealing with single and multi-agent systems with greater uncertainty, wider tolerances, and larger numbers of degrees of freedom than current systems can handle. They will need to work safely and robustly in all settings -- ranging from fully autonomous operation in extreme environments to collaboration with humans at home or at work. Robot arms on mobile bases will have end-effectors that will need to efficiently plan and consistently perform fine manipulation and grasping tasks in unstructured and constrained environments. These robots might have 12 degrees of freedom. Anthropomorphic humanoid robots, on the other hand, could have as many 60 degrees of freedom to control and coordinate. At the other extreme, multi-agent and swarm robotics, while physically decoupled, require the coordination of a few to thousands of agents.

While in the past, control and planning were considered separate problems, modern control and motion planning need to be addressed in unison. Efficient planning methods that consider low-level controllers of their agents (whether arms, rovers, drones, etc.) and their tasks (manipulation, traversal, flight, etc.) will use new techniques from mathematical topology and recent sampling based planning methods to effectively search the relevant high-dimensional spaces that define their environments and interactions.

7.6.1. Task and Motion Planning Under Uncertainty

Robots use sensors to observe the environment and situate themselves in it, and then plan actions to attain a goal configuration. Due to the lack precise sensors, algorithms must be designed so that robots operate safely and robustly in the presence of uncertainty. While progress has been made in recent years, current methods can only handle simple tasks in fairly structured environments. More research is needed to develop algorithms for planning in belief space that can handle realistic problems in unstructured environments. These methods must be capable of real time operation in close proximity to and cooperation with humans. They need to provide safety and robustness guarantees while accommodating incomplete, inaccurate, and intermittent sensor data. Finally, although they have traditionally been studied separately, a principled integration of task and motion planning incorporating uncertainty is required to reach the level of autonomy needed for robots to become useful partners in unstructured settings such as the home.

7.6.2. From Specification to Deployment

Control design relies heavily on idealized, physical models that are typically expressed in terms of differential equations. The resulting controllers then have to be implemented on finite precision, discrete time computational substrates and deployed in real environments with noisy sensors and actuators. What is needed to significantly accelerate the design cycle are tools for automatically bridging these theory-practice gaps in a correct manner, which involves formal methods, hybrid computational models, and real-time adaptation of the control protocols to the dynamically evolving, real environments.

7.6.3. Control and Planning in Constrained Environments

Constraints to robot control planning present themselves in many different forms, whether they are physical constraints on a robot's reach, obstacles that constrain their workspaces, force constraints when interacting with sensitive materials, power/resource constraints of the robot, or dynamical constraints limiting robot actuation. Currently, efforts in constrained optimization approaches have been used to demonstrate effective optimization of these tasks for short tasks and small motions, in static environments with reasonable certainty. However, situations such as in surgery, service, and manufacturing involve long durations, fluid sequences of tasks, and dynamic environments. The next work in constrained optimization for robotics will be to roll the constrained tasks effectively into planning algorithms that provide continuous and connected motions, that can anticipate and react to dynamic constraints, and do so over long periods of time where its performance remains stable.

7.6.4. Manipulation

Manipulation and grasping are fundamental capabilities for operating in the physical world – they are needed to open and close doors and drawers, to pick up, move or push objects, to use tools, to manipulate a steering wheel, or to otherwise reconfigure or interact with the environment. Current algorithms can handle only relatively simple scenarios, such as low degree of freedom problems with small, regular geometries and quasi-static motion. Research is needed to develop grasp planning and metrics for complex and unique geometries. Improved techniques are needed for contact tasks, for manipulating deformable objects, for non-prehensile actions and tool use, and for dynamic motion. Strategies for robustness and failure detection and recovery are required for safe and secure operation.

7.6.5. Dynamic environments

Dynamic environments encompass manipulation tasks in sensitive environments, with humans or other robots, and moving obstacles for which the robot does not have explicit knowledge of their underlying motions. Currently, strides have been made in modeling dynamic environments on a small scale with one or a few objects and agents in the environment generally following known trajectories or having repetitive behaviors that can modeled using simple process models. Robots in these low-dimensional environments can effectively plan over long time horizons. Yet, a challenge that needs to be addressed is scalability (numerous, heterogeneous dynamic objects and agents) and uncertainty (complex or unpredictable dynamics) that may require re-planning and adaptation by robot systems in real-time.

7.6.6. Coordination Among Multiple Agents

Applications of multi-agent coordination (i.e. a set of robots that work together as a collective) appear in manufacturing and warehouse management, robot deployment for network coverage and disaster monitoring, construction robots, and others. Much of the work in multi-agent coordination has been inspired by nature. Evolutionary algorithms and decentralized intelligence have produced complex behaviors but are continuously challenged to converge toward optimal behaviors in short periods of time. Also, these behaviors have generally been applied to homogeneous agents, whereas realistically, deploying heterogeneous agents is more practical and add flexibility to the planning process. Centralized intelligence provides a means for bi-directional communication to and central planning for multiple agents; these methods are challenged to work in real-time, with single agents needing local controllers to react quickly or stabilize in the presence of unexpected events, and can be limited in hardware if each agent must communicate and coordinate between all other agents.

Future research will investigate real-time coordination methods that take advantage of a gradient between a centralized intelligence and local behavior, formal methods that prove convergence to some optimal behavior, and introduction in planning for heterogeneous agents for executing complex sequences of tasks with their neighboring robots.

