March 20, 2013

A Roadmap for U.S. Robotics
From Internet to Robotics
2013 Edition

Organized by

Georgia Institute of Technology
Carnegie Mellon University
Robotics Technology Consortium
University of Pennsylvania
University of Southern California
Stanford University
University of California–Berkeley
University of Washington
Massachusetts Institute of Technology

Sponsored by

Robotics in the United States of America
Robotics - VO

Table of Contents

Overview

Chapter 1

Chapter 2

Chapter 4

Roadmap for Robot Applications in Space...89

Chapter 5

Roadmap for Robot Applications in Defense 105

Overview

Robotics as a Key Economic Enabler

Last year, robotics celebrated its 50-year anniversary in terms of deployment of the first industrial robot at a manufacturing site. Since then, significant progress has been achieved. Robots are being used across the various domains of manufacturing, services, healthcare/medical, defense, and space. Robotics was initially introduced for dirty, dull, and dangerous tasks. Today, robotics are used in a much wider set of applications, and a key factor is to empower people in their daily lives across work, leisure, and domestic tasks. Three factors drive the adoption of robots: i) improved productivity in the increasingly competitive international environment; ii) improved quality of life in the presence of a significantly aging society; and iii) removing first responders and soldiers from the immediate danger/action. Economic growth, quality of life, and safety of our first responders continue to be key drivers for the adoption of robots.

Robotics is one of a few technologies that has the potential to have an impact that is as transformative as the Internet. Robotics is already now a key technology for inshoring of jobs by companies such as Apple, Lenovo, Tesla, Foxconn, and many others and citizens who used to have to rely on family or nurses for basic tasks such as shaving, preparing a meal, or going to the restroom are having a higher degree of independence. In the aftermath of the earthquake in Fukushima, it was evident that it would be a challenge to get an actual sense of the resulting destruction without the deployment of robots for assessment of the magnitude of the damage and assessment of the environmental impact. A similar use of robot systems was also demonstrated in the aftermath of the well blowout in the Gulf of Mexico.

To fully evaluate the potential of using robotics across the set of available applications, a group of more than 160 people came together in five workshops to identify: i) business/application drivers; ii) the current set of gaps to provide solutions to end-users; and iii) R&D priorities to enable delivery on the business drivers. The meetings were topical across manufacturing, healthcare/medical robotics, service robotics, defense, and space. The workshops took place during the second half of 2012. At each workshop, there was a mixture present of industry users, academic researchers, and government program managers to ensure a broader coverage of the topics discussed. Robotics is one of a few technologies capable of near-term building new companies, creating new jobs, and addressing a number of issues of national importance.

This report is a follow-up to the CCC-sponsored roadmap that was published in May 2009 and presented to the Congressional Caucus on Robotics on May 21, 2009. That roadmap subsequently led to the creation of the National Robotics Initiative (NRI), which is jointly sponsored by NSF, USDA, NASA, and NIH. The NRI was launched in 2011. The present roadmap is an update to the former document in the areas of manufacturing, healthcare/medical, and service robotics. In recognition of the important role that space and defense robotics has both to R&D but also as early adopters, new chapters were added for those areas. These new sections should primarily be seen as identifying areas with dual and multiple-use potential and areas with a clear potential for multi-domain coordination. As such, the space and defense sections are complementary to independent roadmaps developed by agencies within those domains. The update of the roadmap has been organized by the Robotics Virtual Organization

Roadmap Results: Summary of Major Findings

- Robotics technology holds the potential to transform the future of the country and is expected to become as ubiquitous over the next decades as computer technology is today.

- Through adoption of robots in flexible manufacturing, it is possible to generate production systems that are economically competitive to outsourcing to other countries with lower wages.

- A key driver in adopting robotics technology is the aging population that results in an aging workforce but it also poses a number of challenges to the healthcare system.

- Robotics technology has advanced sufficiently to allow for "human augmented" labor that enables acting on the vision of co-workers who assist people with dirty, dull, and dangerous tasks, and it facilitates a new generation of systems for domestic support to improve quality of life for the broader population. In addition, robots have already proven their value in removing first-responders and soldiers from immediate danger.

- Robotics technology offers a unique opportunity to invest in an area that has a real potential for new jobs, increased productivity, and to add to worker safety in the short-term. It will allow an acceleration of inshoring of jobs, and longer-term, will offer improved quality of life in a society that is expected to experience significant aging.

- Each of the areas covered by the roadmap identifies both near- and long-term applications of robotics technology, establishing 5-, 10-, and 15-year goals for critical capabilities required to provide such applications, and identifies the underlying technologies needed to enable these critical capabilities.

- While some critical capabilities and underlying technologies are domain-specific, the systems effort identified a number of critical capabilities that are common across the board, including robust 3-D perception, planning and navigation, human-like dexterous manipulation, intuitive human-robot interaction, and safe robot behavior.

Area Specific Conclusions

Manufacturing

The manufacturing sector represents 14% of the GDP and 11% of the total employment. Close to 70% of the net export from the U.S. is related to manufacturing. The sector represents a very important area to the general economic health of the country.

Over the last four decades, tremendous progress has been achieved in the use of robots, and in particular, in the automotive sector. More recently, the electronics sector has taken over as the dominant sector. The sale of robotics for manufacturing grew 44% during 2011, which is a clear indictor of the revitalization of the production system in the U.S. Robots have been used as a facilitator to inshore manufacturing for companies such as Apple, Lenovo, Samsung, and Foxconn. The use of robots is shifting from big companies such as GM, Ford, Boeing, and Lockheed Martin to small- and medium-sized enterprises to enable burst manufacturing for one-off products. The continued progress in this area relies on further progress in the area of integrated design, integration from design to manufacturing, new methods for cyber-physical systems integration, and a higher degree of computer-mediated manufacturing. Finally, there is also a need to educate a new generation of workers for the factory floor and to provide clear career paths for young people entering the field of manufacturing.

Medical Robots

Over the last few years, we have seen 40+% annual growth in the number of medical procedures performed using robots. They were initially introduced for use in cardiothoracic, gynecology, and urology. More recently, a broader set of procedures is pursued across orthopedics, neurology, and general surgery. Use of robots for surgery can reduce the number of complications by 80% and also allow a significant reduction in the time for hospitalization as well as a significantly faster return to the workforce and to a normal life. Due to the aging of society, the number of expected surgical procedures is expected to double over the next 15 years. More than 140,000 people in the U.S. are diagnosed with stroke every year. For stroke victims, it is important to receive early support for rehabilitation. Use of robots has been demonstrated to allow for faster and more complete recovery from stroke. It is essential to continue to develop and deploy robot systems for improvement in medical procedures and to reduce the overall cost of care.

Healthcare

More than 11 million people live with severe disabilities and need personal assistance. The cost of assistance from a certified nurse is $24K annually and a service dog typically costs $15K plus there is a significant wait. In addition, more than 5.4 million people live with dementia and are in need of cognitive support for their daily lives. Robots have recently been demonstrated both for rehabilitation and as a replacement for service dogs. The robots can provide significant support as part of daily life for mobility, for basic tasks such as getting out of bed, preparing a meal, personal hygiene, etc. Robotics technology allows for a significant improvement in quality of life and a reduction in the cost of support by allowing people to live independent for a longer period of time.

Service Applications

Robots are used both in professional and domestic service applications. More than 6 million autonomous vacuum cleaners have been purchased, and more than 200,000 autonomous lawn mowers are used worldwide. Robots have also been deployed for personal security applications. Professional service applications include inspection of power plants and infrastructure such as bridges. Service robots are also used in logistics applications such as delivery of beddings, meals, and pharmaceuticals at hospitals. The annual growth in professional service robots is 30%, and in domestic service applications, the growth is 20+%. U.S. companies have dominated this area, and it is considered important to maintain the momentum.

Space

Over the last decade, we have seen tremendous progress in science exploration of Mars through use of robotics systems such as Spirit and Opportunity. The systems have enabled extended missions on a faraway planet without deployment of astronauts. More recently, robots have also been deployed on the international space station to explore how some of the menial tasks can be performed by a robot in comparison to use of astronauts. Repetitive, high-precision, and extended tasks are all examples of where a robot may offer an advantage over use of humans. In addition, progress in these advanced applications offers important insight into how the same systems can be used in daily lives, which is one of the reasons that NASA and GM have teamed up to design and deploy the Robonaut system in the ISS. Going forward, there is no doubt that unmanned space exploration, will, in many ways, benefit from the use of robotics technologies to go where no man has gone.

Defense

At the height of the intervention in Iraq and Afghanistan, more than 25,000 robotics systems were deployed with a fairly even divide between ground and aerial systems. Unmanned aerial systems allow for extended missions, and the risk to the pilot is eliminated. Today, more than 50% of the pilots entering the Air Force become operators of remotely piloted systems rather than becoming regular airplane pilots. The opportunity for deployment in civilian airspace is explored through a new FAA initiative. The dual-use opportunities are tremendous. In a decade, airfreight may be transported coast-to-coast or transoceanic by remotely piloted aircrafts. Ground robots offer added safety for the warfighter in dismantling improved explosive devices and in gaining early intelligence before entering unknown territory. The increased distance between the warfighter and/or first responders is of tremendous value.

Further Information

www.robotics-vo.us

Contact: Henrik I. Christensen, PhD
KUKA Chair of Robotics
Georgia Institute of Technology
Atlanta, GA 30332-0180
Phone: +1 404-385-7480
Email: hic@cc.gatech.edu

Roadmap for Robotics in Manufacturing

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 healthcare, 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 proposal for a "National Network for Manufacturing Innovation" (NNMI).

"Effective use of robotics will increase U.S. jobs, improve the quality of these jobs, and enhance our global competitiveness."

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 keeps 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 would otherwise go offshore; 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 have already been recognized by companies such as Apple, Lenovo and Tesla in their setup of new factories in the U.S. Through utilization of robotics and automation, the expectation is that such in-shoring will continue to flourish.

This white paper summarizes 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.

1. Introduction

Image courtesy of Seegrid

This document summarizes the activities and results of a workshop on manufacturing and automation robotics that was carried out under the auspices of the Robotics-VO, an organization sponsored by the National Science Foundation. The workshop was one of five organized to update *A Roadmap for U.S. Robotics: From Internet to Robotics* [NRR 09]. The objective of the workshop was an update of the roadmap considering progress over the last 4-5 years. The research agenda proposed in this report will lead to a significant strengthening of 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.”

The Manufacturing and Automation Robotics Workshop was held on October 2, 2012 in Washington D.C. (http://robotics-vo.us/node/201). The goal was three-fold: First, to determine the strategic importance of robotics and automation technologies in manufacturing industries in the U.S. economy (Section 2); second, to determine applications where robotics and automation technologies could increase productivity (Section 3); and third, to determine research and development that needs to be done in order to make robotics and automation technologies cost effective in these applications (Section 4). To achieve this, whitepapers describing current uses and future needs of robotics in industry were solicited from professionals responsible for manufacturing in their companies. A number of academics were also invited to ensure broad coverage across research and practice.

2. Strategic Importance of Robotics in Manufacturing

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 14% of the U.S. GDP and about 11% of the total employment [WB-11,BEA-11]. Fully 70% of the net export of the U.S. is related to manufacturing [State09], so the sector represents an area of extreme importance to the general economic health of the country. Within manufacturing, robotics represents a $5B industry in the U.S. that is growing steadily at 8% per year. This core robotics industry is supported by the manufacturing industry, which provides the instrumentation, auxiliary automation equipment, and the systems integration adding up to a $20B 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 14% of the U.S. GDP, and 10% of U.S. employment [WB06]. 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 more than 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.

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 in the world has dropped significantly to only 28%. 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-08]. 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 nanoscale 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 [NAE07]. India and China are systematically luring back their scientists and engineers after they are trained in the U.S. According to [NAE07], 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 has been investing $600M into robotics and cognitive systems as part of the 7th Framework Programme. In the Horizon 2020 program, that investment will be completed 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. The non-defense U.S. federal investment 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. In addition, during 2011, Tesla Motors in California opened up a factory for the manufacturing of alternative fuel cars using heavy automation to enable all of the manufacturing to remain the United States.

2.2 Growth Areas

The Department of Commerce and the Council on Competitiveness [CoC08, CoC10, DoC04] 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**	**Growth**
Robotics–Manufacturing, Service, and Medical	20%	0-90%
IP Companies	25%	15-32%
Entertainment/Toys	6%	4-21%
Media/Games	14%	2-48%
Home Appliances	1%	-5-6%
Capital Equipment	9%	-2-18%
Automotive	2%	-11-31%
Logistics	22%	6-92%
Automation	6%	2-12%

Consolidated annual growth rates over a set of 280 U.S. companies for the period 2004-2011

Current growth areas for manufacturing include logistics, 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.

2.3 "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 U.S. companies and industries that currently make up a large percentage of the USGDP 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 robots as a threat. For example, two-thirds of the owners of iRobot's autonomous vacuum cleaner have named their Roomba, and one-third admit to taking their Roomba to visit friends.

2.4 A Vision for Manufacturing

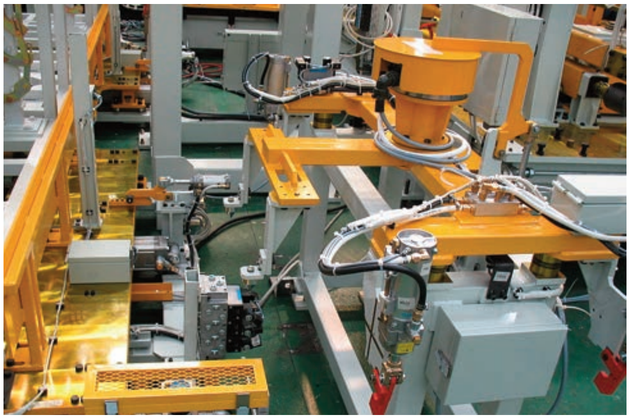

Image courtesy of Mixabest via Wikimedia Commons

U.S. manufacturing today is where database technology was in the early 1960s, a patchwork of ad hoc solutions that lacked the rigorous methodology that leads to scientific innovation. In 1970, Ted Codd, an IBM mathematician, invented relational algebra, an elegant mathematical database model that galvanized federally funded research and education leading to today's $14 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 ManufacTuring." The result, as with databases and computers, would be higher quality, more reliable products, reduced costs, and faster delivery [GK07, 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 others' strengths (e.g., humans 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 the 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 would support millions of workers in this vital sector of the U.S. economy. Such investments would benefit healthcare, agriculture, and transportation, and would 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.

3. Research Roadmap

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 and discussed in Section 4). Integration of all the parts of this roadmap into a cohesive program is essential.

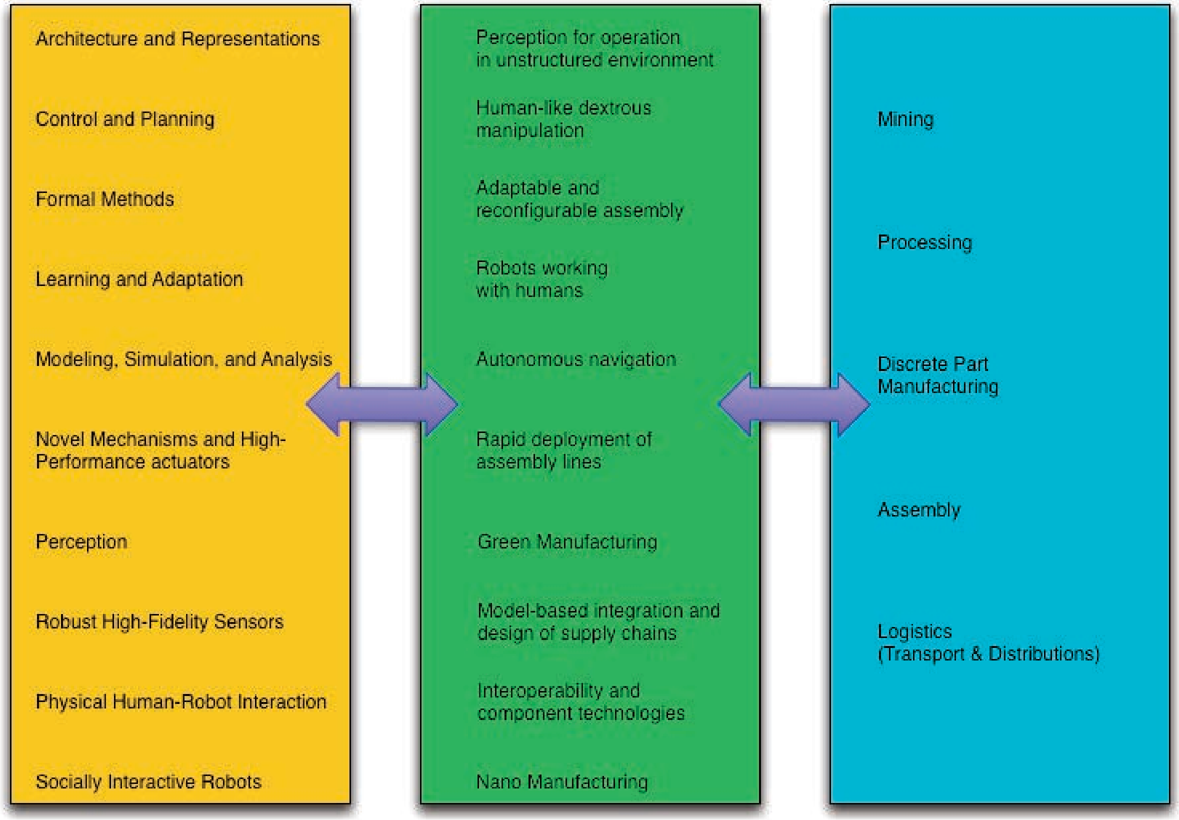

Figure 1: *The roadmap process: Research and development is needed in technology areas that arise from the critical capabilities required to impact manufacturing application domains.*

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 5-, 10-, and 15-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 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 costs $22k and can be programmed directly by demonstration with little or no training. 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 in a foreign country is found to suffer from serious quality control problems. The U.S.-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.

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. In Section 4, we describe some promising research directions that could enable us to meet these milestones.

3.3.1 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 workcells 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.

3.3.2 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 an 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 years:** 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 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).

3.3.3 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 the 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.

3.3.4 Humanlike 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:

Image courtesy of Seegrid

- **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 ten or more 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.

3.3.5 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 Manufac*Turing*." If achieved, as with databases and computers, this 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.

3.3.6 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.

3.3.7 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 necessitates 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 environments 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.

3.3.8 Intrinsically Safe Robots Working with Humans: The Democratization of Robots


Image courtesy of Rethink 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 that number 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.


Image courtesy of energy.gov via Flickr

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 workplace while maintaining consistent performance.

- **15 years:** Systems that can recognize, work with, and adapt to human or other robot behaviors in an unstructured environment (e.g. construction zones or newly configured manufacturing cells).

3.3.9 Education and Training

The U.S. can only take advantage of new research results and technology if there is a workforce well-trained in the basics of robotics and the relevant technologies. This workforce should have a wide range of skill and knowledge levels—from people trained at vocational schools and community colleges to operate high-tech manufacturing equipment, to BS- and MS-level developers trained to create robust, high-tech manufacturing equipment, to PhD-level basic researchers trained to develop and prove new theories, models, and algorithms for next-generation robots. To train the best workforce, the educational opportunities must be broadly available. The roadmap for the workforce is as follows:

- **5 years:** Each public secondary school in the U.S. has a robotics program available after school. The program includes various informational and competitive public events during each session, and participants receive recognition comparable to other popular extra-curricular activities.

- **10 years:** In addition to the 5-year goal, every 4-year college and university offers concentrations in robotics to augment many Bachelor's, Master's, and PhD degrees.

- **15 years:** The number of domestic graduate students at all levels with training in robotics is double what it was in 2008. Ten ABET-approved BS programs in Robotics and 10 PhD programs in Robotics are active.

Products designed without considering the manufacturing processes and their associated constraints are typically more expensive than strictly necessary to manufacture. Design for Manufacturing, and more specifically Design for Automation, is a critical product design competency necessary to ensure that the promise of automation is realized and to transition advanced manufacturing techniques into U.S. corporations.

- **5 years:** A continuing education and coaching program on how to design for manufacturing and automation.
- **10 years:** Design for Automation awareness incorporated into degree programs at every 4-year college and university offering engineering degrees. Design for Automation specialization preparation available at a set of research universities.
- **15 years:** The knowledge of design for automation is widespread across the workforce. Continuing education and coaching programs are available on various dimensions of automation requirements and degree requirements are in place to ensure the training of the next generation workforce.

4. Research and Development: Promising Directions

Achieving the **critical capabilities** described in Section 3 above and listed in the center column of Figure 1 requires basic research and development of the **technologies** listed in the left column of Figure 1. These technologies are briefly motivated and described below along with promising research directions. Note that each one supports more than one critical capability. For example, the "Perception" technology directly impacts "Operation in Unstructured Environments," "Intrinsically Safe Robots Working with Humans," "Autonomous Navigation," and "Humanlike Dexterous Manipulation."

4.1 Learning and Adaptation

One of the biggest barriers to the use of robots in factories is the high cost of engineering the workcells (i.e., the design, fabrication, and installation of jigs, fixtures, conveyors, and third-party sensors and software). These engineering costs are typically several times the cost of the primary robotic hardware. Robots must be able to perform their tasks in environments with greater uncertainty than current systems can tolerate. One possible way to achieve this is through learning by demonstration. In this case, a human performs the task several times without the engineered environment while the robot observes. The robot then learns to mimic the human by repeatedly performing the same task safely and comparing its actions and task results to the human's. Robots could also adapt by monitoring their actions, comparing them to nominal parameterized task representations, and adjusting the parameters to optimize their performance. It is also possible for robot systems to use "Iterative Learning" techniques to improve performance beyond that of human demonstrations in terms of speed and reliability.

4.2 Modeling, Analysis, Simulation, and Control

Modeling, analysis, simulation, and control are essential to understanding complex systems, such as manufacturing systems. Future manufacturing systems will require models of parts or subassemblies undergoing intermittent contact, flexible sheet-like materials, linkages with closed chains, systems with changing kinematic topologies, and relevant physics at the micro- and nano-scales. To leverage these to design improved manufacturing systems, models and the resulting simulation techniques need to be validated experimentally and combined with search and optimization techniques. With improved models and simulation techniques and with improved high-performance computing, we will have the ability to simulate all aspects of manufacturing systems from the extraction of raw materials, to the production of parts, to the assembly and testing.

4.3 Formal Methods

In some domains, mathematical models and the tools of logic have been used to guide specification, development, and verification of software and hardware systems. Because of the high cost of application, these *formal methods* have been used in significant manufacturing efforts primarily when system integrity is of the utmost importance, such as spacecraft and commercial aircraft. However, it is not only the cost that prevents formal methods from common use in the development of manufacturing (and many other engineered) systems. Lack of use is also related to the limitations of the framework for representing important manufacturing operations, such as the assembly of parts, which can be viewed as hybrid systems with disjunctive nonlinear inequality constraints of many continuous variables.

4.4 Control and Planning

Robots of the future will need more advanced control and planning algorithms capable of dealing with systems with greater uncertainty, wider tolerances, and larger numbers of degrees of freedom than current systems can handle. We will likely need robot arms on mobile bases whose end-effectors can be positioned accurately enough to perform fine manipulation tasks despite the base not being rigidly anchored to the floor. These robots might have a total of 12 degrees of freedom. At the other extreme are anthropomorphic humanoid robots that could have as many 60 degrees of freedom. Powerful new planning methods, possibly combining new techniques from mathematical topology and recent sampling-based planning methods may be able to effectively search the relevant high-dimensional spaces.

4.5 Perception

Future factory robots will need much improved perception systems in order to monitor the progress of their tasks and the tasks of those around them. Beyond task monitoring, the robots should be able to inspect subassemblies and product components in real time to avoid wasting time and money on products with out-of-spec parts. They should also be able to estimate the emotional and physical state of humans, since this information is needed to maintain maximal productivity. To do this, we need better tactile and force sensors and better methods of image understanding. Important challenges include non-invasive biometric sensors and useable models of human behavior and emotion.

The large cost of engineering of workcells derives mostly from the need to reduce uncertainty. To remove this cost, the robots must be capable of removing uncertainty through high-fidelity sensors or actions that reduce uncertainty. Sensors must be able to construct geometric and physical models of parts critical to an assembly task and to track the progress of the task. If a human is doing this task partly or wholly, then non-invasive biometric sensors must also determine the state of the human. Grasping actions and assembly strategies that previously depended on expensive tooling should be redesigned so that they take advantage of compliance to remove uncertainty.

4.6 Novel Mechanisms and High-performance Actuators

Improved mechanism and actuators will generally lead to robots with improved performance, so fundamental research is needed on these topics. Development in actuators has historically focused on mechanical performance metrics such as accuracy, repeatability, and resolution. However, as robotics is applied to applications in novel domains such the manipulation of parts on the nano- and micro-scales, materials-sensitive environments such as those surrounding MRI scanners, and environments shared with humans, the designs (including material choices) of actuators and mechanisms will have to be rethought. Further, the adoption of robotics as manufacturing partners emphasizes the need for fundamental research in safe actuation. New mechanisms for human augmentation include exoskeletons, smart prosthetics, and passive devices. These systems will require high strength-to-weight ratios, actuators with low emissions (including noise and electromagnetic), and natural interfaces between the human and the mechanisms.

4.7 Human-Robot Interaction

In the manufacturing environment, the primary importance attached to interaction between humans and robots is safety. Beyond that, the practical fact of the matter is that robots are used for the benefits they bring in terms of cost, efficiency, and productivity. If it is determined that collaborative activities between robots and humans are cost-effective and appreciably more productive than activities involving either group working alone, then these work strategies will be adopted and refined. Therefore, any activities involving human-robot interaction must provide favorable results toward that end.

Designing robot systems with the end-user in mind, as well as the product/task, will result in human-robot interaction that is not only safer but also more cost effective, efficient, and productive. Simple, clear interfaces and observable, transparent behaviors can make working with robots as intuitive as working with fellow human workers. Both humans and robots will need to provide indicators of intent (verbal and non-verbal) that are easily understood. When robots are collaborating with humans, they must be able to recognize human activities to maintain proper task synchrony. Similarly, humans must be able to read and recognize robot activities in order to interpret the robot’s understanding. (For example, a robot that has drilled a set of holes in the right place, but at the wrong depth, may be telling you that the task was not clearly specified. If only that message were conveyed in some form other than the row of faulty holes!) Finally, robots must be easy to train and easy to learn how to use them. Learning aids should be built into the robot system both for initial use, maintenance, learning, and error diagnostics/fault recovery.

These situations suggest the need for new sensing systems with higher bandwidths and resolutions than those available today; the use of sensing systems that capture biometric data of human workers that has previously been ignored in robot control; and the design of communication options between humans and robots that include natural language, gesture, vision, and haptics.

4.8 Architecture and Representations

New manufacturing robots must be intelligent enough to productively share space with humans and other robots and to learn how to improve their effectiveness with experience. To support such learning, robot operating systems, and the models and algorithms behind them, must be sufficiently expressive and properly structured. They will need ways to represent the various manipulation skills and relevant physical properties of the environment to incorporate their impact on task execution. There should be continuous low-level perception-action loops whose couplings are controlled by high-level reasoning. Robots will exploit flexible and rich skill representations in conjunction with observation of humans and other robots to learn new skills autonomously. Robots will need new methods of representing environmental uncertainties and monitoring tasks that facilitate error recovery and skill enhancement based on these errors.

4.9 Measurement Science

Research results have typically been difficult to transition from university laboratories to industrial implementations. In order to assure that the research spotlighted in this document achieves fruition in manufacturing shop floors, there needs to be a community-wide effort to develop the underlying measurement science to assess and promote progress and technology transfer. Measurement science is a term used to describe the infrastructural underpinnings essential for technology maturation. Elements under the broad measurement science umbrella include fundamental metrology, performance metrics, test methods, reference artifacts and data, reference architectures, and critical technical inputs to standards. Benchmarks and testbeds help researchers assess progress, replicate experiments, and validate new technologies. Ultimately, a well-founded characterization of a new technology's expected performance under realistic conditions reduces the risk of adoption by industry and stimulates progress.

4.10 "Cloud" Robotics and Automation for Manufacturing

Manufacturing systems require reliable sensing and interacting with complex, dynamic, high-dimensional environments and systems. In 2010, a new paradigm emerged, "Cloud Robotics," which shifts the demanding processing and data management to the Cloud. Summarized as: "No robot is an island" (Steve Cousins of Willow Garage), this idea is gaining attention at major companies such as Google and Cisco.

Drivers include: rapidly expanding and improving wireless networking, availability of vast and rapidly-expanding on-demand (elastic) computing clusters, huge data centers that can collect and accumulate shared datasets, "big data" techniques, "Internet of Things" where devices and objects have RFIDs or internal servers, crowdsourcing, open-source sharing of data and code, regular backups, software updates, and security patches.

The Google self-driving car indexes a vast library of maps and images; it exemplifies "Cloud Robotics;" which treats the Cloud as a vast and rapidly expanding resource for massively parallel computation and real-time sharing of vast data resources. A (non-robotics) example: Apple's Siri system: speech understanding using both local and remote processing and confidences are used to decide if remote analysis is required. Each instance and outcome is saved for global incremental learning.

Sharing data across manufacturing equipment operating in different facilities can lead to better diagnostics and adaptation in manufacturing settings leading to lower process variation and greater manufacturing efficiency. Robots performing manipulation tasks or just navigating in unstructured environments can learn from each other. ROS-like tools can lead to a paradigm where the role of expensive hardware is minimized, hardware costs are greatly reduced, and software and hardware architectures that are modular and extensible leading to more robust robots by facilitating rapid design-test-redesign iterations. Sharing data across manufacturing equipment operating in facilities across the country can lead to better diagnostics and adaptation in manufacturing settings leading to lower process variation and greater manufacturing efficiency. Robots performing manipulation tasks or just navigating in unstructured environments can learn from each other. ROS-like tools can lead to a paradigm where the role of expensive hardware is minimized, hardware costs are greatly reduced, and software and hardware architectures that are modular and extensible leading to more robust robots by facilitating rapid design-test-redesign iterations.


Image courtesy of Google via Flickr

Cloud Robotics has the potential to significantly improve performance in at least five ways:

1) Offering a global library of images, maps, and object data, often with geometry and mechanical properties.

2) Massively-parallel computation on demand for sample-based statistical modeling and motion planning.

3) Inter-robot sharing of outcomes, trajectories, and dynamic control policies.

4) Human sharing of "open-source" code, data, and designs for programming, experimentation, and hardware construction, in particular the rising popularity of the ROS system.

5) Detecting problems and requesting on-demand human diagnostics/guidance.

Challenges/Issues: Security, Privacy, Latency/QoS due to Intermittent/Bursty Communication.

5. References

[BEA11] Bureau of Economic Analysis, U.S. Department of Commerce, 2011.

[BEA07] Bureau of Economic Analysis, U.S. Department of Commerce Press Release, April 24, 2007. www.bea.gov/newsreleases/industry/gdpindustry/2007/gdpind06.htm.

[CoC08,CoC10] Council on Competitiveness, Competitiveness Agenda-New Challenges, New Answers, November 2008, 2010 (www.compete.org).

[Cloud12] Cloud Robotics summary with links to resources, recent papers, presentations, etc: http://goldberg.berkeley.edu/cloud-robotics/.

[DoC04] 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 O shore on Technology Development Paths in the Automotive and Optoelectronic Industries. Ph.D. Thesis. M.I.T. Cambridge, MA: 2006.

[GK07] Goldberg, K., Kumar, V, "Made in the USA" can be Revitalized, San Jose Mercury News: Op-Ed, 24, October 2007.

[NAE07] Rising Above The Gathering Storm: Energizing and Employing America for a Brighter Economic Future, National Academy of Engineering, 2007.

[WB06] Where is the Wealth of Nations? The International Bank for Reconstruction and Development, The World Bank, 2006.

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

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

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

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

6. Contributors

This report has its origins in presentations and discussions at a workshop on manufacturing and automation robotics that took place October 2, 2012 in Washington, D.C. The report is part of the Robotics-VO study on Robotics. The National Science Foundation (NSF) sponsors the Robotics-Virtual Organization (Robotics-VO). The present report has been authored by the workshop organizers and does not necessarily reflect the opinion of the Robotics-VO or NSF. The responsibility of the report lies entirely with the authors.

The workshop organizers were Henrik I. Christensen, Ken Goldberg, Vijay Kumar, and Elena Messina. The workshop had broad participation across academia and industry as shown in the list of participants below:

Jeff Baird
Adept Technology

Gary Bradski
Industrial Perception

Michael Branicky
Case Western Reserve University

Rodney Brooks
Rethink Robotics

Jenny Burke
The Boeing Company

Henrik I. Christensen
Georgia Institute of Technology

John Enright
Kiva Systems

Clay Flannigan
Southwest Research Institute (SWRI)

Philip Freeman
The Boeing Company

Thomas Fuhlbrigge
ABB U.S. Corporate Research

Joe Gemma
Staubli

Satyandra Gupta
University of Maryland

Ginger Hildebrand
Schlumberger

Vijay Kumar
University of Pennsylvania

Peter Luh
University of Connecticut

Matt Mason
Carnegie Mellon University

Elena Messina
National Institute of Standards and Technology

Erik Nieves
Motoman

Curtis Richardson
Spirit AeroSystems

Daniela Rus
Massachusetts Institute of Technology

Mark Spong
University of Texas–Dallas

Jason Tsai
FANUC Robotics America

Stuart Shepherd
KUKA Robotics

Roadmap for Healthcare and Medical Robotics

Motivation and Scope

Several major societal drivers for improved healthcare access, affordability, quality, and personalization can be addressed by robotics technology. Existing medical procedures can be improved and new ones developed, to be less invasive and produce fewer side effects, resulting in faster recovery times and improved worker productivity, substantially improving both risk-benefit and cost-benefit ratios. Medical robotics is already a major success in several areas of surgery, including prostate and cardiac surgery procedures. Robots are also being used for rehabilitation and in intelligent prostheses to help people recover lost function. Telemedicine and assistive robotics methods are addressing the delivery of healthcare in inaccessible locations, ranging from rural areas lacking specialist expertise to post-disaster and battlefield areas. Socially assistive robotics (discussed in Section 2.4) efforts are developing affordable in-clinic and in-home technologies for monitoring, coaching, and motivating both cognitive and physical exercises addressing the range of needs from prevention to rehabilitation to promoting reintegration in society. With the aging population a dominating demographic, robotics technologies are being developed toward promoting aging in place (i.e., at home), delaying the onset of dementia, and providing companionship to mitigate isolation and depression. Furthermore, robotics sensing and activity modeling methods have the potential to play key roles in improving early screening, continual assessment, and personalized, effective, affordable intervention and therapy. All of the above pursuits will have the effect of maintaining and improving productivity of the workforce and increasing its size, as well as enabling people with disabilities, whose numbers are on the rise, to return to the workforce. Today, the U.S. is the leader in robot-assisted surgery and socially assistive robotics for continued quality of life aimed at special-needs populations and the elderly. However, other countries are fast followers, having already recognized both the need and the promise of such technologies.

> ***"Robotics technologies are being developed toward promoting aging in place, delaying the onset of dementia, and providing companionship to mitigate isolation and depression."***

Participants

The workshop contributors consisted of experts in surgical robotics, prosthetics, implants, rehabilitation robotics, and socially assistive robotics, as well as representatives from industry. All participants contributed insights from their communities and areas of expertise; many common interests and challenges were identified, informing the roadmap revision effort.

Workshop Findings

The spectrum of robotic system niches in medicine and health 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). Technical challenges increase with the complexity of the environment, task, and user (dis)ability. The problem domains identified are those of greatest societal need and largest predicted impact: surgery and intervention; replacement of diminished/lost function; recovery and rehabilitation; behavioral therapy; personalized care for special needs populations; and wellness and health promotion. Those problem domains involved the following set of capabilities that present technological and research challenges: physical human-robot interaction and interfaces; social human-robot interaction and interfaces; robot-mediated health communication; automated understanding of human state and behavior during robot interaction; large-scale and long-term modeling of users' changing needs; quantitative diagnosis, assessment, and training; information map-guided interventions; high dexterity manipulation; sensor-based automated health data acquisition; and secure and safe robot behavior. In addition, key technology deployment issues were identified, including: reliable and continuous operation in human environments, privacy, security, interoperability, acceptability, and trust. The few funding opportunities for interdisciplinary integrative projects that bring together expertise in engineering, health (and business), and develop and evaluate complete systems in human subjects studies were identified as the cause for a lack of critical mass of new, tested, and deployed technological innovations, products, and businesses to create an industry.

1. Introduction

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, medical and health 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 room.

Robotics is already beginning to affect healthcare. Telerobotic systems such as the da Vinci Surgical System are being used to perform surgery, resulting in shorter recovery times and more reliable outcomes in some procedures. 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 MIT-Manus (commercially, InMotion) 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 have already proven more effective than conventional approaches, especially in assisting recovery after stroke, the leading cause of permanent disability in the U.S. 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.

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.

1.2 Societal Drivers

There are numerous societal drivers for improved healthcare that can be addressed by robotic technology. These drivers lie, broadly, in two categories: broadening access to healthcare and improving prevention and patient outcomes.