	5 years	10 years	15 years
Integrated Task and Motion Planning under Uncertainty	Real-time algorithms for simple tasks in structured environments	Real-time algorithms for realistic tasks in structured environments or for simple tasks in unstructured environments	Real-time algorithms for realistic tasks in unstructured environments
Constrained Optimization	Integrated physical / system constraints into controllers and motion planners for short time horizons	Integrated physical / system constraints into controllers and motion planners for long time horizons	Commercial integration towards automated and optimized robot agents operating independently under known constraints.
Manipulation	Robust grasping of simple geometries in realistic scenarios, manipulation of simple deformable objects	Grasping of complex and unique geometries, simple tool use	Robust grasping and manipulation of complex and deformable objects and effective tool use
From Specification to Deployment	Computationally feasible models for correct by design deployed, real-time control code	Formal verification of applied safety-critical robotics applications	Broad commercial penetration of robotic solutions in safety-critical applications
Dynamic Environments	Modeling of dynamic processes beyond simple linear or repetitive models; control and planning around these models.	Estimation and planning for high dimensional spaces with many dynamic elements	Application to semantic models of dynamic objects and environments to complex robot behaviors such as non-prehensile tasks
Planning for multiple agents	Automatic generation of distributed control algorithms from high-level specifications.	Robustly deployed multi-robot systems in realistic environments.	Commercially available, user-friendly, modular multi-robot solutions.

7.7. Human Robot Interaction

Across many application domains, future robots are expected to work in human environments, side by side with people. Interactions between robots and their users will take many forms, from a trained factory operator supervising several manufacturing robots, to an older adult receiving care from a rehabilitation robot. The users will also vary substantially in background, training, physical and cognitive abilities, and readiness to adopt technology. Robotic products are expected to not only be intuitive, easy to use, and responsive to the needs and states of their users, but they must also be designed with these differences in mind, making *human-robot interaction* (HRI) a key area of research in the Roadmap. To cover the range of different human-robot interactions and the needs of specific application and use domains, this research area is divided into eight challenge areas:

7.7.1. Interface Design.

Robotic systems or products in some domains, such as search and rescue, space and underwater exploration, manufacturing, and defense, may present information to or require input from their users via computer interfaces. Human-robot interaction in these domains may be distal, with the robot and its user not co-located, and also distributed, with an operator potentially supervising many robots. Such uses will require interfaces that can effectively and intuitively present critical information to the user and provide clear and easily interpretable controls. User interfaces must also be scalable so that they accommodate increasing task complexity and human-robot team complexity. While computer interfaces have substantially matured in the last three decades, robot interfaces are still in their infancy. Extensive research is needed to develop principles and guidelines that are directly applicable to human-robot interfaces.

7.7.2. Perceiving, modeling, and adapting to humans.

Robotic systems and products must have the ability to perceive and recognize the behaviors, intent, and cognitive and affective states of their users. Endowing robots with such capabilities will require extensive research in order to develop sensing and perception techniques for human physiology and behavior, as well as models for interpreting what is perceived in the context of the task, application, and domain. Robots will need to be able to recognize and understand what they are perceiving, and also be able to estimate user intent in order to proactively plan their actions and contributions to collaborative interaction scenarios. Additionally, robots must be able to adapt to changes in user behavior and to variability across individuals, which will require the development of models of adaptive action and customizability.

7.7.3. Sociability.

Applications of robotic technology in service, health, manufacturing, and other domains where the robot plays the role of a social partner — companion, coach, teacher, peer, caretaker — draw on human social interaction as the primary paradigm for human-robot interaction. Socially interactive robots, therefore, must be designed with the capability to understand human social behavior as well

as display social behaviors that follow accepted norms of human interaction. Capabilities such as complex dialogue, the ability to interpret and generate rich nonverbal cues, and being able to understand and express emotion are key to achieving the required level of sociability. Research toward enabling these capabilities will develop detailed recognition, synthesis, and dialogue models for verbal and nonverbal interaction for generalized social skills as well as models that help the robot adapt to changing norms across contexts.

7.7.4. Collaborative Systems.

Robotic assistants in work environments, healthcare settings, and mission-critical contexts, such as robots designed for manufacturing, rehabilitation, and space exploration, will collaborate in human-robot teams to carry out assembly, rehabilitation, and exploration. In these roles and settings, robots must not only have detailed models of the tasks that they are engaged in, but must also be able to recognize, track, infer, and contribute to the task actions of their human counterparts. They must also employ these capabilities in dynamic settings, where task contingencies arise, task models are changed, and humans deviate from task plans or perform unexpectedly. Research into enabling robot systems to serve as effective collaborators must develop methods and models that will enable robots to perceive and interpret changing task environments and human counterparts as well as adaptively plan their contributions to the task.

7.7.5. Robot-mediated communication.

In addition to their roles as agents, assistants, and collaborators, robotic products will also serve as media for communication that facilitates interaction among people. Telepresence robots, for example, enable distant individuals to interact and communicate in substantially richer ways than it is possible with text, audio, and video. Pilot deployments of telepresence robots have already shown promise, enabling doctors to conduct patient rounds at a hospital and children with chronic conditions to attend classes. Research is needed to better integrate these systems into the environments they will inhabit, improve the naturalness of interfaces for users at both sides (remote operator and co-present individuals), and address issues of privacy, safety, and fair and equitable access and use. Additionally, autonomy and well-designed user interfaces must be an integral part of these technologies to facilitate greater mobility, smoother communication, and widespread adoption without taking control away from their primary users.