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, healthcare costs would be lowered due to improved quality, fewer complications, shorter hospital stays, and increased efficiency of treatment.

Image courtesy of Ekso Bionics

Economic and population factors must be considered. In the United States, over 15% of the population is uninsured [Census: Income, Poverty, and Health Insurance Coverage in the United States: 2007]; many others are under-insured. The situation prevents individuals from receiving needed healthcare, sometimes resulting in loss of function or even life, and also prevents patients from seeking preventative or early treatment, resulting is worsening of subsequent health problems. Access to healthcare 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 ageing 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 healthcare 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 healthcare. 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 healthcare. Demographic studies show that the U.S. population will undergo a period of significant population aging over the next several decades. Specifically, by 2030, the U.S. will experience an approximately 40% increase in the number of elderly, Japan will see a doubling in the number of people over the age of 65, and Europe will have a 50% increase in the number of elderly. 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 1,660,290 new cancer cases (excluding the most common forms of skin cancer) will be identified in the U.S. in 2013. Furthermore, the probability of developing invasive cancers increases significantly with age [ACS Cancer Facts and Figures 2013].

These trends are producing a growing need for personalized healthcare. 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 in order 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 U.S. who 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 healthcare.

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). 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 healthcare are unaffordable. Robotics technologies can improve safety and monitoring to avoid missing 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 reduction in available labor/workforce.

Finally, keeping technology leadership in the broad domain of healthcare is a key goal, given the size of the U.S. population and its age demographics.

2. Strategic Findings

2.1 Surgical and Interventional Robotics

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.
- 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 are now being widely used in the operating room or interventional suite. Surgical robots are beginning to realize their potential in terms of improved accuracy and visualization, as well as enabling of new procedures.

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 have 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 accuracy 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 the course of 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 the surgeon's field of view, and recognition of surgical motions and patient state to evaluate performance and predict health outcomes.

Image courtesy of Aethon

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 orthopaedic 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 needle steering 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.

2.2 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 loads 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.

2.3 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 healthcare 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.
- 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 of 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 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 in order to prevent infection and speed healing. Such systems can serve both as force multipliers in heath care delivery, providing more care to more patients, and also as a means of delivering personalized, customized care to all patients.

2.4 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, and TBI as discussed in Section 1.2), the available healthcare needed to provide supervision and coaching for such behavior therapy is already lacking and on a recognized steady decline.

Socially assistive robotics (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 the elderly, 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 tools for children, the elderly, stroke patients, and other special-needs populations requiring personalized care.

2.5 Personalized Care for Special Needs Populations

The growth of special needs populations, including those with physical, social, and/or cognitive disorders, which may be developmental, early onset, age-related, or may occur at any stage of life, presents an increasing need for personalized care. Some of the pervasive disabilities are congenital (from birth), such as cerebral palsy and autism spectrum disorder, while others may occur at any point during one's lifetime (traumatic brain injury, stroke), and still others occur later in life but persist longer with the extended lifespan (Parkinson's Disease, dementia, and Alzheimer's Disease). In all cases, these conditions are life-long, requiring long-term cognitive and/or physical assistance associated with significant resources and costs.

Physically and socially assistive robotic systems of the types described above have the power to directly impact the user's ability to gain, regain, and retain independence and be maximally integrated into society. The most major of those recognized today include mobility, facilitating independence, and aging in place.

Physical mobility aids, ranging from devices for the visually impaired to the physically disabled, and from high-end intelligent wheelchairs to simpler, self-stabilizing canes, expand accessibility to goods and services and decrease isolation and the likelihood of depression and the need for managed care. Robotics technologies promise mobility aids that can provide adjustable levels of autonomy for the user, so one can choose how much control to give up, a key issue for the disabled community. Intelligent wheelchairs, guide-canes, and interactive walkers are just a few illustrative areas being developed.

> ***"In addition to physical/mechanical aid, special needs populations stand to benefit significantly from advances in socially assistive robotics."***

With the fast-growing elderly population, the need for devices that enable individuals with physical limitations and disabilities to continue living independently in their own homes is soaring. This need is augmented by the needs of the smaller but also growing population of the physically disabled, including war veterans. Complex systems for facilitating independence, such as machines that aid in manipulation and/or mobility for the severely disabled, and those that aid complex tasks such as personal toiletry and getting in/out of bed, are still in the early stages of development but show promise of fast progress. At the same time, mobile robotics research is advancing the development of mobile manipulation platforms, toward machines capable of fetching and delivering household items, opening doors, and generally facilitating the user's ability to live independently in his/her own home. The delay (or elimination, if possible) of the need for moving an individual to a managed care facility significantly decreases the cost and burden on the individual, family, and healthcare providers. It also greatly diminishes the likelihood of isolation, depression, and shortened lifespan.

In addition to physical/mechanical aid, special needs populations stand to benefit significantly from advances in socially assistive robotics (discussed in the previous section), which provide personalized monitoring, companionship, and motivation for cognitive and physical exercises associated with life-long health promotion.

2.6 Wellness/Health Promotion

Improved prevention and patient outcomes are broad and fundamental goals of healthcare. Better, more effective, accessible, and personalized ways of encouraging people to eat right, exercise, remain socially active, and maintain mental health, would significantly decrease many urgent and chronic health issues.

In spite of its fundamental importance, health promotion receives less attention and significantly fewer resources than health intervention. Research funding is justifiably aimed at efforts to identify causes and seek cures for diseases and conditions, rather than on their prevention, with the exception of vaccine research in specific sub-areas (e.g., cancer, AIDS). However, prevention-oriented research and its outcomes have the potential to most significantly impact health trends and the associated major costs to society. Insurance companies are particularly motivated to promote prevention, and to invest in technologies that do so. While they are not positioned to support basic research, they are willing to support evaluation trials of new technologies oriented toward prevention and health promotion.

Robotics technologies are being developed to address wellness promotion. Many of the advances described above also have extensions and applications for wellness. Specifically, robotic systems that promote, personalize, and coach exercise, whether through social and/or physical interaction, have large potential application niches from youth to the elderly, and from able-bodied to disabled, and from amateurs to professional athletes. Wearable devices that monitor physiologic responses and interact with robotic and computer-based systems also have the potential to promote personalized wellness regimens and facilitate early detection and continuous assessment of disorders. In this context, robotics is providing enabling technologies that interoperate with existing systems (e.g., laptop and desktop computers, wearable devices, in-home sensors) in order to leverage advances across fields and produce a broad span of usable technologies toward improving quality of life.

3. Key Challenges and Capabilities

3.1 Motivating Exemplar Scenarios

3.1.1 Surgery and Intervention


Image courtesy of Intuitive Surgical

A pre-operative diagnostic test indicates that a patient may have cancer in an internal organ. That 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 small incision. An imaging and navigation system guides the surgeon through the surgery and provides the surgeon three-dimensional views of the anatomy, with cancerous tumors clearly highlighted. The system gives the surgeon the sense that he or 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—thus acting as a digital assistant. The end result: the cancerous tissue is removed with very little impact on surrounding healthy tissue and the patient recovers quickly enough to return to work within the week, with little pain and scarring, and the burden of cancer lifted from the patient's mind.

3.1.2 Replacement of Diminished/Lost Function

A young person loses an upper limb in an accident. A robotic prosthesis with a dexterous hand that replicates the functionality of the lost limb is custom made to fit the patient through medical imaging, rapid prototyping processes, and robotic assembly. The prosthesis is seamlessly controlled by the patient's thoughts, using a minimally or non-invasive brain-machine interface. The patient can control all the joints of his or her artificial hand, and receives multimodal sensory feedback (e.g., force, texture, temperature), allowing him to interact naturally with the environment. Of particular importance to the user are being aware of the limb's motion even in the dark, feeling the warmth of a loved one's hand, and being able to perform complex manipulation tasks like tying his or her shoes.

3.1.3 Recovery and Rehabilitation

A patient is still unable to perform the tasks of daily living years after a stroke, and begins robot-assisted therapy in the clinic. The robotic device applies precisely the necessary forces to help the patient make appropriate limb movements, sometimes resisting the patient's motion in order to help him learn to make corrective motions. Data are recorded throughout therapy, allowing both the therapist and the robot to recommend optimal strategies for therapy, constantly updated with the changing performance of the patient. This precise, targeted rehabilitation process brings the patient more steady, repeatable, and natural limb control. Simultaneously, neuroscientists and neurologists are provided with data to help them understand the mechanisms of the deficit. Outside of the clinic, a home robot nurse/coach continues to work with the patient to motivate continued exercises while projecting appropriate authority and competence but not impeding the user's autonomy and self-efficacy. This effectively shortens convalescence time and sees the user through recovery.

3.1.4 Behavioral Therapy

A robot works with a child with neurodevelopmental disorders (e.g., autism spectrum disorder and other disorders) to provide personalized training for communication and social integration in the home. The robot interacts with the child in a social way, promoting social behaviors, including turn-taking in play, joint attention, pointing, and social referencing. The robot provides appropriate levels of challenge and motivation to retain the child's engagement. It becomes a trusted peer as well as a social catalyst for play with other children, first in the home, and then in the school lunchroom, and eventually on the public playground. Throughout, the robot collects quantitative data on the child's behavior that can be analyzed both automatically and by healthcare providers for continuous assessment and delivery of personalized therapy/treatment/intervention.

3.1.5 Personalized Care for Special Needs Populations

Personalized robots are given to the elderly and physically and/or cognitively disabled users (e.g., due to Alzheimers/dementia, traumatic brain injury). They are capable of monitoring user activity (from task-specific to general daily life) and providing coaching, motivation and encouragement in order to minimize isolation and facilitate activity and (re)integration in society. Robots send wireless information to summon caretakers as needed, and can be used to continually assess and look for warning signs of disorders or worsening conditions (e.g., decreased sense of balance, lessened social interaction, diminished vocalizations, lack of physical activity, increased isolation from family/friends) that trigger the need for early intervention and change in treatment.

3.1.6 Wellness and Health Promotion

Affordable and accessible personalized systems that monitor, encourage, and motivate desirable health habits, including proper diet, exercises, health checkups, relaxation, active connection and social interaction with family and friends, caring for pets, and so on are purchased as easily and readily as current personal computers, and easily configured for the user and made interoperable with other computing and sensory resources of the user environment. For example, robot pets monitor the amount of physical activity of a overweight diabetic user to promote increased physical activity, and motivate required reporting of dietary practices and health checkups, sharing appropriate information updates with the family and the healthcare provider, as well as with the insurance company whose rates adjust favorably in response to adherence to a healthy and preventive lifestyle.

3.2 Capabilities Roadmap

To address the healthcare challenges noted in Sections 1 and 2 and to achieve the exciting scenarios described immediately above in Section 3.1, we have developed a list of major capabilities that a robotic system must have for ideal integration into medicine and healthcare. These capabilities, in turn, motivate research into the technologies described in Section 4.

3.2.1 Physical Human-Robot Interaction and Interfaces

Almost all branches of medicine involve physical interaction between the clinician and the patient, ranging from cancer-removal surgery to post-stroke physical therapy. Robots can improve the efficacy of such treatments in three main ways: by enhancing the physical interaction between clinician and patient as it transpires; by helping clinicians safely practice diagnostic and interventional skills; and by directly delivering care to the patient. All such applications need intuitive interfaces for physical interaction between humans and robots and require advances in the core robotics areas of sensing, perception, and action. A great variety of sensing and perception is required, including recording the motions and forces of the user to infer intent, creating biomechanical models of the human body, and estimating the forces between the robot and the user. The reciprocal nature of interaction means that the robot will also need to provide useful feedback to the human operator, whether that person is a clinician or a patient. In addition to haptics (force and tactile cues), systems that also involve vision, hearing, and other senses also must be studied.

Systems involving physical human-robot interaction are difficult to design, in part because humans are extremely unpredictable from the perspective of a robot. Unlike a passive, static environment, humans dynamically change their motion, strength, and immediate purpose. These changes can be as simple as physiologic movement (e.g., a patient breathing during surgery) or as complex as the motions of a surgeon suturing. During physical interaction with a robot, the human is an integral part of a closed-loop feedback system, simultaneously exchanging information and energy with the robotic system; thus, the human cannot simply be treated as an external system input. In addition, the loop is often closed with both human motion and visual feedback, each with its own errors and delays, which can potentially cause instabilities in the human-robot system. Given these problems, how do we guarantee safe, intuitive, and useful physical interaction between robots and humans? There are several parallel approaches to solving these problems: modeling human behavior and dynamics; sensing the human's physical behavior in a very large number of dimensions; and developing robot behaviors that will ensure appropriate interaction no matter what the human does. Great strides have been made in these areas over the last two decades, yet there are still no existing systems that provide the user with an ideal experience of physically interacting with a robot in any domain of medicine. Five-, ten-, and fifteen-year goals for this capability focus on increasing complexity and uncertainty of the task at hand.

- **In 5 years:** New devices and algorithms will enable more effective two-way exchange of information and energy between the human and the robot. In surgical robotics, systems will be able to provide the full suite of physical feedback to the surgeon as they control the robotic instruments. The interface will provide rich haptic feedback including forces as well as complementary information such as surface texture and environmental compliance of the remote patient's tissue, with similar information available during simulated training sessions. Robotic devices for rehabilitation will be able to output a wide range of impedances, from completely un-encumbering (zero mass/stiffness/friction) to very high impedance with the ability to entirely support the patient's weight. Orthotic and prosthetic devices will restore lost functionality, such that the human user is able to conduct basic daily tasks without assistance. Understanding of desired human motion based on external sensors and brain-machine interfaces is essential for prosthesis design, requiring an appropriate mapping between human thoughts and the actions of a robotic prosthetic limb.

- **In 10 years:** Human-robot interaction will be made intuitive and transparent, such that the human's intent is seamlessly embodied by the robotic system. Interfaces should be automatically customized to the specific user to maximize intuitiveness of the interface. Interfaces will estimate the user's intent, rather than simply executing the user's commands that may be subject to human imperfections. Surgical robots will enable outcomes better than what can be expected in open surgery by eliminating non-useful information from the interface and from the motion command. In rehabilitation applications, robots will interface with patients to provide assistance and/or resistance along appropriate degrees of freedom, and should provide backdrivable or compliant behavior that transmits appropriate physical feedback regarding patient behaviors. Patients will feel feedback that makes them aware of movement errors and spasticity, encourages smooth repetitive movements, and is engaging and challenging as needed. Orthotic and prosthetic devices will enable functionality that begins to match that of the original biological capabilities.

- **In 15 years:** Assistance from robotic systems will enable the human user to become better than human. By sensing a human's movement and inferring intent, robots will be able to provide context-appropriate forces to a human operator, such as a rehabilitation patient using a robot to regain limb function and strength after a stroke. The robot will limit applied force or motion to levels that are useful and intuitive for the user. It will provide virtual constraints and other physically meaningful elements to help execute tasks accurately, or to provide feedback helpful for training, motor learning, and musculoskeletal adaptation. Surgical teleoperators will enable a human or multiple humans to control non-anthropomorphic manipulation systems with velocity limits, degrees of freedom, and kinematics that deviate significantly from that of a human, with the limitations of the robotic system intuitively conveyed to the human user(s). Both surgical teleoperation systems and rehabilitation robots will provide training to the human user throughout the physical interaction such that the human learns how to be a better user of the robotic system while simultaneously learning how to rely less on the robotic system. Orthotic and prosthetic devices will enable functionality that surpasses the original biological capabilities.

3.2.2 Social Human-Robot Interaction and Interfaces

Robots that provide social and cognitive support are beginning to appear in therapy, health, and wellness applications. 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.

Research challenges in achieving socially assistive robots 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, addressing bystanders, and so on.

> ***"Socially assistive robots motivate their users to pursue healthy behaviors, engage them in a therapy program, and provide an easy to use natural interface."***

A research goal that is particularly pertinent to healthcare 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 response 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.

- **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 multimodal 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, as well as shape the interaction to match the role and need of individual users.

3.2.3 Robot-Mediated Health Communication

Remote robotic telepresence 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 the elderly 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.

Enabling cost-efficient and effective robot-mediated health communication requires that 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. Thus, any advances toward robots operating autonomously in human environments (e.g., navigation and dexterous manipulation) will also impact telepresence robots. Robotics challenges unique to teleoperation and remote presence are highlighted next.

Embodiment of the remote mentor or collaborator, in contrast to video conferencing via a laptop, enhances the remote clinician's effectiveness by allowing him or her to navigate the environment, access more information for building common ground, and even to provide physical assessment and assistance. An important open question is the effectiveness of mediated human-human interactions, comparing health communication done through robotic telepresence to video and traditional face-to-face visits.

Designing effective and intuitive interfaces for remote users to control telepresence robots is an open challenge. For instance, it is imperative to investigate who can operate a telepresence robot, when to allow access to it, and where and in what ways the operator can control it. An operator's role and relationship to the person being visited impacts what health information is communicated (e.g., general health discussions versus personalized medication information including dosage and schedule). Due to

HIPAA regulations and the private nature of health communications, additional measures will be necessary to protect the privacy of end-users and people in the same space as the robot. For example, the operator may see high-quality video and hear full-quality audio in the telepresence robot's immediate surroundings, but diminished quality beyond.

With technical advances in autonomous robots, telepresence robots could also 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 has to communicate to the end users whether it is being operated by a remote human or autonomously or by some combination of the two. The interface available to the user might vary across scenarios and across different levels of autonomy.

The following are a set of milestones in the research challenges of this area:

- **In 5 years:** We will have a better understanding of the healthcare benefits of robotic telepresence compared with other forms of physical presence, cyber-presence, and telepresence. Additionally, we will have advances in end-user privacy controls that will pave the way for field studies.

- **In 10 years:** Advances in manipulation and navigation capabilities move telepresence robotics into the realm of semi-autonomous control. These developments include advances in the design of intuitive novel control interfaces for remote users for effectively varying autonomy.

- **In 15 years:** We will have extended field evaluations of telepresence robots deployed in a variety of health scenarios. These robots will operate semi-autonomously with effective and intuitive control interfaces for remote users, afford socially appropriate physical and social interaction modalities, and offer advanced patient privacy controls.

3.2.4 Automated Understanding of Human State and Behavior During Robot Interaction

Medical and healthcare 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, as well as application of statistical methods for user modeling based on this multimodal data. Fundamental mechanistic models of how robot interaction affects state and behavior are still in their infancy; further development of these models will enable more effective design of the control algorithms for medical and healthcare 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 healthcare: 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.


Image courtesy of VGo via Flickr

Physiologic data sensors are typically wearable sensors and devices that provide real-time physiologic signals (e.g., heart rate, galvanic skin response, body temperature). 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 5 years:** Robots will be able to have the ability to capture human state and behavior (aided with wearable sensors) in controlled environments (e.g., physical therapy sessions, doctor's offices) with known structure and expected nature of interactions. Data from such sessions will begin to be used to develop models of the user that are useful for developing general schemes for optimizing robot interactions.