7.7.6. Shared Autonomy.

Introduction of autonomy into local operation or tele-operation to assist users in performing tasks is referred to by the research community as “shared control.” Shared control promises more precise, more capable, and sustained operator performance across a variety of domains. For example, shared control can enable a surgical robot to reduce tremor in transferring the surgeon’s operations, assist a paralytic patient perform sustenance actions, and a telepresence robot to navigate in a remote location without risking the safety or comfort of its occupants. In these scenarios, the robot may not have a model of its user’s goals and relies primarily on predictions of goals to modify direct control signals in a way that maximizes the success of the operation. Research in this area must develop method to enable robots to utilize a broad range of control signals from users, including input through a control interface to brain-computer interfaces and to facilitate varying levels of shared control in complex scenarios.

7.7.7. Long-term interaction.

In the US, the average ages of an automobile, a computer, and industrial machinery are 11.5, 5.6, and 10.5 years, respectively (numbers for 2014–2015). Robots are expected to have a lifetime similar to these products, providing assistance to and maintaining interactions with their users for approximately a decade. Unlike these products, however, robots will have the capability to learn about and adapt to the changing needs, behavioral patterns, and capabilities of their users over time. Very little research has been conducted on how robots can maintain such long-term interactions, learn new tasks, user models, interaction strategies, and transfer what they have learned about their users to their replacements. Extensive research, development, and longitudinal field testing is necessary to enable capabilities for long-term interaction and understand human expectations and responses to products and systems that are used over the course of a decade or potentially longer.

7.7.8. Safety.

The integration of robots into the human environment across several key domains will require not only that robots are designed with an acceptable level of intrinsic safety, but also that appropriate standards, legal guidelines, and social norms regarding their use by humans are developed. The development of standards such as ISO 10218.6 and ISO/TS 15066:2016 can guide industry regarding appropriate levels of safety for robots functioning in human environments. Extensive research must be carried out, however, to develop design guidelines for what norms robots must follow, for example, in navigating in human environments, interacting with their users, and resolving conflict (e.g., in order to determine right of way, liability, and so on) with other robots and users. While the development of the legal framework in which human interaction with robots is considered is beyond the scope of robotics research, the research community can inform the development of appropriate standards and guidelines for user expectations, making social norms an integral part of the design of robotic products and systems.

7.7.9. Research Roadmap

Challenge	5 years	10 years	15 years
<i>Interface design</i>	Basic generalizable standards and design guidelines for robot interfaces developed.	Accessibility standards and principles for scalability (different users, task complexity, long-term interaction) developed.	Authoring and programming tools for robot interfaces support standards, scalability objectives, and address safety, privacy, and security.
<i>Perceiving, modeling, and adapting to humans</i>	Robots can recognize and adapt their actions to basic human behavior, task actions, and intent in explicitly modeled tasks and environments.	Robots can learn and update user models on the fly and handle perception, modeling, and adaptation in semi-structured tasks and environments.	Robots can perceive, model and adapt to complex user behaviors, actions, and intent in semi-structured tasks and environments, and transfer learned models across domains and environments.

Challenge	5 years	10 years	15 years
<i>Sociability</i>	Robots maintain basic dialogue, recognize basic user states, and express internal states in controlled and semi-controlled settings (e.g., a classroom, rehabilitation center).	Robots can handle contingencies in dialogue, adapt to user states, and seamlessly integrate social behaviors in semi-controlled settings (e.g., the lobby of a corporate building).	Robots can adapt their interactions to diverse groups of users using dialogue and social interaction strategies in uncontrolled settings (e.g., public spaces, field and disaster settings).
<i>Collaborative systems</i>	Robots serve as effective collaborators in explicitly allocated plans, handling necessary perception, manipulation, and communication.	Robots can learn and infer task models and contribute to them on the fly, adapt to task contingencies (e.g., missing tools), and recognize and adapt their actions to changing task-relevant objects and environments.	Robots can not only recognize but also predict contingencies, user error, and changing capabilities of human collaborators and take action toward preventing or minimizing their effects.
<i>Shared autonomy</i>	Shared autonomy systems can integrate multiple explicit forms of user input to perform user-guided actions and provide users with flexibility in choosing from among different levels of autonomy.	Shared autonomy systems can recognize implicit user input to determine user goals and follow an active learning approach to engage the user in supporting these goals and choosing the level of autonomy required for the performance.	Shared autonomy systems can integrate and fuse various forms of implicit and explicit user input, model user goals and error on the fly, and vary levels of autonomy as necessary while communicating with their users.
<i>Long-term interaction</i>	Robots will maintain interactions with, learn from, and adapt to their users in the timeframe of months in semi-controlled settings.	Robots will function in semi-controlled settings in the timeframe of a year and accommodate the needs of multiple users with varying degrees of capability.	Robots will maintain adaptive functionality in uncontrolled environments in the timeframe of several-years with an arbitrary number of users with varying degrees of capability.

Challenge	5 years	10 years	15 years
<i>Safety</i>	Safety standards will be detailed to address domain-specific risks, and their applicability and effectiveness to deployed robot systems will be tested.	Norms (user expectations beyond inherent safety) regarding how robots should function in the human environment will be developed, applied to deployed robot systems, and tested for effectiveness.	Norms and standards will be refined based on an understanding of long-term use and interaction with deployed systems toward seamlessly integrating robots into society and enabling safe, effective, and acceptable robot functioning in uncontrolled environments with diverse groups of users.

7.8 Multi-Agent Robotics

From a robustness and effectiveness vantage-point, there is strength in numbers, i.e., in a number of applications, being able to deploy teams of robots, as opposed to a single robot, has distinct advantages. If a robot breaks down, the mission will still continue, by spreading robots over a large spatial domain, more effective coverage is achieved, or by distributing capabilities across multiple platforms, more flexible and adaptive robot systems are produced.

Applications of multi-robot systems appear in manufacturing and warehouse management, robot deployment for network coverage and disaster monitoring, construction robots, agricultural robotics, and others. Much of the work in multi-agent coordination has been inspired by nature. Evolutionary algorithms and decentralized intelligence have produced complex behaviors, e.g., for making teams of robots assemble geometric shapes, covering areas, tracking boundary cubes, or finding and tracking intruders. Despite these recent advances, a number of research issues remain to be solved for large teams of robots to be robustly deployed in real environments over sustained periods of time.

7.8.1. Distributed Control and Decision Making

Fundamentally, what makes multi-agent robotics challenging is that the individual agents may be ill-informed in the sense that the information they have access to is limited to what they can measure themselves and what neighboring agents may share. As such, distributed decision making algorithms must be developed that take this limited information into account in such a way that the desired

global behaviors emerge from a collection of local rules. Although a number of anecdotal examples of such mechanisms are present in the literature, we still lack an effective framework for going from high-level specifications of what the team should be doing, to low-level, distributed control algorithms. The development of such a framework is key for multi-robot systems to become truly flexible and adaptive tools across a number of application domains.

7.8.2. Mixed Decentralized/Centralized Information Exchange Mechanisms

Traditionally, small teams (up to three or so) of robots have been controlled by a centralized decision maker while larger teams are acting truly decentralized in that only locally available information is available. However, the reality is oftentimes a mix of these two extremes, where a centralized, cloud-like node may gather information over time and intermittently inject centralized information into the system. As of yet, there are no effective abstractions or systematic algorithms for describing or taking advantage of hybrid, centralized/decentralized information exchange mechanisms where the information flows may be operating at different time scales. Subsequently, no clear understanding exists for what information needs to be shared in a centralized manner and what information can be kept local.

7.8.3. Human-Swarm Interactions

Although HRI (Human-Robot Interactions) has matured as a discipline, the question of how a human operator should be interacting with teams of robots is still largely unresolved. Consider for example a farmer trying to interact with a team of small, autonomous tractors or a pilot trying to control a large swarm of unmanned aerial vehicles. In these two scenarios, it is fundamentally not understood what constitutes effective interaction modalities, both from a cognitive workload and from a bandwidth management perspective. As we are expected to see an increasing number of robots entering our factories, homes, hospitals, and farms, these Human-Swarm Interaction questions must be addressed for people to be able to effectively use the robots to their fullest extent.

7.8.4. Heterogeneous Networks

One of the benefits of deploying teams of mobile robots is that one can distribute capabilities across the different robots. Some robots may be able to fly, thereby providing a high-level overview of the environment, while others may be able to slither through and inspect narrow tubes. The resulting system is a heterogeneous team, where heterogeneity may refer to dynamical configurations, sensing capabilities, spatial footprints, or behavioral strategies. But, beyond isolated examples (mainly air-

ground coordination), we do not fully understand how to take advantage of heterogeneity in a fundamentally sound manner. What types of robots are needed, given a particular task? How heterogeneous is a collection of robots? How heterogeneous should a team be to provide maximal flexibility in terms of the tasks it can perform?

7.8.5. Communication and Sensing in Multi-Robot Systems

In order to realize the full potentials of a robotic network, we need a foundational understanding of the interplay between sensing, communications and navigation in these systems. For instance, communication between the agents or to the outside nodes can be degraded due to several factors such as path loss, shadowing or multipath fading. Furthermore, path planning at every agent not only affects its information gathering (sensing) but also impacts its connectivity maintenance (communications). This multi-disciplinary nature makes designing robust decision-making strategies for successful task accomplishment in robotic networks considerably challenging and an open problem, when considering realistic communication links, resulting in the emerging research area of communication-aware robotics.

7.8.6. Roadmap for multi-agent robotics

	5 Years	10 Years	15 Years
Distributed Control and Decision Making	Automatic Generation of Distributed Protocols from High-Level, Global Specifications	Robustly deployed large-scale teams in real environments	Commercial penetration across a number of industries, including agriculture, manufacturing, warehousing, and environmental monitoring
Mixed Decentralized/Centralized Information Exchange Mechanisms	Foundational understanding of how to distributed information requirements across centralized and decentralized channels.	Federations of multi-robot teams learning and collaborating with each other via cloud-based architectures	Internet-of-Robotic-Things

Human-Swarm Interactions	Rich examples of effective human-swarm interaction modalities	Commercially available, human-centric swarm systems	Swarm-robotics as a user-friendly service
Heterogeneous Networks	Characterizations of heterogeneity across a number of dimensions, including functional, spatial, and temporal	Heterogeneous solutions to complex tasks, as opposed to a priori assumptions about team compositions	Robustly deployed, mix-and-match collaborative robotic teams
Communication and Sensing in Multi-Robot Systems	Effective models for trading off mobility, sensing, and communications	Automatic co-optimization of power consumption across sensing, mobility, and communications modalities	Low-cost, widely available multi-robot systems.

8. Workforce Development

8.1. Introduction

Popular news reporting has helped create the impression that robotics technologies are proceeding at a rapid rate and in many industries will increasingly replace human workers. As with the entire history of industrial manufacturing and automation, there will continue to be shifts from unskilled to skilled labor. However, in reality the expanded use of robotics and automation will likely create many off setting skilled labor and engineering types of jobs. The workforce challenges that face all sizes of industry, but especially small to medium size enterprises, are complex and have limited the successful and cost effective utilization of robotics and automation. Applications demand extensive use of engineers and specialized technicians to implement, start up, and sustain such operations. Unique combinations of embedded computing, software, and electronics skills, to name a few, are required and are generally expensive and in short supply. New ideas and programs are needed to address this the future workforce in this sector.