- **In 10 years:** Robots will be able to automatically classify human state and behavior from lightly instrumented users (lightweight sensors), in less structured settings (e.g., doctor's offices and homes with less-known structure), visualize those data for the user and the healthcare provider, and choose appropriate interactions for individual users based on the classification.

- **In 15 years:** Robots will be able to detect, classify, predict, and provide coaching for human activity within a known broad context (e.g., exercise, office work, dressing) with minimal use of obtrusive sensors. The robot will be able to provide intuitively visualized data for each user, based on the user's needs. Decisions for robot interactions based on the ongoing classification of state and behavior will use algorithms validated as effective in rigorous experimental studies.

3.2.5 Large-Scale and Long-Term Modeling of Users' Changing Needs

The need for robot systems to have user-specific models is especially important in healthcare domains. Each user has specific characteristics, needs, and preferences to which the robot must be attuned. Furthermore, these typically change over time as the person gets accustomed to the robot, and as his or her health state changes, over the short term (convalescence), medium term (rehabilitation), and lifelong (lifestyle changes, aging). To be effective, robotic systems interacting with human end-users must be able to learn user-specific models that differentiate one person from others and are adaptive over time.

To achieve this, the robot must take advantage of data integrated from a vast array of sources. Physical information about the user, both from wearable and remote sensors, such as cameras, provides data related to the user's current state. In addition, the Internet can provide user data, such as personal or medical history and past performance, as well as data about other, similar users. Challenges include integrating a spectrum of such multimodal information, in light of uncertainty and inconsistencies over time and among people. Machine learning, including robot learning, has been adopting increasingly principled statistical methods. However, much work remains in addressing complexities of uncertain real-world data (noisy, incomplete, and inconsistent), multimodal data about a user (ranging from signal-level information from tests, electrodes, and wearable devices, to symbolic information from charts, patient interviews, etc.), long-term data (over months/years of treatment), and data from multiple users (such as Internet-based compendia).


Image courtesy of Georgia Institute of Technology

User models should leverage both historical data from a single user (longitudinal) as well as data from multiple users (lateral). For instance, rather than basing an interaction on a fixed measure (such as what level of glucose is considered "normal"), a robot should base its decisions on models that are specifically related to a particular user—from past measurements and/or other users with similar characteristics. The models should be of sufficient resolution and accuracy to enable intelligent, rational decisions, whether by itself, caregivers, or the users themselves. In addition, it is critical that the models be created efficiently as new data arrives, to facilitate "just in time" user models.

The intersection of healthcare and social media creates a new avenue of data to enhance user models. The "Quantified Self" movement, self-tracking of health-related information (e.g., exercise routines, pedometer readings, diet plans), is newly popular, and a growing number of people are sharing this information via social media. Additionally, analyzing a user's interaction with social media provides an indication of personality and relationships. Tracking the changes in interactions over time may provide a unique way to understand mental changes not apparent from physiological sensors.

Taking these challenges into account, an adaptive, learning healthcare robot system would model the user's health state, compared to some baseline, in order to adjust its delivery of services accordingly. Such robots could generate quantitative metrics to show positive health outcomes based on health professional-prescribed convalescence/intervention/therapy/prevention methods.

- **In 5 years:** Robotic systems will use increasing amounts of multi-source real-world health data to generate user-specific models in spite of noisy, inconsistent data. User models will be of sufficient resolution to enable the system to adapt its interaction with the user to improve task performance within a particular context (e.g., specific exercise).

- **In 10 years:** Adaptive and learning systems will be extended to operate on multimodal, long-term data (months and more) and large-scale, Internet-based data of users with similar characteristics to generate user-specific models that highlight differences over time and between users. Robot systems will analyze social media to model user personality and relationships. Learned user models will support comprehensive interaction in extended contexts (e.g., daily activity).

- **In 15 years:** Adaptive and learning systems should actively collect data (from both sensors and the Internet) that are likely to be useful in differentiating between alternate user models. Sensor data and social media will be analyzed to model changes in the user's state of mind. Learned user models will support robot systems and healthcare providers in detecting and reacting to changes in the user's psychological and physical health in a timely manner. These comprehensive models of user health state will be used to continue to optimize human-machine interaction for improved health practices.

3.2.6 Quantitative Diagnosis, Assessment, and Training

Robots coupled to information systems can acquire data from patients in unprecedented ways. They can use sensors to record the physiologic status of the patient, engage the patient in physical interaction in order to acquire external measures of physical health such as strength, and interact with the patient in social ways to acquire behavioral data (e.g., eye gaze, gesture, joint attention) more objectively and repeatedly than a human observer could. In addition, the robot can be made aware of the history of the particular health condition and its treatment, and be informed by sensors of the interaction that occur between the physician or caregiver and the patient. The same quantitative diagnosis and assessment paradigm can provide information on the performance of a medical care provider, such as a surgeon or therapist, by tracking the provider's fatigue, stress, repeatability, accuracy, and procedural outcomes as new skills are acquired, or through routine assessment over a period of time.

Quantitative diagnosis and assessment requires sensing of the human user (patient or provider), the application of stimuli to gauge responses, and the intelligence to use the acquired data for diagnosis and assessment. When diagnosis or assessment is uncertain, the robot can be directed to acquire more appropriate data. The robot should be able to interact intelligently with the physician or caregiver to help them make a diagnosis or assessment with sophisticated domain knowledge, not necessarily to replace them. As robots facilitate aging in place (i.e., in the home), automated assessment becomes important as a means to alert a caregiver, who will not always be present, about potential health problems.

Many technological components related to diagnosis and assessment, such as micro-electromechanical lab-on-a-chip sensors for chemical analysis and "smart clothing" that records heart rate and other physiologic phenomena, borrow from ideas in robotics or have been used by robots in diagnosis and assessment. Others, such as using intelligent socially assistive robots to quantify behavioral data, are novel and present new ways of treating data that had, to date, been only qualitative. These new sensing paradigms, coupled with miniaturization of sensors and ubiquity of wireless communication, guarantee an ever-increasing volume of data to drive this quantitative paradigm.

An important challenge is to convert quantitative data into clinically-relevant performance metrics that demonstrate repeatability, accuracy, stability, and perhaps most importantly, validity in the clinical setting. The value of this quantitative approach becomes obvious when these performance metrics can be used to close the loop and drive the personalization of the training regimen of the healthcare provider, or treatment protocol for the patient.

Each of the stages of the diagnosis/assessment/training process needs to be improved individually, and the process then improved as an integrated whole. These stages include acquiring quantitative data by controlled application of stimuli, reducing the dimension and extracting critical details, and performing and potentially altering actions to achieve a better informed diagnosis/assessment and outcome. In some settings, this process is self-contained (i.e., administered within a controlled session) while in others, it may be an open-ended procedure (i.e., administered in a natural environment, such as the home).

To date, this cycle has been deployed in a qualitative sense, with the caregiver in the loop, determining the appropriate actions. The quantification and automation engendered in a robotic system facilitates closing of the loop, leading to significant improvements in terms of efficiency, reliability, and repeatability in both the delivery of patient care and the acquisition of procedural skills. Achieving this sophisticated process requires reaching several major milestones.

- **In 5 years:** A robot will be able to collect relevant behavioral and biophysical data reliably in nonlaboratory settings. Off-line analysis of these data (such as physiological state, movement and sensorimotor capability, eye gaze direction, social gestures) will lead to the computation of clinically relevant performance metrics. Validation of these metrics with significant subject pools ensures accurate diagnosis and assessment using the quantitative paradigm. Optimal ways of relaying the information to the robot system, the patient, and caregiver will be developed. Integration of multimodal physiological sensing and visualization of data are essential.

- **In 10 years:** We will be able to access biophysical signals using external hardware instrumentation and have direct analysis of both biophysical and movement behaviors to provide detailed diagnosis and/or assessment. Robotic devices are used to adaptively alter the stimulation to excite the appropriate behaviors and extract appropriate data, from the motor to the social. Algorithms for automatically extracting salient behaviors from multimodal data enable for data segmentation and analysis, for aiding quantitative diagnosis. These quantitative metrics inform the customization of the training or therapeutic protocol.

- **In 15 years:** We can accomplish connecting and easily accessing biophysical signals with wearable or implantable devices in real time. This is linked to integrated unencumbered multimodal sensing and intuitive data visualization environment for the user and caregiver. Real-time algorithms enable not only off-line but also online quantitative analysis of such data to inform in situ diagnosis as well as long-term patient tracking or skill acquisition. Systems are developed for in-home use and detection of early symptoms of pervasive disorders, such as autism spectrum disorder, from behavioral data. Similarly, the progression of degenerative motor disorders, such as Parkinson's or muscular dystrophy, can be monitored. Finally, by closing the loop, adaptive training algorithms based on the quantitative assessments will enable personalized protocols for procedural training for surgeons, or individualized rehabilitation regimens for cognitive and sensorimotor impairments.

3.2.7 Context-Appropriate Guidance and Variable Autonomy

Robots can provide context-appropriate guidance to human patients and caregivers, combining the strengths of the robot (accuracy, dexterity at small scales, and advanced sensory capabilities) with the strengths of the human (domain knowledge, advanced decision-making, and unexpected problem-solving). This shared-control concept is also known as a human-machine collaborative system, in which the operator works “in-the-loop” with the robot during task execution. As described earlier, humans (both patients and caregivers) represent uncertain elements in a control system. Thus, for a robot to provide appropriate assistance, it is essential that a robot understand the context of the task and the human behavior for a range of healthcare and medical applications.

In prosthesis control, low-level robotic controllers are needed to automatically adjust the coordinated behavior of the artificial limb to support the high-level activities desired by the user. The user must designate subsets of these coordinated behaviors with which to achieve a desired task. This model of collaboration establishes a command hierarchy in which the human selects the high-level action plan and the robot carries out those instructions. The resulting behavior should be so intuitive that the human operator does not even notice that some autonomy is taking place (or learns to rely and depend on it).

In surgical systems, human-robot collaboration joins the expertise of a surgeon with the super-human precision and accuracy of a robot to provide greater patient safety. Surgical systems may subscribe to a similar command hierarchy with the human surgeon commanding the high-level surgical plan and the robot performing precise micro-movements. In this case, additional autonomy can be delegated to the robot to maintain rigid constraint boundaries on movement to protect vital anatomical structures during surgery.

> ***“In surgical systems, human-robot collaboration joins the expertise of a surgeon with the super-human precision and accuracy of a robot.”***

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.

In assisted living, human-robot collaboration can address multiple tasks by allowing for shifts in the robot autonomy as circumstances require. The robot might influence behavior and provide support through social mechanisms—offering suggestions on activities, assisting in communication tasks, and providing cognitive support like reminders about medication. Under more critical circumstances (e.g., falls or extended inactivity), the robot might need to shift to a completely autonomous and proactive role in order to notify emergency responders. This variable autonomy within a single integrated system can provide more robust and flexible behavior while supporting desired modes of interaction with the target user.

- **In 5 years:** We will have integrated human-robot systems that rely on a fixed command hierarchy and roles, that operate in well-controlled environments or with limited interfaces, that rely on structured aspects of the environment, and that provide measurable gains in user performance.

- **In 10 years:** We will have the ability to shift roles based on user delegation or task context. The robots will model the complementary role of the human user in order to amplify human resources by providing customized support across greater time scales.

- **In 15 years:** We will adapt to context-driven variable autonomy under unstructured environmental settings such as the home. We will also assemble relevant historical data and consultations with expert caregivers to support a wide range of activities of daily living.

3.2.8 Information Map-Guided Interventions

Minimally invasive approaches to surgery reduce situational awareness because the relevant anatomy and tools cannot be directly seen and felt by the clinician. As minimally invasive surgery moves from straight rigid instruments to snake-like devices, this problem is exacerbated. Existing techniques to address this problem focus, for example, on computer-based fusing of pre-operative images with intra-operative images or rely on the clinician mentally fusing multiple sensor displays with memory-based models.

These approaches fall short since they fail to incorporate all available information and do not capitalize on a key capability of robotics to integrate sensor data, models, and other types of information into the planning and execution of motions. In medicine, a rich set of information sources is available. These include imaging technologies such as ultrasound, magnetic resonance imaging (MRI), spectroscopy, and optical coherence tomography (OCT). These modalities provide 3D geometric models of the anatomy enabling precise localization of diseased tissue as well as sensitive tissue to be avoided (e.g., nerves and blood vessels). Imaging techniques such as elastography also provide the capability to assign mechanical properties, e.g. (stiffness) to the 3D geometric model enabling surgical planning through the prediction of tool-tissue interaction forces. Additional sources of information include both instrument-mounted sensors (e.g., pressure, force) and patient-mounted sensors (e.g., heart rate, blood pressure, oxygen saturation). Furthermore, national databases can provide information on patients who underwent similar procedures including anatomy, surgical description as well as outcome. Finally, the available information also can include a patient-specific surgical plan updated in real time.

It is only through the creation and exploitation of an Information Map incorporating all of this data that the full potential of robotic minimally invasive surgery will be achieved. Such a map will provide the individual clinician with more detailed and up-to-date information than was ever available to their most experienced colleagues. Furthermore, the Information-Map-driven robot will endow the individual clinician with a skill level higher than that of the best experts in their field. Substantial technical challenges must be overcome to reach these goals. These include (1) creating techniques for building and updating the Information Map as a procedure progresses, (2) developing methods for immersing the clinician inside the information map so as to maximize his or her situational awareness resulting in effortless and precise control of robot with respect to the anatomy, (3) developing robots that are compatible with imaging and sensing tools (e.g., MRI and ultrasound), and (4) deriving robot control algorithms that maximize robot situational awareness with respect to the patient and with respect to surgeon intent. The latter will enable the robot to adapt clinician-prescribed motions toward desired surgical targets and away from delicate structures while also facilitating autonomous execution of surgical subtasks. Milestones toward these goals are given below.

> ***"Minimally invasive approaches to surgery reduce situational awareness because the relevant anatomy and tools cannot be directly seen and felt by the clinician."***

- **In 5 years:** Standardized techniques will be available for designing robots compatible with both MRI and ultrasound. This will involve actuator and mechanism design utilizing materials that enable robot motion during imaging, minimize artifacts that obscure both the robot and surrounding tissue, and satisfy physical interference constraints (e.g., fitting the robot and the patient inside the MRI bore). Also, robot-assisted techniques will enable automatic transformation of multimodal data to patient-specific physical models appropriate for guiding snake-like robots to reach desired targets while avoiding delicate structures.

- **In 10 years:** Algorithms will construct and update Information Maps using real-time sensor data as well as surgical databases. Surgeon interfaces connecting the clinician to both the Information Map and the robot will provide effortless control for tasks as diverse as brain tumor removal or the navigation of miniature swimming robots in the vasculature.

- **In 15 years:** Robot control algorithms will detect a surgeon's intent and, utilizing real-time information-map-based surgical planning, filter commanded robot motions so as to enhance both procedural outcome and safety. These algorithms will also enable semi-automated and automated surgical assistants that use fully real-time image-to-model generation (including geometry, mechanics, and physiological state).

3.2.9 High-Dexterity Manipulation

High-dexterity manipulation offers benefits both for minimally invasive surgical applications and for applications in which assistive robots physically interact with people. The potential for minimally invasive surgery, when compared to open surgery, is profound: reduced pain, quicker recovery, and reduced cost are all benefits that both the patient and society will gain. Unfortunately, minimally invasive surgery has not met this potential because of the limitations of existing instruments. Many are rigid and straight, limiting access to locations within line of sight of the incision. Others are flexible and buckle easily, limiting the locations that can be reached as well as the forces that can be applied. Many attempt to mimic open surgical technique but lack the requisite dexterity. In contrast, robotic technology can enable snake-like steerable instruments that possess controllable flexibility and thus allow a surgeon to reach deep into the anatomy through a small incision.

Access is not enough—minimally invasive instruments must also provide the dexterity necessary to perform such surgical tasks as retraction, dissection, tissue removal and tissue approximation. While dexterity enhancement through robotics has already been proven for laparoscopic-style surgery, there are substantial challenges to providing comparable dexterity at the tip of steerable snake-like robots and to inventing surgical techniques and tools optimized for these delivery platforms. These challenges span mechanism design, manufacturing, actuation, sensing, controls, and motion planning. Creating highly dexterous miniature devices is especially challenging, but is necessary to produce robots that can enter the body through tiny incisions, move with minimal damage through body lumens and along tissue planes, and produce the dexterous motions and forces needed for delicate surgical repairs. Taken to the limit, miniature tetherless robots will be able to either pass through or reside in the body to be activated as needed for modulating physiological functions or performing repairs. Existing capsule endoscopes and sensing smart pills are simple diagnostic examples of this approach.

For overcoming physical and cognitive limitations, assistive robotic technology has the potential to provide highly effective, individualized care to people with diverse needs. For example, a wearable robot could progressively rehabilitate the hand of a stroke patient by providing assistance as needed. Similarly, a mobile robot that manipulates everyday objects in a patient's home could assist with daily living tasks, or enable a doctor to make a house call via the robot. To achieve full potential, assistive robots will need to be capable of operating in human-centric environments. Furthermore, they will need to dexterously and safely manipulate objects in proximity with humans as well as to manipulate the human body itself. These capabilities entail great challenges, since, ultimately, caregiving robots would benefit from human-level dexterity or better. Progress will require innovations across robotics, from mechanism design to artificial intelligence. For example, as with humans, tactile sensing may be critical to robot dexterity, enabling robots to sense the world and carefully interact with it.

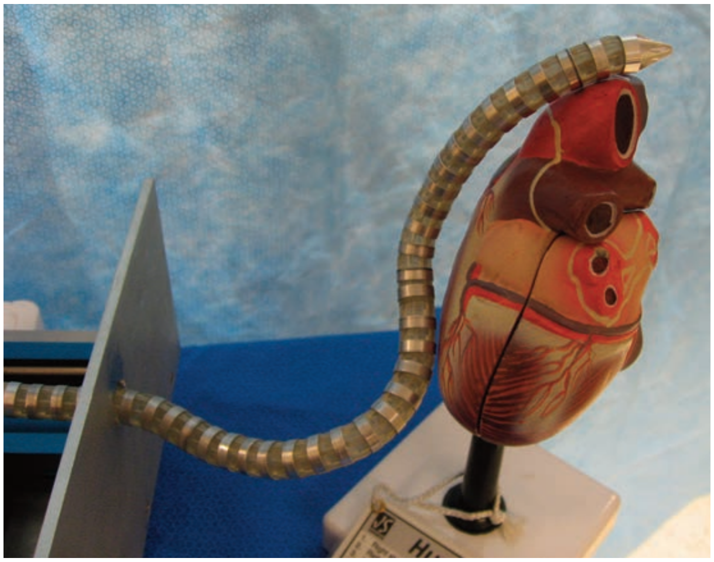
Image courtesy of Carnegie Mellon University

A natural set of milestones for dexterous manipulation is to consider capabilities linked to specific applications. Milestones for both surgical manipulators and assistive robots are given below.