8.2. Strategic Findings

From their early beginnings as tele-operators and manipulators, robotic systems have served to *extend the reach of humans* in interacting, manipulating and transforming the world around us. Since then, the enormous growth in numbers, diversity and complexity has been driven by their usefulness in enhancing human manipulation capabilities over various spatial and temporal scales (nano to mega) and for automation of the 4D (dull, dirty, dangerous and dumb) tasks. At the same time, the applications-oriented field of robotics and intelligent machines has offered a means of tangible embodiment of ideas and algorithms for a host of scientific disciplines – system design, control engineering, computational science, and artificial intelligence – among others. The diversity of application arenas stands as testament to its interdisciplinary nature and ultimately its immense potential – however setting the scope of robotics activities and developing forward-looking roadmaps for the field is a useful exercise, especially given the blurring of the boundaries between robotics and its constituent scientific disciplines.

These drivers are already revolutionizing the robotics landscape with attention being focused on developing all facets of the research, development, educational and ultimately logistics and commercial infrastructure to support this enterprise. And the commercial industry-led driving impetus behind these ventures bodes well for the long-term success of these efforts.

However, in the midst of these developments, another revolution has been silently underway that is fundamentally transforming the landscape of robotics. The archetypical robotic system is being transformed into a system-of-systems created as heterogeneous collections of physical and information resources coupled together by intricate connections and interactions. The enablers of this revolution are the same science and technology drivers that have made robotic devices smaller, smarter, easier to use and more connected with each other, with people, and with the environment. None of these changes are remarkable in of themselves but the net effect has been to enable new capabilities and remove barriers in ways that previously inconceivable. The ramifications of this transformation are so immense that to (mis)quote Thomas Friedman [1]: “Overnight, while the world slept the robotics field has been flattened”.

Over the past 5 years, Bill Gates’s vision of a “robot in every home” paving the way emergence of a new robotics industries with even greater potential to revolutionize the way we live. There are the striking parallels between the personal-computer and the personal-robot industries in their early years – in terms of the fragmented state-of-existence (diversity of platforms/software), the inflexible operational paradigms (monolithic solutions) and the newer hardware and software trends (modularity, open-source) that paved the way for the revolution. And, as in the personal-computer industry, the evolutionary pressures from *evolving application focus, rapid technological/scientific progress* and *technological crosspollination* are driving and constantly reshaping the landscape.

- First, the archetypical PUMA manipulators, central to the manufacturing-floor automation of the heavy industries, could be viewed as the equivalent of the mainframes of the past era. Today, the growth in robotic systems is focused more the non-manufacturing application arenas and principally in the service robotics sector. Even here, the high-cost specialized devices from computer-assisted surgeries (da Vinci comes to mind), space explorations (NASA Mars Rover), and military robots in hostile combat environments (disposal of roadside bombs in Iraq), and robots to assist the search of trapped miners, form only a small fraction principally due to low-volumes. The most significant growth comes from the low-cost, high-market volume domestic and personal robotics market.
- Second, technological advances in sensing/actuation/computing on one hand and improved fundamental scientific understanding and algorithmic implementation on the other have contributed significantly to the growth of robotic systems of various shapes, sizes and functionality. Modularity and standardization in hardware, software and tools and the coupling of commercial interest with open-source movement is beginning to reshape the robotics arena much in the same way that personal-computing has been transformed.
- Last, the technological cross-pollination that occurs with each new round of innovation, improves not only existing robotic systems but opens up other avenues where intelligent mobile robots can be employed, effectively creating new markets. For example, the student developing advances for a robotic unmanned ground vehicle could go on to develop your neighbor’s robotic lawn mower. Or for that matter, the years of research on safe and stable teleoperation can serves to enhance driving feel in tomorrow’s drive-by-wire automobile!!

8.3. Near Term Opportunities and Factors Affecting Immediate Deployments

The rapid increase in the number of formal undergraduate and graduate degree programs in robotics in recent years motivates the need to develop a model robotics curriculum. Such a curriculum consisting of unified courses and projects utilizing standard robotics software and hardware will accelerate the creation of robotics programs to support the ever-increasing demand from the industry for engineers with multidisciplinary skills.

In the past decade, major attention has been given to STEM initiatives seeking to attract K-12 students into all science and engineering areas, especially minorities and women. Typical initiatives seek to make K-12 students aware of the wide range of STEM opportunities, and often use “robotics” as the attention getting focus. Such efforts are critically important but they will not necessarily produce the results needed to assure the skilled labor force that will enable the expanded use of increasingly complex robotics and automation solutions. Some are beginning to talk about “K-to-Gray”, bringing attention to need to consider the entire “life cycle” of workforce issues. New ideas and programs are needed to attract younger workers to the opportunities in robotics and

automation, and to seek ways to better transfer and couple the knowledge base of experienced workers into the emerging workforce – this need applies to all aspects of the robotics and automation future workforce from technicians to engineers and computer scientists. Specifically, new ideas and programs are needed to integrate the engineering and skilled labor training domains so that future technician training keeps track with the rapidly advancing use of intelligent systems and robotics. Better communications and collaborations are needed between professional societies, universities and community colleges are needed to assure that skilled technicians are being trained in critical areas such as mechatronics and embedded computer controls. Such interactions exist in many instances, but in general the depth, comprehensiveness, and real-time integration must go to another level

8.4. Contributors

The discussions at the respective workshops was summarized by an editorial team

Bill Hamel
University of Tennessee

Nancy Amato
Texas A&M University

Venkat Krovi
Clemson University

Taskin Padir
Northestern University

9. Shared Infrastructure in Robotics

As robotics continues to expand to more application domains, the development and maintenance of suitable experimental facilities are becoming bottlenecks in the innovation process. In fact, there is a significant gap between the theoretical foundations that are being broadly pursued, and a focused, application-driven transition from small-scale experiments to robust and high impact deployments. This gap is both scientific and practical. By having researchers from different institutions, disciplines, and backgrounds come together around a common testbed, there is potential to accelerate innovation and to build on past findings in a more effective manner than what is currently done. The development and maintenance of meaningful, large-scale robotic testbeds is a resource-intensive undertaking, which is why it is particularly well-suited to a shared and even remote-access format.

Some specialized robot testbeds exist (e.g., UMass Lowell's NERVE Center, the Southwest Research Institute's Small Robotic Vehicle Evaluation and Applications Group, courses for response robots and manufacturing robots at NIST in Gaithersburg, MD, and Texas A&M's Disaster City), but these testbeds lack a variety of robot systems; researchers must bring their own robot systems, limiting the ability to test the generalizability of algorithms and preventing people without robot systems from being able to test their theoretical results in the real world. Georgia Tech's Robotarium, currently in development, will result in a robot testbed for swarm robotics, with both testing environments and robot hardware. In order to accelerate the development and effective testing of robot systems, shared community resources of testbeds for a variety of application domains with a variety of robot systems must be developed throughout the country, each with a particular application focus (e.g., agriculture, marine, manufacturing, medical). To maximize the use of available resources, existing facilities could be expanded to create a comprehensive shared infrastructure (i.e., one with both testbeds and shared robot systems) while developing testbeds for application domains where no such facilities yet exist.

9.1. Flexible Research Platforms

In order to be a truly useful remote-access research testbed, it is vitally important that the testbed is structured in such a manner that it allows for a number of different research questions and experiments to be pursued. Moreover, the testbed itself must evolve over time to remain relevant to the changing research trends and directions. It must be possible to automatically specify experimental setups and scenarios, which calls for research to be done on modular interoperable hardware (plug-n-play etc.) as well as software (robotic experiment description languages etc.) to facilitate their inclusion into larger eco-systems and downstream commercialization. Given that this infrastructure is a community resource, each facility should have a user committee to allocate site usage and suggest facility updates. While each facility will have a specific application focus, there should be collaboration between these facilities in order to prevent duplicating efforts and to share best practices.

9.2. Community Consensus Validation Benchmark Frameworks

Various research groups have developed in-house methods for quantitative performance assessments of both robotic systems and the human-robot interaction. Individual groups have begun the process to collect and share their data sets and best-practices. However, efforts remain fragmented and disconnected due to the lack of realistic and relevant test environments (physical and virtual benchmarks). A multipronged validation regimen (e.g. supporting both virtual and physical testing; staged evaluation of components, subsystems and systems; device vs. user testing) is crucial. The

development of such frameworks for open-access creation, collection and curation of the appropriate reference environments and data-corporuses against which quantitative performance can be assessed would significantly speed the process of technology development as well as transfer. Past efforts in the robotics community have strengthened the argument that potentially posing these as a competition or grand challenge could help focus the energy of both the academic and industrial communities [Citation], while also opening doors for subsequent standardization efforts. Some robotics domains have already been developing standard test methods and metrics through ASTM, IEEE and ASME; these efforts should continue while efforts in new application domains begin. The Robotics-VO could provide oversight for community discussions to develop these shared resources.

9.3. Reference Open-Access Testbeds

The enormous growth in the field has created an explosion in the number and variants of the solutions presented. For example, the range of manipulator arms, mobile bases and grippers commercially available can be mind-boggling. It is becoming increasingly difficult for a researcher working on grasping algorithms to obtain access to a variety of manipulators with different types of grippers in order to evaluate the effectiveness of their algorithms. However, the lack of access to truly industry-grade test-beds with interoperable hardware and software modules is beginning to impede innovation. Efforts at creating open-source platforms are underway and represent a good starting point. A broad and inclusive program needs to be supported through roadmapping workshops and study-groups to facilitate development of open-source community-vetted standards. A good example in the robotics software arena is the ROS Framework [Citation]. The accessibility to such plug and play frameworks will let research groups focus on their subtopics while still contributing to a broader coherent community effort. Additionally, system interoperability and synergistic technical tools (e.g. programming, hardware, communication) are critical and will benefit academia and industry alike for hastening robotic system research and development.

	5 Years	10 Years
Flexible Research Platforms	Coordination framework among academic researchers to create modular shareable hardware (CAD repositories) and software (APIs) for the next-gen robotic systems (e.g. wearables, soft-suits etc.)	
Community Consensus Validation Benchmark Frameworks	Coordination with Professional Societies (IEEE, ASME) to host competitions and roadmapping workshops	Involvement of Standards organizations (IEEE, ISO, ANSI, ASME)
Reference Open-Access Testbeds	Robot testbeds spread across the nation. Standup a coordination framework (perhaps like Robotics VO or CPS VO)	Transition of viable hardware/software to pre-competitive TRL hardening via NNMI

9.4. Contributors

The issue of shared infrastructure was discussed extensively at the two workshops and the main observations were summarized by an editorial team

Magnus Egerstedt
Georgia Institute of Tech.
Henrik I. Christensen
UC San Diego

Venkat Krovi
Clemson University

Holly Yanco
Univ. of Mass. - Lowell

10 Legal, Ethical, and Economic Context

The Roadmap to Robotics is primarily a technical document. Its central purpose is to describe the present and anticipated state of the art in robotics in the United States and to help the American government set levels and priorities for support.