- **In 5 years:** Snake-like robotic instruments will enable surgeons to perform simple natural orifice translumenal endoscopic surgical procedures in the abdomen via the stomach. Robot assistants will aid healthcare workers in safely moving patients in and out of hospital beds.
- **In 10 years:** Snake-like surgical robots will be capable of high-dexterity surgical tasks throughout the body and should also be miniaturized to enable precise microsurgical repairs. Tetherless centimeter-scale robots will be introduced that can perform interventional tasks inside the body such as removing polyps or modulating blood flow. Assistive robots will interact with impaired individuals to perform self-care tasks, such as grooming, hygiene, and dressing.
- **In 15 years:** Groups of tetherless millimeter- and micron-scale robots will be able to both swim through bodily fluids and bore through tissue to perform highly localized therapies. Assistive robots will autonomously perform general care-related tasks in human-centric environments with only high-level supervision.

3.2.10 Sensor-Based Automated Health Data Acquisition

We are approaching an age of nearly pervasive perception. Cameras are cheap, and getting cheaper, and image analysis algorithms are getting better. The networking infrastructure continues to improve. For whatever reason (home security, petcams, etc.), it is likely that significant parts of our lives will be observed by the resulting sensor network. The network itself continues to expand with effective and common sensors embedded in our cell phones (accelerometers, cameras, and GPS). Add to this the rapid growth in more conventional medical imaging, and the possibility of other biosensors, such as wearable monitors or ingested cameras and instrumented toilets, and it becomes technically feasible for each of us to have a detailed record covering nutrition, behavior, and physiology.

Similar trends are true in more acute care settings. Surgical robots provide the capability of creating a complete record of a surgical intervention. Digital endoscopy can be recorded for both robotic and non-robotic surgery. Most sensors used in hospital settings are now digital, and can be recorded for later analysis. The same is true of devices for rehabilitation or prosthetics. These data sources will provide baseline data for large populations, and thus will enable comparative studies to evaluate the impact of new robots for surgery, new haptic feedback for rehabilitation, or new neural interfaces for prosthetics. In short, these sources of data will enable quantitative and objective evaluation of research in medical robotics.

The key challenge is to create a common infrastructure to collect, organize, transmit, and store diverse sources of data. For example, a consolidated data record may store anonymized image information from a patient, key physiological readings, and data collected during a robotic surgery. To fully exploit these data sources, ontologies for data fields must be established, methods of anonymization, federation, and data sharing must be established, and open data sets related to common research problems must be established.

By aggregating over the entire population, we will have a database vastly more detailed and broader in scope than anything we have seen in the past. Such a database enables a new level of medical research based entirely on historical data. At present, medical studies are targeted to address specific issues or hypotheses, and the cost of these studies restricts their scope and duration. There are also some types of data, such as behavior patterns in one's normal life, which are very difficult to obtain at present. A large-scale database enables more open-ended research, identifying patterns or correlations that may never have been suspected. It also brings a new level of personalized healthcare, providing speedier and more accurate diagnoses, as well as a source of advice on lifestyle choices and their likely consequences.

- **In 5 years:** We will establish common and open infrastructure for data collection, building on recently developed models such as ROSbag, but extended to a broader class of medical robots and devices.

- **In 10 years:** We will create sharable data sets for key research areas, including robotic surgery, prosthetics, rehabilitation, and in-home living.

- **In 15 years:** We will create cloud-based analysis frameworks, with baseline performance of existing algorithms, to enable rapid design, development, and evaluation cycles for medical robotics research.

3.2.11 Secure and Safe Robot Behavior

The challenge of guaranteeing safe robot behavior is as old as the field of robotics itself. However, safety takes on a new dimension when direct physical interactions with human users, often vulnerable ones, constitute the core of the robot's purpose. In social human-robot settings, recognizing and providing appropriate response to human behavior (e.g., knowing the difference between inadvertent human behavior and specific intent) represents a new technical challenge. At another extreme, a surgical robot may manipulate a razor-sharp scalpel inside the abdomen or brain with obvious negative consequences of software or hardware failure.

One area of research is to investigate basic mechanisms for safe manipulation in the vicinity of sensitive or easily injured tissue. Smart "virtual fixtures" can provide a force field to protect tissues designated as particularly sensitive by the physician, and later automatically identified. Smart tissue handling will give the robot capability to "gently but firmly" retract internal organs with an understanding of the delicate tissues of an elderly patient and the limits to which they can be safely stretched. To make these approaches a practical reality, research is needed to identify quantitative models of safe tissue interactions that will drive intrinsically safe medical human-robot interaction.


Image courtesy of Georgia Institute of Technology

System security is another area where research is needed. As medical robots become networked and teleoperated, the consequences of a hacker blocking or taking over the communication link between a practitioner and the robot are extremely adverse. As a result, the communication link between the two must be highly secure. Additional issues in the medical robotics space include critical latency requirements, secure video transmission, and bi-directional authentication. Furthermore, having a mobile and sensorized robot in the home can provide a valuable resource if a person living alone is unresponsive and injured, but rules for granting access to outside parties will be difficult to define and to explain. It is also useful and increasingly feasible for robotic and prosthetic aids to record a history of their usage for diagnostic purposes, but this may disclose intimate details about their owner's health and behavior whose privacy and legal status may not be protected by physician-patient privilege.

Future robots must not only react to, but also anticipate, dangerous situations. In the home environment, the robot must be able to anticipate dangerous behavior or conditions (i.e., create virtual constraints) and respond safely to any urgent condition. When human-robot contact is involved, research is focusing on inherently safe mechanisms at the mechanical and hardware level to facilitate safety well before the software level. Surgical instruments must retain the capability for intentional therapeutic damage to tissue while still guaranteeing a high degree of safe handling of tissues that need to be preserved.

Safety of behavior has implications beyond the variables of physical interaction. While socially assistive robotics do not typically involve any physical contact between the robot and the user, the interaction may result in unwanted emotions such as strong attachment or aversion. While no such responses have yet been observed, the possibility of emotional stress or injury must be taken into account in the context of safe system design.

- **In 5 years:** We will exhibit cost-effective, inherently safe actuation, and light-weight/high-strength robot bodies in surgical and socially assistive robotics for in-clinic and in-home testing for specific tasks. We will fully characterize and theoretically counter security vulnerabilities present in remotely controlled surgical devices.

- **In 10 years:** We will create affordable and safe standardized translational research platforms (both hardware and software) for safe in-clinic and in-home robot evaluation with heterogeneous users (healthcare providers, family, patient). We will collect longitudinal data on safety and usability. We will test secure communication links suitable for secure telemedical interventions.

- **In 15 years:** We will achieve safe deployment of robot systems in unstructured environments (e.g., homes, outdoor settings) involving human-machine interaction in real-time with unknown users, with minimal training and using intuitive interfaces. We will deploy proven safe surgical robots with partial autonomy capabilities which achieve greater precision than human surgeons. We will develop secure telemedical interventions mediated by the open Internet.

3.3 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 healthcare 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 medical and health robotics 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. The lack of a supporting industry makes progress in medical and health robotics slow.

To create a health robotics industry, resources must first 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.

4. Basic Research/Technologies

Achieving the application-oriented capabilities described above will require significant progression of basic robotics research and the resulting technologies. This section describes the basic robotics research necessary to advance medical and health robotics.

4.1 Architecture and Representations

Robot control architectures encapsulate organizational principles for proper design of programs that control robot systems. One of the most complex fundamental problems that architectures address is the integration of low-level continuous perception-action loops with high-level symbolic reasoning through the use of appropriate data representations. The development of robot control architectures has reached a new level of complexity with medical and health robotics systems, because such systems must interact, in real time, with complex real-world environments, ranging from human tissue to human social interactions. Such systems and interactions feature multimodal sensing, various types of embodied interactions, and challenges for data representation and manipulation on a time-scale necessary for timely response. To address these challenges, architectures must be developed to facilitate principled programming for agile, adaptive systems for uncertain environments involving direct physical and/or non-physical interactions with one or multiple human users. For human-robot interaction, architectures must also account for modeling cognitive systems, skill and environment representations, reasoning about uncertainty, hierarchical and life-long skill learning and user modeling, real-time social interaction (including speech/language and physical activity interaction), and failure recovery, among others.

4.2 Formal Methods

Formal methods are mathematical approaches for the specification, development, and verification of systems. In medical and health robotics, they enable numerous core capabilities. One set of areas is robust modeling, analysis, and simulation tools for multi-scale systems. Formal methods allow optimal system integration, so that we can design systems based on robotic technologies whose components work with each other in a completely predictable fashion. For medical robots that interact directly with human caregivers and patients, controller designs, planners, operating software, and hardware should be verified and validated as safe when using formal methods. At this time, most work in formal methods does not incorporate uncertainty to the extent that is needed for medical and healthcare robotics. A related goal is the use of formal methods in the design and modeling the behavior of systems that work with humans, including formal modeling of human behavior and human-robot interaction.

4.3 Control and Planning

Control, defined here as the computation of low-level robot commands (such as how much torque a motor should apply), is an essential component of all physical robots. In medical robotics, a particularly important aspect of control is contact/force control. In this form of control, we usually want a robot to maintain contact with the environment with a given force (e.g., applying force to a patient in a rehabilitation scenario, contacting soft tissue during palpation, grasping an object with a prosthetic limb). Maintaining stable, safe contact is challenging because of time delays and imperfect dynamic models (especially models of friction). All of these problems need to be addressed through improvements in robot design, modeling, and control, all in parallel. Thus, developments in force/contact control are essential to the advancement of robots in contact with uncertain environments.

For any robot to function autonomously or semi-autonomously, it must use a plan to decide a course of action. Examples of plans in medical and health robotics include a plan for how to help a patient out of bed, how to reach a tumor in an organ, and how to motivate a patient to exercise. In medical and health robotics, plans must be adaptable to human inputs (e.g., that of a surgeon, caregiver, or patient) and uncertain environments (e.g., soft tissue, a living environment, or a patient being rehabilitated). While planning has been an extremely successful component of robotics research, much existing work relies on detailed knowledge of the environment and is designed for completely autonomous systems. Planning considerations for medical and health robotics require new approaches for operation in uncertain environments and with human input.

4.4 Perception

> ***"Planning considerations for medical and health robotics require new approaches for operation in uncertain environments and with human input."***

Robot perception, which uses sensor data and models to develop an understanding of a task, environment, or user, is a crucial component of all medical robots. In image-guided surgery, image data must be analyzed and transformed into useful information about particular features, such as organs, obstacles (e.g., the pelvic bone in urologic surgery), and target areas (e.g., a tumor embedded in the liver). Such perception often requires not only sensor data, but also information from an "atlas," which records features identified in many similar patients, so as to guide the process of recognizing important features in a particular patient. The output of the perception system can be used to develop a surgical plan, create a simulation, and provide real-time feedback to a human operator. Another form of perception relevant to healthcare is interpreting tactile, force and contact sensor data in order to build models of humans, robots, and environments, and the interaction between them. For example, if a prosthetic hand is holding a cup using a low-level control system (to lessen the human attention required), it essential to process data that allows the hand to determine whether the cup is being crushed or slipping out of the grasp, and how much liquid it contains.

A related issue is that robotic systems for healthcare must also understand some aspects of how human perception functions. For example, in image-guided surgery, information should be presented to the human operator in a manner that is intuitive, has appropriate level of detail and resolution, and does not distract from the task at hand. Another example is for applications in brain-controlled prostheses and some forms of robot-assisted physical rehabilitation. For such systems, understanding how humans will interpret feedback from the robot is key to the selection of sensors and the way their data are presented. Finally, in socially assistive robotics, the robot needs to understand enough about human perception in order to present information in socially appropriate, intuitive, and understandable ways. All these contexts require better models of human perception and will allow the interaction between humans and robots to be optimized.

Finally, a key challenge for systems that interact with a user is real-time perception and understanding of the user's activity in order to enable effective human-machine interaction. Natural, unconstrained human behavior is complex, notoriously unpredictable, and fraught with uncertainty. The development of wearable sensors and predictive models is necessary for facilitating solutions to human behavior perception and understanding, as discussed in Section 4.9, to follow.

4.5 Robust, High-Fidelity Sensors


Image courtesy of Ekso Bionics

Several types of sensing are especially important for medical robotics: biocompatible/implantable sensors, force/tactile sensing, and sensors that allow tracking and navigation. These sensors, along with perception algorithms, are often necessary to give state of a caregiver/physician, the patient, and (in some cases) the environment.

Biocompatible/implantable sensors would be a great catalyst to major advancements in this field. The close physical interaction between robots and patients requires systems that will not harm biological tissues or cease to function when in contact with them. In surgery, mechanisms must be designed that will not unintentionally damage tissues, and sensors need to be able to function appropriately in an environment with wetness, debris, and variable temperature. For prosthetics, sensors and probes must access muscles, neurons, and brain tissue and maintain functionality over long periods without performance degradation. These sensors and devices must be designed with medical and health robotics applications in mind, in order to define performance requirements.

When robots work in unstructured environments, especially around and in contact with humans, using the sense of touch is crucial to accurate, efficient, and safe operations. Tactile, force, and contact data is required for informed manipulation of soft materials, from human organs to blankets and other objects in the household. It is particularly challenging to acquire and interpret spatially distributed touch information, due to the large area and high resolution required of the sensors. Current sensors are limited in robustness, resolution, deformability, and size.

Improved tracking and navigation systems (particularly through non-contact means, such as imaging and wireless tracking/magnetic devices) are also important to enhance the performance of medical robots. Since many medical robots will need to rely on pre- and intra-operative data that is distributed and potentially changing with time, these sensing modalities can further support robotics as they become more accurate and precise, more accessible, and, in some cases (like ionizing radiation in fluoroscopy), be less harmful to patients. The use of biomarkers and contrast agents that work in conjunction with medical imaging are also important components of a medical robotics toolkit. Finally, imaging and tracking modalities must be compatible with actuation and mechanism technologies.

4.6 Novel Mechanisms and High-Performance Actuators

For systems ranging from ultra-minimally invasive surgery robots to human-size prosthetic fingers, medical robots need very small actuators and mechanisms with high power-to-weight ratio. These designs will allow us to build robots that are smaller, use less power, and are less costly. This enables greater effectiveness, as well as dissemination to populations in need. We will highlight below two examples of how advances in mechanisms and actuators could improve medicine.

In surgery, novel mechanisms are needed to allow dexterity of very small, inexpensive robots that can be mechanically controlled outside the body. Since many mechanisms are difficult to sterilize, surgery would benefit from disposable devices constructed from inexpensive materials and made using efficient assembly methods. As mentioned earlier, the capability of image-guided surgery relies (for some imaging methods) on specially designed, compatible robots that eliminate electric and magnetic components. This places particular constraints on actuators, which are electromechanical in most existing robots.

Advanced prostheses also motivate significant improvements in mechanisms and actuators. The design of robot hands with the dexterity of human hands, and arms and legs with the strength of human arms and legs, is especially challenging considering the volume and weight constraints demanded by the human form. Mechanisms that use novel topologies, enabled by kinematics theory and a deep understanding of material properties, need to be developed and incorporated. Another important concern for prosthetics is how they will be powered. The power-to-weight ratio of conventional (electromechanical) actuators is inferior to many other potential technologies, such as shape memory/superelastic alloys and direct chemical to mechanical energy conversion (e.g., monopropellants). However, many new actuator technologies are problematic because of safety reasons, slow reaction times, and difficulties in accurate control. There is a need to continue to explore and develop these and other potential robot actuators.

Image courtesy of Georgia Institute of Technology

4.7 Learning and Adaptation

As discussed in Section 3.2.4, the ability for a system to improve its performance over time, and to improve the user's performance, are key goals of medical and health robotics. Toward this end, dedicated work is needed in machine learning applied to real-world uncertain and multimodal medical and health data and moving beyond specific narrow domains toward more comprehensive user health models. Such learning algorithms must ensure guaranteed levels of system performance (safety, stability, etc.) while learning new policies, behaviors, and skills. This is especially important in long-term and life-long user modeling and task learning, both major goals of assistive systems. Growing efforts in the domain of learning and skill acquisition by teaching, demonstration, and imitation need to be directed toward real-world medical and health domains, again using real-world uncertain and time-extended data for grounding in relevance. In general, learning and adaptation to users, to environments, and to tasks and time-extended interactions should become a standard component of usable and robust intelligent robotic systems of the near future.

4.8 Physical Human-Robot Interaction

Physical human-robot interaction is inherent in most medical applications. As described earlier, such interactions require appropriate sensing, perception, and action. Sensing the human could be based on conventional robot sensors or biocompatible/implantable sensors such as brain-machine interfaces. Such sensor data must be combined with modeling to enable perception. Modeling and/or simulation of human form and function are the basis for the design of robots that come into physical contact with humans. Much work needs to be done in this area, since we do not fully understand what models of humans are useful for optimizing robot design, perception, control and planning.

An important aspect of the physical contact between humans and robots is haptics (the technology of touch). When clinicians or patients use robots to interact with environments that are remote in distance or scale, the operator needs to have a natural interface that makes the robot seem "transparent." That is, the operator of a surgical robot, prosthesis, or rehabilitation robot should feel as if he or she is directly manipulating a real environment rather than interacting with a robot. Haptic (force and tactile) displays give feedback to the user that is akin to what he or she feels in the real world. This haptic feedback can improve performance in terms of accuracy, efficiency, and comfort.

4.9 Interaction Algorithms for Socially Assistive Robots

Effective social interaction with a user (or a set of users) is critically important for enabling medical and health robotics to become useful in improving health outcomes in convalescence, rehabilitation, and wellness applications. The user's willingness to engage with a socially assistive robot in order to accept advice, interact, and ultimately alter behavior practices toward the desired improvements, rests directly on the robot's ability to obtain the user's trust and sustain the user's interest. Toward that end, user interfaces and input devices must be developed that are easy and intuitive for a range of users, including those with special needs. Wearable sensors, wands, and other increasingly ubiquitous interaction modalities will be leveraged and further advanced, along with gesture, facial and physical/movement expression, and other means of embodied communication. Social interaction is inherently bidirectional and thus involves both multimodal perception and communication, including verbal and non-verbal means. Thus, automated behavior detection and classification, and activity recognition, including user intent, task-specific attention, and failure recognition, are critical enabling components for HRI. Research into the role of personality and its expression, as well as automated understanding of emotion and believable expression of emotion through multiple channels (voice, face, body) are necessary in order to facilitate real-time believable human-machine interaction.