It is clear to the authors, however, that the development of robotics in the United States and elsewhere takes place against a backdrop of law, policy, ethics, and economics—among other social, cultural, and political forces. The purpose of the following chapter is to acknowledge this broader context. The chapter raises some of the more pressing non-technical challenges for robotics and directs the reader’s attention to ongoing efforts and resources to address these issues, where such efforts exist.

This chapter is not meant to be comprehensive, nor does it purport to articulate a consensus in the legal, policy, ethical or other communities as to what official policy toward robotics should be. Rather, we aim to raise certain key challenges that have repeatedly surfaced in the literature, in workshops, and in public discourse.² In addition, we articulate our commitment as a community to participate in and support this dialogue, which is by necessity deeply interdisciplinary, as well as to recommend that government and academia work to actively remove barriers interrogating robotics’ broader societal context.

The remainder of this chapter consists of short discussion of key issues followed by our recommendations.

10.1. Safety

Robots have to be safe. But how safe is safe enough? There are many possible configurations, but a key role for government is to help set the safety thresholds or standards for a variety of robotic systems with the capacity to do physical harm to people or property out in the world. Thus, the Federal Aviation Administration will have to set safety thresholds for delivery of goods using unmanned aerial systems and the National Highway Transportation Safety Administration will have to set expectations around autonomous vehicles. Having set these thresholds or standards, techniques are then needed to test and validate that they are being met.

Special considerations may arise where robots are performing task usually performed by people with specialized training. Each profession that today certifies its own professionals will need to confront whether and how their standards can be translated into technical systems performing comparable

² These include the National Science Foundation and Department of Homeland Security *Policy for Automation* workshop, the *Future of AI: Opportunities and Challenges*, the White House series *Preparing for the Future of Artificial Intelligence*, the Stanford University *AI 100* inaugural report, and the annual robotics law and policy conference *We Robot*, with more efforts in progress and on the horizon. Several of the authors of this report have participated in these and other efforts to identify and address the legal, policy, economic, and ethical issues robotics and AI may present. These efforts are focused on the United States; there are, if anything, more and longer standing efforts abroad, especially Europe (Italy, the UK, Germany), Japan, and South Korea.

tasks. While it may not make sense to give autonomous vehicles driving tests, clearly the medical profession will need to sign off on robots that eventually perform surgery autonomously.

For robotics to remain as safe and accountable as possible, there should also be a role for independent researchers. Academics and others may be in a position to help determine if systems are behaving or will behave as intended. To secure their participation, however, independent researchers need to know exactly what sorts of activities the law permits. The concern is that existing law—such as the Computer Fraud and Abuse Act, which disallows unauthorized access to many technical systems, and the anti-circumvention provisions of the Digital Millennium Copyright Act—may be read to prohibit research activities that ultimately serve the goals of public safety.³ Lawmakers and enforcement agencies should clarify that, for instance, reverse engineering or otherwise examining software or hardware for the purpose of assessing its safety is permitted under all relevant laws.

10.2. Liability

Wise investment in robotics is likely to mean continued gains in public safety. Robots can perform inherently dangerous tasks, for instance, and perform risky tasks with greater precision. But robots, like humans, may find themselves in situations where harm is unavoidable. Courts and perhaps lawmakers will need to establish liability rules by which to compensate victims of robot-related hazards while preserving incentives for innovation.

Consider, for example, a home robot built by one company that injures a person while running software the robot's owner purchased from another company through a robot app store. From the victim's perspective, a robot built by a company with deep pockets caused an injury. But is it wise or fair to hold the manufacturer of a robot that—like a computer, tablet, or smartphone—is open to third party innovation by design?⁴

Or consider a robot that, alone or in interaction with other systems, causes a kind of harm no one could reasonably anticipate. It should be clear, for example, that the manufacturer of a fully autonomous vehicle will be liable where that vehicle causes a traffic accident by turning without signaling. But now imagine an autonomous vehicle designed to find ways to maximize fuel efficiency through experimentation. The system might perform functions—such as running its engine in an enclosed garage to recharge its battery—that no one intended or anticipated, but which end up causing serious harm. Such events could pose a challenge to tort law, which is premised on the notion that courts should only compensate injuries that are intended or foreseeable.⁵

Closely related to the task of determining liability is understanding the role of insurance. Market forces seems already to be responding to the acceleration of robotics; companies whose business models rely upon the use of robots are better able to secure insurance than they were a decade ago. But there also may be a role for government. The widespread availability of autonomous vehicles,

³ See Peter Stone et al., *Artificial Intelligence and Life in 2030* (Sept. 2106), at 43.

⁴ See Ryan Calo, Open Robotics, 70 *Maryland Law Review* 571 (2011).

⁵ See Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 *California Law Review* 513 (2015).

for example, may present the need to revisit the utility of no-fault insurance, which has been declining in popularity among state lawmakers.⁶

10.3. Impact on Labor

Many commentators have articulated a concern that the risks associated with the use of artificial intelligence to make decisions about consumers and citizens will fall disproportionately upon the vulnerable, i.e., those in society with the least capacity to mitigate technology's effects.⁷ These concerns also apply to robotics. For example, we might worry that greater reliance on robotics by police would disproportionately affect, and come to further alienate, low income or minorities communities.⁸

The prospect that robots may take low or high skill jobs is of particular concern to the public. The concern is sometimes overstated. Even a simple analysis reveals that robots will both displace and create jobs at an individual level. There will be diminished need for captains, pilots, or truck drivers if companies automate long-haul transportation. At the same time, the burgeoning unmanned aerial vehicle or drone industry is already hiring. There is also evidence that automation in manufacturing has, to date, correlated with job *creation* in the United States.⁹ Finally, many analyses of robotics' impact on labor all but ignore the extensive and growing area of human augmentation. In contrast to automation, augmentation aims to enhance human abilities and create collaborations with machines, so that people are empowered, not replaced. Examples include rehabilitation robotics, socially assistive robotics, and collaborative robotics, to name a few.