> ***"Modeling and/or simulation of human form and function are the basis for the design of robots that come into physical contact with humans."***

4.10 Modeling, Simulation, and Analysis

A variety of models are important for medical and health robotics applications. We can divide these into two main categories relevant to medical and health robotics: people modeling (from tissue biomechanics to human cognitive and physical behavior) and engineered systems modeling (including information integration/flow, and open architectures and platforms). The models can be of biomechanics, physiology, dynamics, environment, geometry, state, interactions, tasks, cognition, and behavior. The models can be used for many tasks, including optimal design, planning, control, task execution, testing and validation, diagnosis and prognosis, training, and social and cognitive interaction.

> ***"Research into the role of personality and its expression is necessary in order to facilitate real-time believable human-machine interaction."***

We now provide some specific examples of models needed for medicine and healthcare. In teleoperated (remote) surgery with time delays, models of the patient are required to allow natural interaction between the surgeon and the remote operating environment. Tissue models, in general, are needed for planning procedures, training simulators, and automated guidance systems. These are just beginning to be applied in needle-based operations, but more sophisticated models would enable planning and context-appropriate guidance for a wider variety of procedures, such as laparoscopic surgery and cellular surgery. Models that are sufficiently realistic to be rendered in real time would enable high-fidelity surgical simulations for general training and patient-specific practice conducted by surgeons. For assistive healthcare robots, we need models of human cognition and behavior in order to provide appropriate motivational assistance. Physical models of a patient's whole body are also needed for a robot to provide physical assistance for tasks such as eating or getting out of bed.

As another example, consider a rehabilitation system that uses robotic technology for early and accurate diagnosis. Such a system would need models of the patient and hisor her deficit in order to design appropriate treatments and accurately assess outcomes. (Ideally, the model of the patient would change after treatment.) Such models are also needed for robotic technology to participate in and augment diagnosis. For understanding human activity in context, such as assessing the accuracy and effectiveness of rehabilitation exercises or daily activity, complex models are needed which effectively capture the user's abilities (based on baseline assessment, age, level of deficit, etc.), and can be used to classify and analyze activity being performed (effectively recognize exercise from other activity) combined with the user's state (is the heart rate in the right range, is the user unduly frustrated, etc.) in order to assess progress (is exercise performance improving, is endurance increasing, is accuracy improving, etc.) and provide appropriate coaching. Both activity and physiologic state are complex signals that require modeling to facilitate classification and prediction. Both population models and individual models are needed for addressing challenging problems of online real-time human state and activity detection, classification, and prediction.


Image courtesy of Intuitive Surgical

5. Roadmap Process

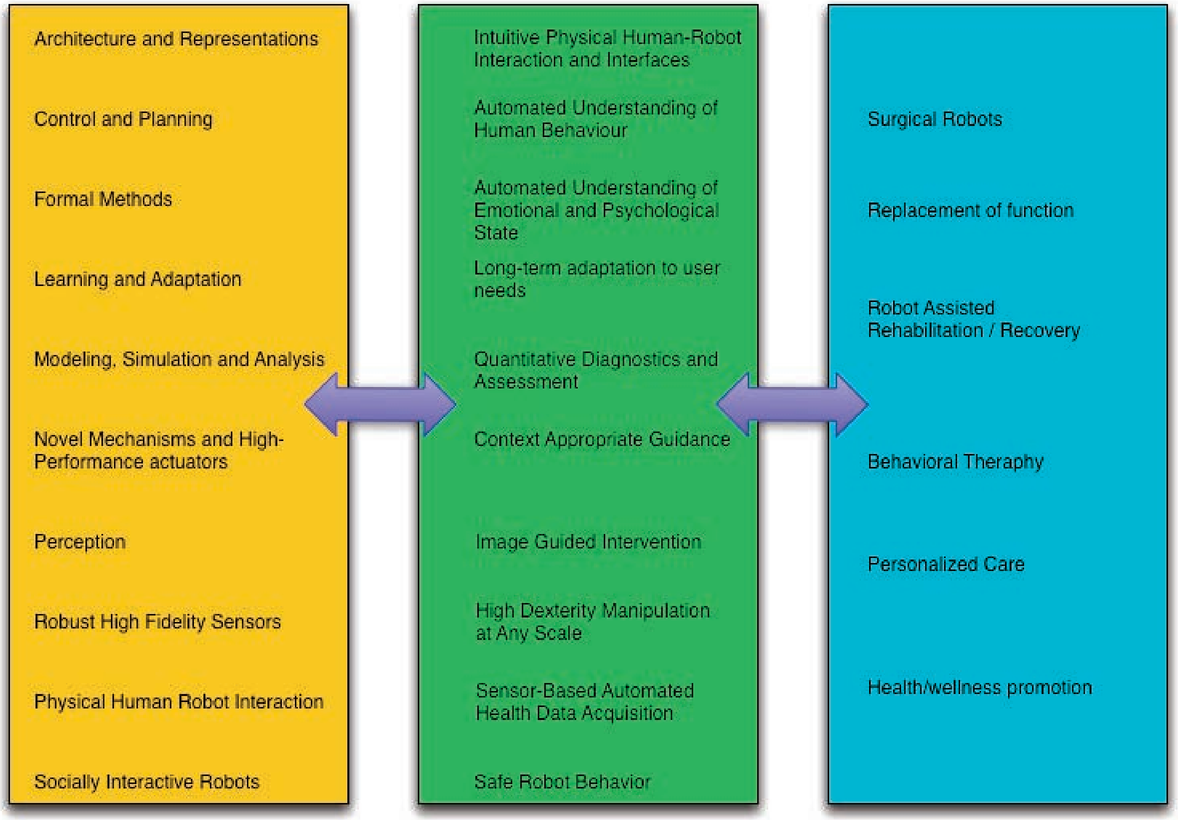

Figure 1: *The roadmap process: Research and development is needed in technology areas that arise from the critical capabilities required to impact healthcare and medical application domains.*

6. Contributors

This document is based on "A Research Roadmap for Medical and Healthcare Robotics," a chapter of the 2009 report *A Roadmap for U.S. Robotics*, and a workshop held on July 23, 2012 at the University of Southern California–University Park Campus. The workshop was sponsored by the U.S. National Science Foundation (NSF) and the USC Viterbi School of Engineering.

Listed below are 26 researchers, clinicians, and industrial representatives who attended the workshop or otherwise contributed to this document. This document was edited by Maja Matarić, Allison M. Okamura, and Henrik I. Christensen, and co-edited by Andrea Thomaz.

Jake Abbott
University of Utah

Howie Choset
Carnegie Mellon University

Henrik I. Christensen
Georgia Institute of Technology

Mihir Desai
University of Southern California

Simon DiMaio
Intuitive Surgical

Henrik I. Christensen
Georgia Institute of Technology

Pierre Dupont
Harvard Children's Hospital

Rod Grupen
University of Massachusetts–Amherst

Gregory Hager
Johns Hopkins University

Blake Hannaford
University of Washington

John Hollerbach
University of Utah

Charlie Kemp
Georgia Institute of Technology

Venkat Krovi
State University of New York–University at Buffalo

Katherine Kuchenbecker
University of Pennsylvania

Art Kuo
University of Michigan

Gerald Loeb
University of Southern California

Maja Matarić
University of Southern California

Bilge Mutlu
University of Wisconsin–Madison

Marcia O'Malley
Rice University

Allison Okamura
Stanford University

David Rinkensmeier
University of California–Irvine

Brian Scassellati
Yale University

Nabil Simaan
Vanderbilt University

Reid Simmons
Carnegie Mellon University

Andrea Thomaz
Georgia Institute of Technology

Kate Tsui
University of Massachusetts – Lowell

Michael Zinn
University of Wisconsin–Madison

Roadmap for Service Robotics

1. Introduction

Service robotics is defined as those robotics systems that assist people in their daily lives at work, in their houses, for leisure, and as part of assistance to the handicapped and elderly. In industrial robotics, the task is typically to automate tasks to achieve a homogenous quality of production or a high speed of execution. In contrast, service robotics' tasks are performed in spaces occupied by humans and typically in direct collaboration with people. Service robotics is normally divided into professional and personal services.

Professional service robotics includes agriculture, emergency response, pipelines, and the national infrastructure, forestry, transportation, professional cleaning, and various other disciplines. [Professional service robots are also used for military purposes but their application in this area is not included in this report.] These systems typically augment people for execution of tasks in the workplace. According to the IFR/VDMA World Robotics, more than 110,000 professional robots are in use today and the market is growing rapidly every year. Several typical professional robots are shown in Figure 1.

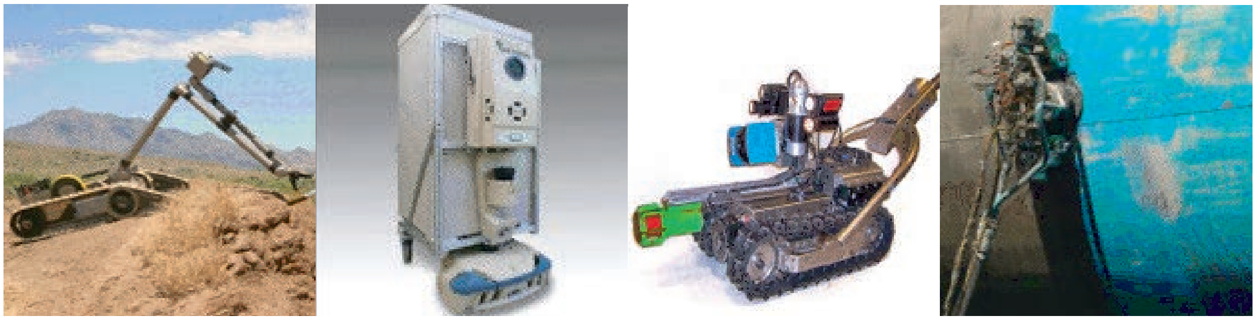

PackBot™ Tactical Robot *AETHON Logistics R Robotobot* *Responder™ Pipeline Robot* *Envirobot™ Paint Stripping Robot*

Figure 1: *Typical service robots for professional applications.*

Personal service robots, on the other hand, are deployed for assistance to people in their daily lives in their homes, or as assistants to them for compensation for mental and physical limitations. By far, the largest group of personal service robots consists of domestic vacuum cleaners; over 6 million iRobot Roombas alone have been sold worldwide, and the market is growing 60%+ each year. In addition, a large number of robots have been deployed for leisure applications such as artificial pets (AIBO), dolls, etc. With more than 4 million units sold over the last 5 years, the market for such leisure robots is experiencing exponential growth and is expected to remain one of the most promising in robotics. A number of typical personal service robot systems are shown in Figure 2. The total sales in the service robotics market during 2011 were estimated to have a value of $4.2B.

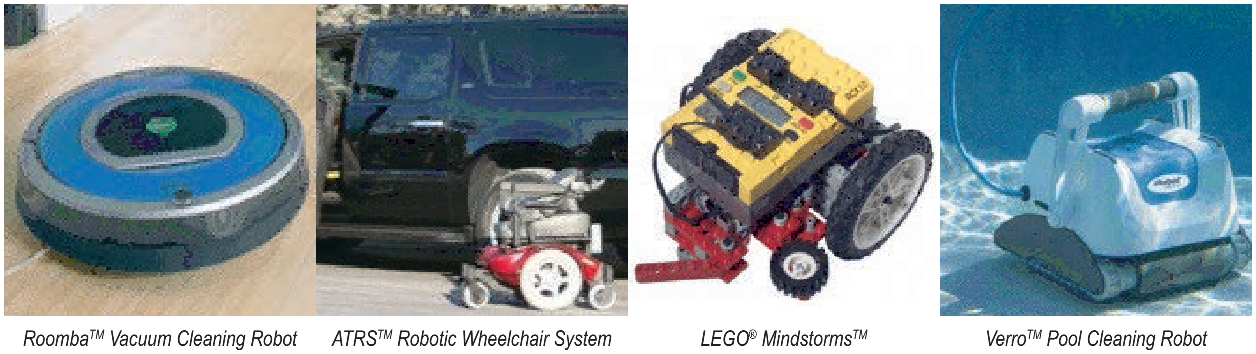

Figure 2: *Typical service robots for personal applications.*

The service robots panel included both professional and personal services and, as such, covered a highly diverse set of applications and problems.

2. Strategic Findings

After much discussion, there was 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. The participants identified several markets where these early commercial solutions are appearing and where service robotics is likely to

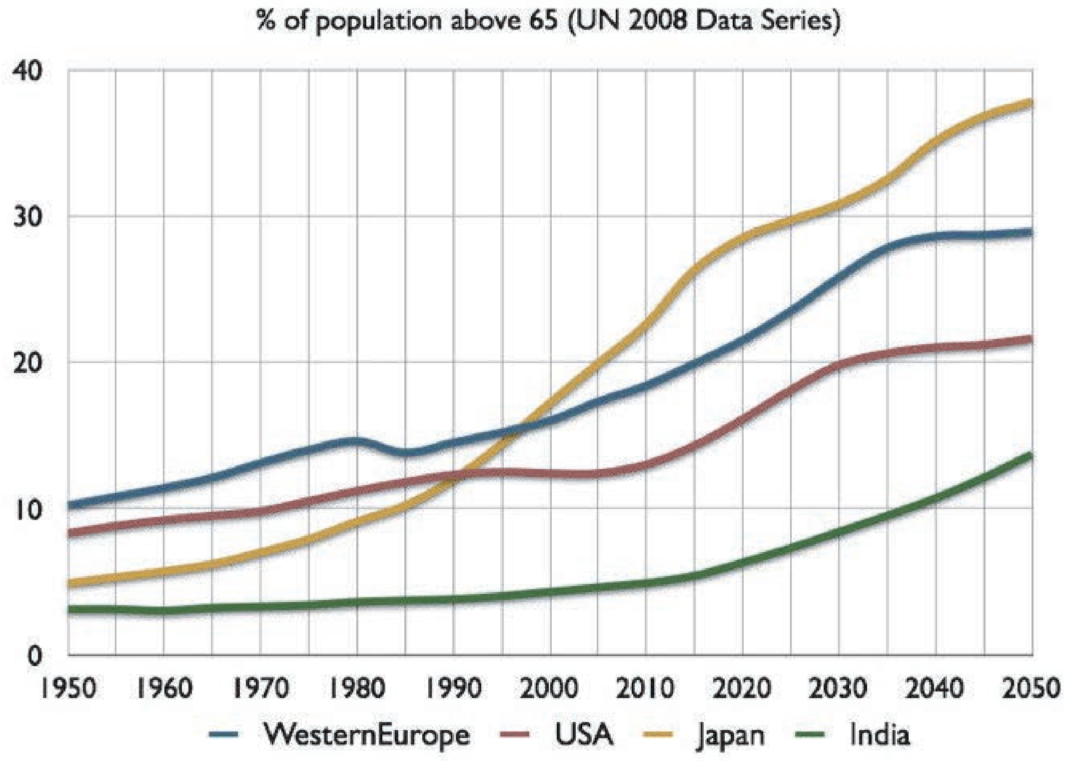

Figure 3: *The changes in demographics in Western Europe, USA, Japan, and India, respectively.*

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 figure 3, 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.

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. For example, a key, primary driver in using robotics technology to automate the automobile factories was the desire to obtain consistent, day-to-day quality and avoid the “built on Monday” syndrome.

2.1 Principal Markets and Drivers

Healthcare & Quality of Life—the current application of robotics technology to provide tele-operated solutions, such as Intuitive Surgical’s *da Vinci®* Surgical System, represents the tip of the iceberg. Robotics technology holds enormous potential to help control costs, empower healthcare workers, and enable aging citizens to live longer in their homes.

Energy & Environment—the attendees identified these two closely linked issues as both critical to the future of our country and ripe for the emergence of robotics technology applications, especially in the areas of automating the acquisition of energy and monitoring the environment.

Manufacturing & Logistics—beyond the traditional application of robotics technology to automate certain assembly line functions, the meeting participants agreed that there is tremendous potential to further automate the manufacture and movement of goods, as fully explored in the parallel roadmapping effort in this area. In particular, robotics technology promises to transform small scale (or “micro”) manufacturing operations and, in the process, help accelerate the transition of manufacturing back to America. This belief has since been substantiated by the formation of a new start-up robotics company, Heartland Robotics, organized specifically for that purpose.

Automotive & Transportation—although we are still decades away from the fully autonomous automobile, robotics technology is already appearing in the form of advanced driver assistance and collision avoidance systems. Public transportation is another area that is expected to 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.

Homeland Security & Infrastructure Protection—participants in the meeting agreed that robotics technology offers tremendous potential for applications in border protection, search and rescue, port inspection and security, and other related areas. In addition, robotics technology is expected to be increasingly used to automate the inspection, maintenance, and safeguarding of our nation’s bridges, highways, water and sewer systems, energy pipelines and facilities, and other critical components of our nation’s infrastructure.

Entertainment & Education—this area, perhaps more than any other, has seen the early emergence of robotics technology enabled products. In particular, robotics has the potential to significantly address the science, technology, engineering, and math ("STEM") crisis facing the nation and to become the veritable "Fourth R" of education. This is evidenced by the tremendous success of FIRST, a non-profit organization founded in 1989 that runs national robotics competitions to inspire young people to be science and technology leaders, and other robotics inspired educational initiatives. Robotics provides kids with a compelling and tactile avenue to learn and apply both the underlying key mathematics and science fundamentals as well as the engineering and system integration principles required to produce intelligent machines to accomplish certain missions.

2.2 Near-Term Opportunities and Factors Effecting Commercialization

Significant investment is required for expanded research and development of robotics technology if the full promise of what can be achieved in each of the above areas is to be realized. As noted above, we are still a long way from the fully autonomous robotics technology required to automate processes to the extent that no human attention or intervention is required. That said, it was the collective opinion of those in attendance that enough progress in robotics technology has been made to enable the development and marketing of a wide variety of initial applications and products in each of these areas to achieve significant levels of "human augmentation."

Such solutions will be capable to varying degrees of automatically performing the following types of functions: monitoring defined, yet dynamic physical environments, identifying objects, detecting changes, or otherwise perceiving the status of their assigned environments, analyzing and recommending actions that should be taken in response to detected conditions, taking such actions in response to human commands, and/or automatically performing such actions within certain pre-authorized boundaries not overridden by human operators.

Examples of such robotics solutions today include teleoperated systems such as the *da Vinci*® Surgical System and autonomous, specialized productivity tools such as the Roomba. As the Internet continues to evolve, it will inspire a natural progression from sensing at a distance to taking action at a distance. This extension of the Internet into the physical world will serve to further blur the boundaries among community, communication, computing, and services and inspire new dimensions in telecommuting and telepresence applications. Hybrid solutions are likely to emerge that enable distributed human cognition and enable the efficient use of human intelligence. Such solutions will combine the robotics-enabled capability to remotely and autonomously perceive situations requiring intervention with the Internet-enabled capability for human operators to take action from a distance on an as-needed only basis.

As referenced above, our aging population will result in a future labor shortage. As workers seek to move up the job hierarchy, there will be a growing need to augment and increasingly automate jobs at the bottom because the workers to perform them may not be readily available and eventually may not exist. While the challenge of achieving fully autonomous solutions in the long run remains primarily technological, the challenge in the near term is one of investing in the science of developing requirements and otherwise determining how to best "cross the chasm." It is one of identifying the right value propositions, driving down costs, developing efficient, effective systems engineering processes, determining how to best integrate such solutions into current or adapted processes, and otherwise addressing the know-how gap of transitioning technology into products.

2.3 Scientific and Technical Challenges

Workshop participants worked in break-out groups to identify technical and scientific challenges pertinent to the applications and business drivers described in the previous section. This section summarizes their findings. We will present the technical and scientific challenges identified by the break-out groups in an integrated manner. The emphasis of this section is on describing the challenges, not on laying out a roadmap towards addressing these challenges—such a roadmap will be outlined in the next section.

2.3.1 Mobility

Mobility has been one of the success stories of robotics research. This success is exemplified by a number of systems with demonstrated performance in real-world environments, including museum tour guides and autonomously driving cars, as in the DARPA Grand Challenge and Urban Challenge. Nevertheless, workshop participants agreed that a number of important problems remain open. Finding solutions to these problems in the area of mobility will be necessary to achieve the level of autonomy and versatility required for the identified application areas.

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 everyday, populated environments that are more unstructured and less controlled, however, these 2D representations will not be sufficient to capture all aspects of the world necessary for common tasks. It will therefore be important to enable the acquisition of three-dimensional world models in support of navigation and manipulation (see next section). 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.

One of the promising technologies towards semantic 3D mapping, as identified by the participants, is using different kinds of sensors for building maps. Currently, robots rely on high-precision, laser-based measurement systems or game-console ranging sensors such as the Microsoft Kinect or the PrimeSense for learning about their environment, using mapping algorithms known as “SLAM” algorithms. The participants identified a desire to move away from lasers to cameras, to further develop the field of “visual SLAM” (VSLAM). This technology relies on cameras, which are robust, cheap, and readily available sensors, to map and localize in a three-dimensional world. Already today, VSLAM systems exhibit impressive real-time performance. Participants therefore believed that VSLAM will likely play a role in the development of adequate and more affordable 3D navigation capabilities.

Participants identified additional requirements for 3D navigation that will be critical to meet the requirements of targeted applications. **Outdoor 3D navigation** poses a number of important challenges that have to be addressed explicitly. Among them is the fact that current 2D environmental representations cannot capture the complexity of outdoor environments nor the changing lighting conditions that cause substantial variability in the performance of sensor modalities. Participants also identified robust **navigation in crowds** as an important mobility challenge.

2.3.2 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 that emphasize planning and control—**perception** becomes an important component of the research agenda toward autonomous manipulation.

> ***"In contrast to most existing methods that emphasize planning and control, perception is an important component of the research agenda toward 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.

Participants suggested that competent **"pick and place"** operations may provide a sufficient functional basis for the manipulation requirements of a many of the targeted applications. It was therefore suggested that pick and place operations of increasing complexity and generality could provide a roadmap and benchmark for research efforts in autonomous manipulation.

2.3.3 Planning

Research in the area of motion planning has made notable progress over the last decade. The resulting algorithms and techniques have impacted many different application areas. Nevertheless, participants agreed that robust, **dynamic 3D path planning** remains an open problem. An important aspect of this problem is the notion of a robot's **situational awareness** (i.e., the robot's ability to autonomously combine, interleave, and integrate the planning of actions with appropriate sensing and modeling of the environment). The term "appropriate" alludes to the fact that complete and exact models of the environment cannot be acquired by the robot in real time. Instead, it will be necessary to reason about the objectives, the environment, and the available sensing and motor actions available to the robot. As a result, the boundary between planning and motion planning is blurred. To plan a motion, the planner has to **coordinate sensing and motion under the constraints imposed by the task**. To achieve task objectives robustly and reliably, planning has to consider **environmental affordances**. This means that the planner has to consider interactions with the environment and objects in it as part of the planning process. For example: to pick up an object, it may become necessary to open a door to move into a different room, to push away a chair to be able to reach to a cabinet, to open the cabinet door, and to push an obstructing object out of the way. In this new paradigm of planning, the **task** and **constraints imposed by the task and the environment** are the focus; the “motion” of “motion planning” is a means to an end. Constraints considered during planning can arise from **object manipulation, locomotion (e.g. footstep planning), kinematic and dynamic constraints of the mechanism, posture constraints, or obstacle avoidance**. Planning under these constraints must occur in **real time**.

Some of the constraints on the robot’s motion are most easily enforced by leveraging sensor feedback. Obvious examples are **contact constraints** and **obstacle avoidance**. The area of **feedback planning** and the integration of **control and planning** are therefore important areas of research towards satisfying the planning requirements identified by the participants. A feedback planner generates a policy that directly maps states to actions, rather than generating a specific path or trajectory. This ensures that sensor, actuation, and modeling uncertainties can adequately be addressed using sensory feedback.

The increased complexity of planning in this context will also require novel ways of capturing **task descriptions**. While in classical motion planning the specification of two configurations fully specified a planning task, the view of planning described here has to handle much richer task representations to address the richness of manipulation tasks and intermediate interactions with the environment.

Participants also perceived the need for formal methods to perform **verification and validation** of the results of planners. Such guarantees may be required to ensure safe operation of robots in environments populated with humans.

2.3.4 Sensing and 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 a 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 obscurants 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, and 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 dextrous manipulation for performing tasks in human environments. Creating contextual models that are **situation-aware** is important to be considered in robotics perception algorithms.

2.3.5 Architectures, Cognition, and Programming Paradigms

The discussions on the topics of mobility, manipulation, planning, and perception revealed that these issues cannot be viewed in isolation but are intricately linked to each other. The question of how to engineer a system to effectively integrate **specific skills from those areas to achieve safe, robust, task-directed, or even intelligent behavior** remains an open question of fundamental importance in robotics. Research towards this objective has been conducted under the name of **architectures, cognition, and programming paradigms**. This diversity in approaches or even philosophical viewpoints may reflect the lack of understanding in the community on how to adequately tackle this challenge. This diversity of viewpoints is also reflected in the diversity of tools currently brought to bear on this issue: they range from imitation learning to explicit programming of so-called cognitive architectures. Some participants felt that a mixture of these would probably be required to achieve the desired outcome.

One of the classical approaches towards the overarching issue of generating robust, autonomous behavior is the **sense/plan/act loop** usually employed by modern control systems. While sense/plan/act has been a constant in robotics research over the last several decades, some participants felt that novel approaches would likely deviate from this approach in its simplest form. Possible alternatives are multiple nested or hierarchical loops, the behavior-based approach, combinations of the two, or possibly even completely novel approaches.

All participants agreed that this area of investigation will require substantial attention and progress on the path towards autonomous robotic systems.

2.3.6 Human-Robot Interaction (HRI)

Given the ultimate goal of deploying mobile and dexterous robots in human environments to enable coexistence and cooperation, substantial progress will be required in the area of human-robot interaction. These interactions could also become an important component in an overarching approach to robust robot behavior, as discussed in the previous subsection. Robots might learn novel skills from their interactions with humans, but under all circumstances, should be cognizant of the characteristics and requirements of their communication with humans.

In addition to the **modes of communication** (verbal, nonverbal, gesture, facial expression, etc.), participants identified a number of important research topics, including **social relationships, emotions** (recognition, presentation, social emotional cognition/modeling), **engagement**, and **trust**. An understanding of these aspects of human-robot communication should lead to an automatic structuring of the interactions between humans and robots where robotic systems' ability to operate independently rises or falls automatically as both the task and the human supervisor's interaction with the system changes.

Progress towards these objectives will depend on **effective input devices** and **intuitive user interfaces**. Participants also advocated the development of a variety of **platforms** to study HRI, including humanoid robots, mobile manipulation platforms, wheelchairs, exoskeletons, and vehicles. Participants identified a **design/build/deploy cycle** in which HRI research should progress. The design process should consider input from a number of relevant communities, including the basic research community and end users. The build process integrates numerous components and research threads into a single system; here there is an opportunity for industry collaborations and technology transfer. Finally, the integrated system is deployed in a real-world context. Participants suggested the notion of a **Robot City** (see next subsection) as a promising idea to evaluate HRI in a real-world context. The cycle is closed by incorporating end user feedback into the experimental design of the next iteration of the design/build/deploy cycle.

2.3.7 Research Infrastructure

Workshop participants felt strongly that rapid progress towards the identified scientific objectives will critically depend on the broad availability of adequate research infrastructure, including hardware and software. To address the research challenges given above, it will be necessary to construct robotic platforms that combine many advanced and interacting mechanical components, providing adequate capabilities for mobility, manipulation, and sensing. Theses platforms will be controlled by a multitude of independently developed, yet interdependently operating, software components. As a result, these integrated robotic platforms exhibit a degree of complexity that is beyond what can easily be designed, developed, tested, and maintained by many independently operating research groups. The lack of standardization of hardware and software platforms may also result in a fragmentation of the research community, difficulties in assessing the validity and generality of published results, and the replication of much unnecessary engineering and integration effort.

> ***"Under all circumstances, robots should be cognizant of the characteristics and requirements of their communication with humans."***

To overcome these challenges, workshop participants advocated coordinated community efforts for the development of hardware and software systems. These efforts should include the development of an **open experimental platform** that would—preferably at low cost—support a broad range of research efforts on the one hand, while enabling **technology and software reuse** across research groups on the other hand. One example of such an open platform is ROS, a robot operating system being developed by Willow Garage that enables code reuse and provides the services one would expect from an operating system, such as low-level device control, implementation of commonly-used functionality, and message-passing between processes. Ideally, such platforms would be complemented by **physical simulation software** to support early development and testing of algorithms without compromising the safety of researchers and hardware. Development efforts could also benefit from **robotic integrated development environments (IDEs)**; these IDEs enforced modularity in software development thereby facilitating reuse and documentation.

Participants noted that research in robotics is rarely thoroughly evaluated and tested in well-defined, repeatable experiments. Other fields, such as computer vision, have greatly benefited from publicly available data sets, which enabled an objective comparison between multiple algorithms and systems. The participants therefore suggested the creation and expansion of **repositories of experimental data**, which could then serve as community-wide benchmarks. However, as much of the research in robotics is focused on the physical interaction between the robot and its environment, electronic data sets are not sufficient. They should be complemented by **skill-specific benchmarks consisting of physical objects**. For example, a number of readily available objects can be selected as a benchmark for grasping research. Furthermore, entire **benchmark environments** were suggested to develop, evaluate, and compare the performance with respect to a particular application or implementation. Such environments could range in size and complexity from a simple workspace (an office desk or a kitchen counter) to an entire room, a house, or an entire city block. In this context, the notion of a **Robot City** was mentioned: a regular urban environment in which all inhabitants are part of the experiment and help in the evaluation process as well as with the definition of adequate requirements for everyday application environments.

Many of the proposed efforts—and in particular **hardware or software integration** efforts—fall outside of the scope of existing funding programs. Participants noted that a policy change in this regard would be necessary to ensure that the availability of research infrastructure does not represent a bottleneck in the progress towards autonomous robotic systems in everyday environments.

2.3.8 Mechanical Hardware

Safety is a critical factor for the deployment of robotic systems in human environments. Inherently safe robots would also enable modes of human-robot interaction that can increase acceptance of robotic technology in everyday life. Participants therefore felt that **inherently safer motors and mechanisms** with increased strength to weight ratio would represent an important enabling technology. In such mechanisms, **variable compliance** would be a desirable property. The concept of variable compliance refers to a mechanisms ability to adjust its behavior to reaction forces when contacting the environment. These reaction forces can be varied for different tasks. Such mechanisms enable safe operation, especially when interacting with humans, as well as flexible, robust, and competent motion when in contact with the environment. Furthermore, **energy efficiency** was identified as a critical concern for many applications, as robots will have to operate without tethers for extended periods of time. Finally, **novel or improved modes of locomotion beyond wheels** are needed to enable safe and reliable operation in indoor and outdoor environments. Outdoor environments oftentimes exhibit highly variable terrain properties while outdoor may contain stairs, ladders, ramps, escalators, or elevators.

> ***"Inherently safe robots would enable modes of human-robot interaction that can increase acceptance of robotic technology in everyday life."***

Participants identified **highly dexterous and easily controllable robotic hands** as an important area for research. Progress in robotic grasping and manipulation very likely will go hand in hand with the development of novel hand mechanisms. At the same time, participants felt that the potential of current hand technology was not fully leveraged by existing grasping and manipulation algorithms. It is therefore conceivable that many interesting and relevant applications can be addressed with available grasping and manipulation hardware.

3. Key Challenges/Capabilities

3.1 Motivating Scenarios

3.1.1 Quality of Life


ATRS™ Robotic Wheelchair System

Robotics technology is expected to make a tremendous contribution to the lives of the elderly and disabled. One such example of an existing application is a revolutionary transportation mobility solution that enables those with limited mobility who use wheelchairs to independently get into and out of their vehicles and remotely load and unload their wheelchairs from a wide range of vehicles. This system makes it possible for those dependent on wheelchairs to transport their wheelchair using an ordinary passenger van and to access it whenever needed without assistance from others, thus offering them a degree of freedom and independence heretofore unavailable. This system provides significant benefits over existing transportation mobility solutions, including lower cost of ownership, ability to use standard crash-tested automotive seats, greater choice of vehicles, no required structural modifications, and the ability to reinstall on subsequent vehicles.

3.1.2 Agriculture


Autonomous Tractor

Robotics technology is expected to impact a myriad of applications in agriculture and address farmers' constant struggle to keep costs down and productivity up. Mechanical harvesters and many other agricultural machines require expert drivers to work effectively, while factors such as labor costs and operator fatigue increase expenses and limit the productivity of these machines. Automating operations such as crop spraying, harvesting, and picking offer the promise of reduced costs, increased safety, greater yields, increased operational flexibility, including nighttime operations, and reduced use of chemicals. A number of such prototype systems and applications, including automated fruit crop spraying and field crop harvesting, have been developed and the technology has now matured to the point where it is ready to be transitioned for further commercialization and field deployment within the next few years.

3.1.3 Infrastructure


Responder™ Pipeline Robot

Robotics technology has tremendous potential to automate the inspection and maintenance of our nation's bridges, highways, pipelines, and other infrastructure. 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. These robotic platforms navigate this critical wastewater infrastructure to inspect sewer pipe unreachable by traditional means and produce very accurate 3D images of the pipe's inside surface. 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.

3.1.4 Mining

Autonomous Haul Truck

Robotics technology is already starting to have a dramatic impact on both the underground and surface mining industries. An innovative belt inspection system that uses a high-speed "machine vision" system and software algorithms to monitor the condition of conveyor belts and help operators detect defects, for example, is in everyday use at several underground coal mines. The patented system is designed to reduce costly downtime caused by the degradation and eventual rupture of conveyor belt splices. On a larger scale, robotics technology is being used to develop autonomous versions of large haul trucks used in mining operations. Caterpillar recently announced that it is developing an autonomous mining haulage system with plans to integrate autonomous haul trucks, each with payload capacities of 240 tons or more, into some mine sites by 2015. The autonomous technology is designed to provide productivity gains through more consistency in processes and minimize environmental impact by both improved efficiency and overall mine safety.

3.1.5 Transportation

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. Meanwhile, carmakers and device manufacturers are creating "smart cars" by installing more powerful computers and sensors, giving drivers a better idea of their environment and car performance.

Top Three Finishers in the 2008 DARPA Urban Grand Challenge

States such as Nevada and Florida have passed legislation that allows deployment of driver-less cars. In Nevada, two cars have been awarded a driver's license (one from Google and one from Audi). The roads they are driving on have increased in capacity by only 5 percent, resulting in 3.7 billion hours of driver delays and 2.3 billion gallons of wasted fuel. To address this issue, highway agencies are attempting to create "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.

3.1.6 Education

FIRST Lego League™ Participants

Robotics has already commenced transforming the American classroom. Robotics puts academic concepts in context and is being used at all levels in K-12 and college education. Robotics provides students with a tactile and integrated means to investigate basic concepts in math, physics, computer science and other STEM disciplines, while enabling teachers at the same time to introduce concepts about design, innovation, problem solving, and teamwork. Robotics curriculums have been developed, teachers have been trained, and scores of competitions are held every year across the country. Perhaps the best-known robotics competition programs are operated by FIRST, a non-profit organization founded in 1989 to inspire young people to be science and technology leaders. As a measure of the growing popularity of robotics competitions, FIRST is expecting over 320,000 students to participate in its competitions in the coming year. Even more significantly, a recent Brandeis University survey found that FIRST participants are more than twice as likely to pursue a career in science and technology as non-FIRST students with similar backgrounds and academic experiences. Although much progress has been made, the surface has only been scratched in terms of the potential impact of robotics in education. To more fully realize this potential, robots need to be made more accessible, affordable and easy to use for both students and teachers.

3.1.7 Homeland Security and Defense

Disaster Site Application

The use of robotics technology for homeland security and defense continues to grow as innovative technology has improved the functionality and viability of search and rescue efforts, surveillance, explosives countermeasures, fire detection, and other applications. Unmanned surveillance, detection, and response systems will be able to make use of robotic platforms, fixed sensors, and command and control networks to potentially monitor and patrol hundreds of miles of rough border terrain, to sniff out and locate chemical/biological/radioactive/nuclear/explosive threats, and survey large perimeters associated with borders, power plants, or airports. Such systems will enable security personnel to automatically detect potential threats, to take a close-in first look from a safe distance, and to provide initial disruption and interdiction at the point of intrusion if necessary. Other "man-packable" robots equipped with instruments including infrared cameras, night vision sensors, and millimeter-wave radar have been used to search for victims at disaster sites, including the World Trade Center. *(Please see the separate chapter on Defense robotics.)*

3.2 Capabilities Roadmap

In the following, we identify the key challenges that have to be met and the key capabilities that have to be developed in order to deliver service robots capable of addressing the aforementioned motivating scenarios. Figure 4 provides an overview of the proposed roadmap and the remainder of this document. The right column in the figure lays out the application areas, many of which are described in the motivating example scenarios in Section 3.1. High-impact advances in these application areas can only be enabled if a number of capabilities for autonomous service robots become available. These capabilities are listed in the middle of the figure and described in more detail in Section 3. To achieve the required level of competency in those areas, sustained investment in research and developments in a number of basic research areas and technologies is required. Figure 4 shows these research areas and technologies in the left column; they are described in more detail in Section 4.