Nevertheless, greater reliance on robotics is likely to have impacts in the short, medium, and long term and that must be managed responsibly.¹⁰ One solution that has been advanced in response to the prospect of widespread automation of jobs is the idea of a universal income, i.e., a basic income for every American subsidized by the gains in productivity and efficiency from automation. Another variant recommends imposing an obligation on employer to pay for retraining of workers displaced by robots.

There are several challenges around universal income, including that it may not be politically palatable, that income guarantees do not resolve other issues—such as idleness or inequity—that follow from unemployment, and that robotics might never so thoroughly transform our economy to permit redistribution of wealth on this scale.¹¹ A requirement that firms provide or subsidize

⁶ See Nora Freeman Engstrom, *An Alternative Explanation for No-Fault's "Demise,"* 61 DePaul Law Review 303 (2012).

⁷ See Kate Crawford and Ryan Calo, *There is a blind spot in AI research,* Nature 538 (Oct. 20, 2016).

⁸ See Elizabeth Joh, *Policing Police Robots,* UCLA Law Review Discourse (forthcoming 2016).

⁹ See International Federation of Robotics, *Positive Impact of Industrial Robots on Employment* (Jan. 2013).

¹⁰ See Andrew McAfee and Erik Brynjolfsson, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (2014).

¹¹ For an early argument that robots are unlikely to cause massive economic change, see Herbert Simon, *The Shape of Automation for Men and Management* (1977).

retraining also brings challenges, such as the potential disincentive to adopt robots where doing so would constitute a net gain for productivity, safety, or both.

We are ultimately hopeful as a community that, if handled well at the level of policy, advances in robotics are likely to improve the overall health, resilience, and well-being of American society.

10.4. Social Interaction

An extensive literature evidences the ways in which people tend to react to anthropomorphic technology such as robots as though the robot were a social entity.¹² Designers of personal and service robots are well aware of this tendency and many have made considerable efforts to ensure a positive and respectful interaction between people and robots. Commentators worry, however, that the propensity people have to form social bonds with robots will prove problematic. Sherry Turkle and others have expressed concern that robotic interaction will substitute for far richer interpersonal relationships, as when an elderly relative is left in the care of a home robot.¹³ Others worry that anthropomorphic robots will be capable of exploiting our social reactions to nudge us toward corporate or other goals at odds with our own.¹⁴

There was widespread agreement among authors—many of whom work in the field of Human-Robot Interaction (HRI)—that both the positive and negative effects of robots on people need to be carefully researched and considered. Indeed, a limitation of current funding models is that too few resources are directed specifically toward studying social impact as itself a technical challenge of robotics. Participants further suggested the establishment of one or more testing facilities designed to emulate the real world. These would consist of instrumented environments where researchers can study human-robot interaction and compare share and compare results. Models for such robot spaces already exist in the United States and abroad.¹⁵

10.5. Privacy and Security

Closely related to the social interaction concern is the set of privacy and security challenges robots inevitably raise. Ryan Calo has argued that robots raise at least three categories of privacy issues: (1) robots make it easier to engage in surveillance, as when police use drones to monitor a protest; (2) robots create access to spaces historically reserved for solitude, as when government or black hat hackers compromise a home robot; and (3) as alluded to above, anthropomorphic technology such as robots occasion in people the perception of being observed. The potential for security vulnerabilities is rendered more acute by the prospect that a compromised robot could cause

¹² E.g., Byron Reeves and Cliff Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places* (1996).

¹³ See Sherry Turkle, *Alone Together: Why We Expect More from Technology and Less from Each Other* (2011).

¹⁴ E.g., Illah Reza Nourbakhsh, *Robot Futures* (2013); Woodrow Hartzog, *Unfair and Deceptive Robots*, 74 Maryland Law Review 785 (2015).

¹⁵ For example, the University of Michigan has a fake town where researchers test driverless cars and the Federal Aviation Administration has designed certain areas of the country for unmanned aerial vehicle testing.

physical harm. Today a number of groups within civil service, industry, academia, and government are working to address some of these privacy and security issues.¹⁶ Government is in a position to better support this important work going forward, including by removing barriers to research.

10.6. Recommendations

To reiterate: this document is primarily a technical roadmap. Its central purpose is to update Congress on the state of the art in robotics and to help policymakers determine where to channel resources in order to realize robotics' great promise as a technology. Robotics develops against the background of a legal, policy, ethical, economics, and social context. This chapter has identified some of the challenges that recur in ongoing discussion of that context.

With this in mind, we conclude by tentatively offering a handful of recommendations aimed at preserving, fostering, and expanding the discussion of how robotics interact with society:

- *Greater expertise in government.* In order to foster innovation in robotics, maximize its potential for social good, and minimize its potential for harm, government at all levels should continue to accrue expertise in cyber-physical systems.
- *Support of interdisciplinary research in government and academia.* Few issues in robotics, or any other context, are amenable to resolution by reference to any one discipline. Government and academia should actively work to support and incentivize interdisciplinary research and breakdown siloes between expertise.
- *Removal of research barriers.* As alluded to above, independent researchers should be assured that efforts to understand and validate systems for the purpose of accountability and safety do not carry legal risk under existing law or doctrine.

10.7. Contributors

The discussions from the workshops was summarized and edited by Ryan Calo, University of Washington.

¹⁶ One example is the National Institute on Standards and Technology working group around privacy for unmanned aerial systems.