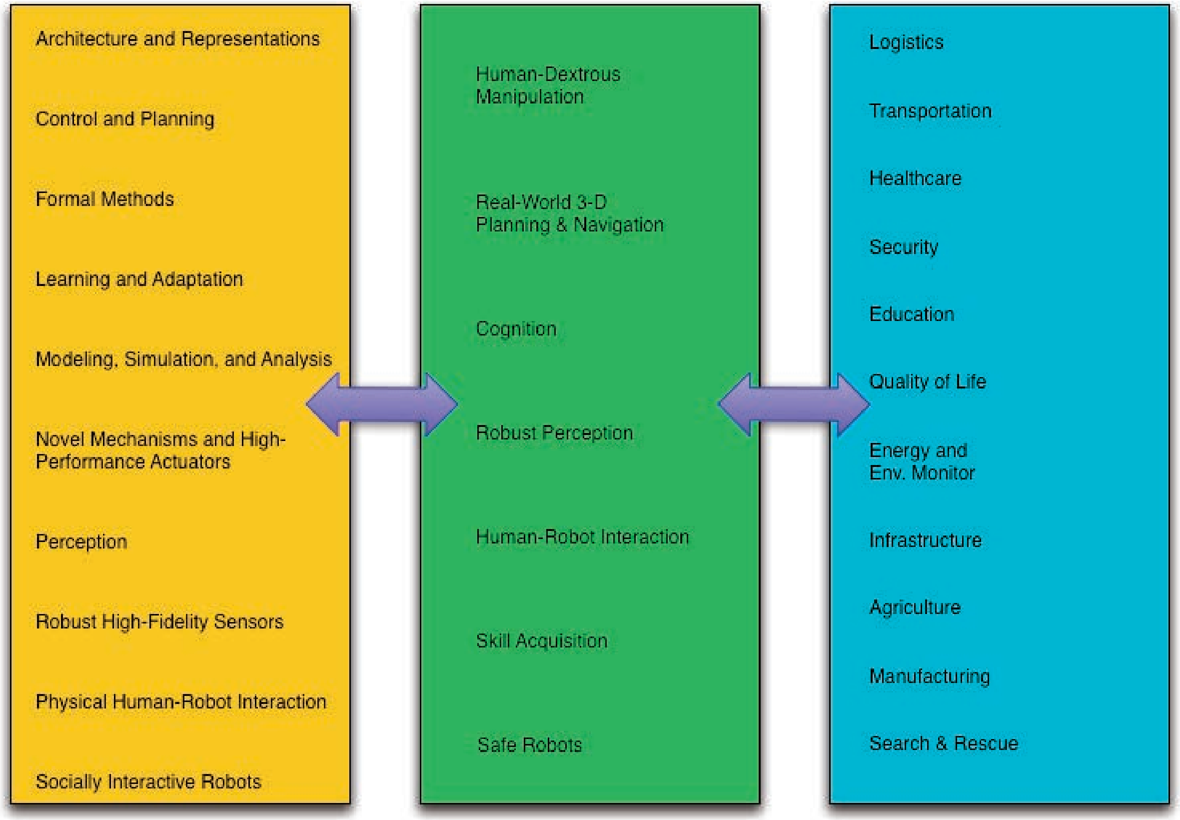

Figure 4: *Overview of the roadmap for domestic and industrial service robotics: Sustained progress in the basic research areas in the left (yellow) column of the figure will enable a number of elementary capabilities, shown in the middle (green) column of the figure. These capabilities, in turn, enable progress in the application areas shown in the right (blue) column.*

3.2.1 Mobility and Manipulation

Autonomous service robots accomplish tasks by moving about their environment and by interacting with their environment. These motions and physical interactions need to achieve a given task by changing the robot's pose and by moving objects in the environment. The accomplishment of a task may require complex sequences of motions and interactions; the robot may have to move from one room to another or it may have to open doors, climb stairs, use elevators, clear obstacles out of its path, remove obstructions, or use tools. To achieve this level of competency, substantial advances at the intersection of perception, control, and cognition must be made. The problems posed by service robotics, however, can only be addressed through integrated solutions to these problems.

Consider the task of a robot walking to a room on a different floor to fetch a box. Depending on the size of the box, finding good grasps, lifting it, navigating it through tight spaces and over steps while avoiding other objects in the environment are challenges that the robot has to overcome. To reason about pushing the box (versus picking it up), the robot has to reason about its own capabilities, the geometry of the scene, constraints imposed by actuation and joint limits, as well as the contact dynamics and friction that arise when moving.

To reason about the world in such a way that the appropriate sequence of actions and motions can be determined, the robot has to be aware of its environment. Not all of the required information can be provided to the robot beforehand, as service robots operate in unstructured and dynamic environments. The robot therefore has to possess capabilities to perceive objects in the environment and estimate their properties. "Semantic mapping" provides the robot with information about the environment that is required to achieve a task. Object detection and recognition and related perceptual skills provide information for semantic mapping, navigation and object manipulation.

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

- **5 years:** 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. Robots create semantic maps about their environment through exploration and physical interaction but also through instruction from humans. They are able to reason about tasks of moderate complexity, such as removing obstructions, opening cabinets, etc. to obtain access to other objects.

- **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.

- **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.

3.2.2 Real-World Planning and Navigation

Since 2009, the area of real-world planning and control for service robotics has advanced in major ways not contemplated in the *A Roadmap for U.S. Robotics* original report. The focus at that time was on operation in highly unstructured situations with limited a priori knowledge of tasks to be performed, navigation space, and obstacles. The robot needed to acquire and reason with sensory data to build a model of the environment in which to do planning, followed by path planning, obstacle avoidance, and task planning. During execution of these plans, additional sensory information would be used as feedback. Even for single or small number of robots in research laboratory settings, planning and control was far from real time.

There has been progress on challenging problems in highly unstructured environments, but the most significant change since 2009 is deployment in several application domains of large fleets of service robots capable of planning and control in real time, where the environment and tasks are more structured. Application domains include logistics and material handling, healthcare, agriculture, and others (Figure 5). These applications are important not only at technical advances, but also demonstration that service robotics has value in solving real-world problems.

Figure 5: *Examples of autonomous mobile service robot fleets deployed in logistics and medical applications (Symbotic, Aethon, Kiva).*

The following attributes help capture the current state of technology:

- Fleet size in the range of 10^3 to 10^4 operating in a single installation.
- Real-time navigation and task planning at rates from 10^4 to 10^5 tasks or transactions per hour.
- Navigation including obstacle and collision avoidance in network topologies defined by maps in 2D and 2.5D, such as storage warehouses and hospital corridors. 2.5D comes from coordination of planning and control task on multiple floors or levels.
- Coordination of tasks across the fleet, to accomplish tasks requiring synchronization of more than one robot.
- Tasks include sensing and manipulation of objects to accomplish tasks for which there is some degree of a priori knowledge, augmented by sensing information when the robot arrives at the local scene.
- Ability to re-plan dynamically to resolve problems, such as task failure, obstacle or collision avoidance, including the presence of humans in the same operating environment.
- Integration of service robot planning and control with higher level IT, such as logistics supply chain or hospital information systems.

Figure 6 captures the capabilities described above for a range of real-time performance metrics and degree of structure in operating environment and tasks. This current roadmap envisions pushing the envelope to the upper right, with increasing capabilities in handling uncertainty and operating complexity.

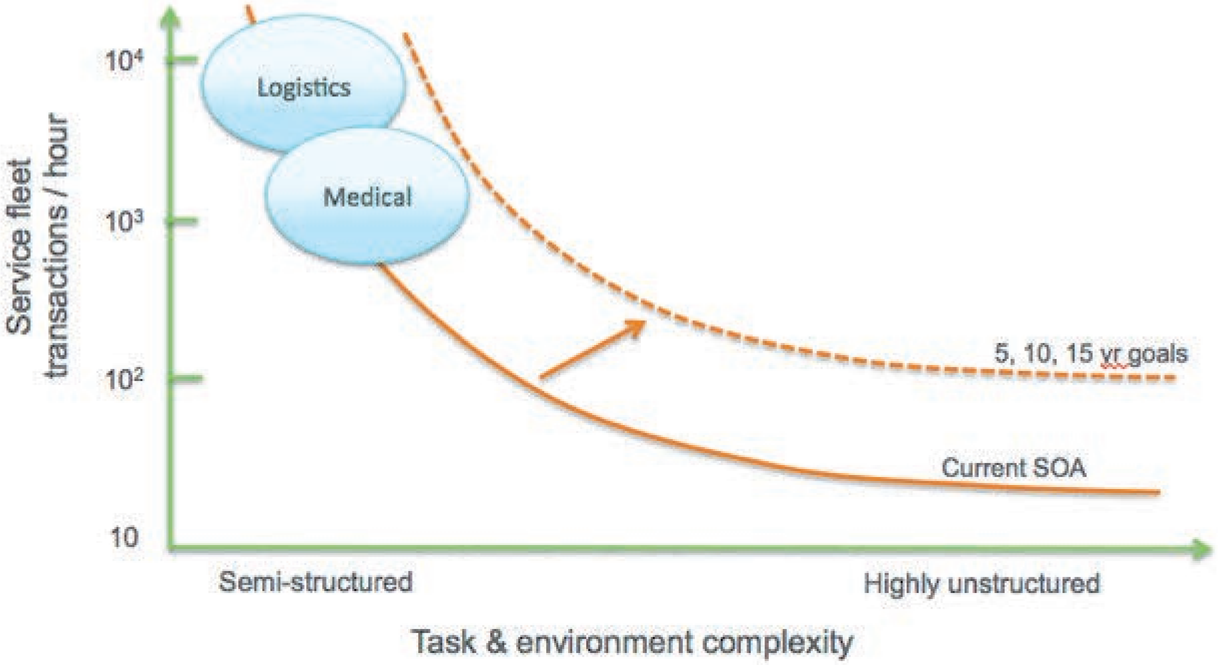

Figure 6: *Real-time performance metrics for service robot fleets will increase over the next decade.*

The following are initiatives that could lead to such advances:

- **Stochastic Planning**—Consider a deterministic planning and control problem where there is no uncertainty in the model. A simulation of the system using the proposed planning and control strategy can only validate the approach within the limits that the simulation is itself representative of the real world. Real-world field trials would reveal areas where the simulation needs to be improved to better represent reality, such as treating parametric data as statistical distributions instead of constants, or adding new dynamic components or failure modes to the model. Further simulation will probe the limits of robustness of the strategy. Over the next five years, a research challenge could be the development of more direct methods for developing planning and control algorithms using characterization of uncertainty from the beginning. A truly robust planner operating in the real world shall be able to build conditional plans taking advantage of statistical knowledge about the environment available a priori. Anticipatory planning will increase both safety and efficiency, as the robot(s) will operate according to plans with built-in uncertainty awareness.

- **Automatic Detection of Current Model Validity**—In the next 5-10, years robots are expected to operate over extended periods of time. In highly dynamic situations, statistical models used at planning time may no longer correctly characterize the environment as time unfolds. Robust methods to detect on-the fly if the system is still operating within the same model or not shall become an integral part of the control loop. A goal for the next 5-10 years is to empower robots with the ability to integrate sensor readings acquired while executing a plan in order to update the underlying statistical model and autonomously decide when and what to re-plan. The goal is the development of systems capable of generating and updating plans enabling uninterrupted, loosely supervised operation for periods of months.

- **Imitation and Transfer Learning**—Service robotics in fleet applications inherently accumulate a rich history of "failures"—uncertainty intrudes in a manner that requires frequent re-planning. Even if the occurrence rate for each class of re-plan is low, we will still see these problems recur. An alternative to solving the re-planning problem fresh each time, using learning, it should be possible to recognize previous occurrences of the problem, and even more important which re-plans proved to be most effective. A 10-year goal is the integration of traditional planning techniques with robust machine learning methods for recognizing re-plan problems that have occurred before and selection of optimal solutions, both in real-time. Numerous results have become available in the field imitation learning and transfer learning. Although planning is eventually performed in a continuous domain characterized by infinite execution trajectories, from a more abstract point of view, it is possible to envision recurrent building blocks, as well as nuisances. The ability to share and reuse plans, as well as contingencies encountered while executing complex action sequences, is needed. In 5-10 years robots shall be able to contribute their plans to local knowledge bases shared with other robots, and to efficiently identify plans, or parts thereof, that can be reused in lieu of solving complex problems from scratch. In the long term, large, efficiently searchable repositories of robot plans shall become available on a planetary scale.

- **Human Operator Supervisory Control**—Consider the analogy to air traffic controllers who, today, have "situation awareness" of hundreds of aircraft in their space, the latter being flown by human pilots with whom they can exchange sensory and command information (Figure 7). This situation will change in the near future with the introduction of drones into a common airspace. In contrast, in current logistics and medical rover fleets, the supervisor has displays of the current state of the system and augmented by limited sensory data, but as fleet sizes get into the range of 10^3 to 10^4, with 10^4 to 10^5 transactions per hour, the supervisor's situation awareness is limited. To understand the situation requires not just the system state, but awareness of the plans for each robot, how they interact with each other and the uncertain changing environment. Advancement in the area could be measured by a "fan out"—the ratio of supervisors to rovers—from current levels around 10 to 10^2 in five years, 10^3 in ten years. It is likely that hierarchical models for the system will be used to enable the supervisor to zoom into situations at appropriate levels, with adjustable autonomy at levels not needing supervisor attention.

Figure 7: *Air traffic controller analogy to human supervision of large fleets of mobile service robots.*

3.2.3 Cognition

In service robotics, there is a need to operate in non-engineered environments, to acquire new skills from demonstration by users, and to interact with users for tasking and status reporting. Cognitive systems enable acquisition of new models of the environment and training of new skills that can be used for future actions. Cognition is essential for fluent interaction with users and deployment in domains where there are limited opportunities for user training. In addition, an added degree of intelligence for coping with non-engineered environment is essential to ensure system robustness.

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

- **5 years:** Demonstration of a robot that can learn skills from a person through gesture and speech interaction. In addition, acquisition of models of a non-modeled indoor environment.

- **10 years:** A robot that interacts with users to acquire sequences of new skills to perform complex assembly or actions. The robot has facilities for recovery from simple errors encountered.

- **15 years:** A companion robot that can assist in a variety of service tasks through adaptation of skills to assist the user. The interaction is based on recognition of human intent and re-planning to assist the operator.

3.2.4 Robust Perception

Service robots operate in relatively 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 focussed 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, and 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 such as homes, highways, hospitals, and warehouses. 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 hundreds 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 systems 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 of 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.

3.2.5 Physical, Intuitive HRI, and Interfaces

Deployment of service robots, both in professional and domestic settings, requires the use of interfaces that makes the systems easily accessible for the users. Diffusion of robotics to a broader community requires interfaces that can be used with no or minimal training. There are two aspects to interfaces: physical interaction with users and people in the vicinity and the command interface for tasking and control of the robot. The physical interaction includes body motion to move/nudge objects and people and non-contact interaction such as change of motion behavior to communicate intent or state. The interface aspect is essential to tasking and status reporting for operators to understand the actions of the robot.

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

- **5 years:** Demonstration of a robot where task instruction is facilitated by multimodal dialog for simple actions/missions and robots that can communicate intent of actions by the body language.

- **10 years:** Demonstration of a robot where programming by demonstration can be used for complex task learning such as meal preparation in a regular home.

- **15 years:** Demonstration of a robot that can be programmed by an operator for complex mission at a time scale similar to the actual task duration.

3.2.6 Skill Acquisition

Service robots must possess the ability to solve novel tasks with continuously improving performance. This requires that service robots be able to acquire novel skills autonomously. Skills can be acquired in many ways: they can be obtained from skill libraries that contain skills acquired by other robots; skills can be learned from scratch or by composing other skills through trial and error; skills can also be learned through observation of other robots or humans; furthermore, they can be taught to a robot by a human or robotic instructor. But skill acquisition also requires the robot to identify those situations in which a skill can be brought to bear successfully. Skills can be parameterized; learning and selecting appropriate parameters for a variety of situations is also included in the capability of skill acquisition. The ability to transfer skills from one domain to another, or to transfer experience acquired with one skill to another skill, can be expected to provide substantial advances in skill acquisition. Advances in perception, representation, machine learning, cognition, planning, control, activity recognition, and other related areas will enable adequate capabilities in skill learning.

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

- **5 years:** Robots can learn a variety of basic skills through observation, trial and error, and from demonstration. These skills can be applied successfully under conditions that vary slightly from the ones under which the skill was learned. Robots can autonomously perform minor adaptations of acquired skills to adapt them to perceived differences from the original setting.

- **10 years:** As perceptual capabilities improve, robots can acquire more complex skills and differentiate specific situations in which skills are appropriate. Multiple skills can be combined into more complex skills autonomously. The robot is able to identify and reason about the type of situation in which skills may be applied successfully. The robot has a sufficient understanding of the factors that affect the success so as to direct the planning process in such a way that chances of success are maximized.

- **15 years:** The robot continuously acquires new skills and improves the effectiveness of known skills. It can acquire skill-independent knowledge that permits the transfer of single skills across different tasks and different situations and the transfer of skills to novel tasks. The robot is able to identify patterns of generalization for the parameterization of single skills and across skills.

3.2.7 Safe Robots

Today, safety for robots is achieved through a clear separation of the workspaces for humans and robots or through operation at speeds that do not represent a risk to humans in the proximity of the system. As the operation of humans and robots become more and more intertwined, there will be a need to explicitly consider operation at higher speeds while operating in direct proximity to people.

Therefore, there is a need to consider standards for safety to enable certification. The current, limited set of standards for safety certification of both professional and personal robots, constrains innovation, reduces the pace of adoption, and adds costs.

> ***"Safety is a multidimensional issue extending beyond technology and including a number of governmental and industry standards."***

While technologically, safety involves several aspects including the need for advanced perception capabilities to detect objects and persons and predict possible safety hazards, control systems that react to possible dangerous situations, and inherently safe actuation mechanisms to ensure that contact with a person or objects causes little or no damage. However, safety is a multidimensional issue extending beyond technology. It includes a number of governmental and industry standards as well as independent certification and liability exposure. These non-technical elements need to progress such that that clear standards exist for both professional and personal robotics providing all stakeholders the visibility needed for rapid innovation and adoption.

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

- **5 years:**
  a. Safety standards for all categories of service robotics has been defined and accepted worldwide, which specifies allowed impacts and energy transfer.
  b. Inherently safe (hardware and software) professional mobile robots, with manipulation, operating in cooperation with trained humans in all professional environments (manufacturing, hospitals, labs, factory floor, warehouse, etc.)
  c. Inherently (hardware and software) personal mobile robots, without manipulation, operating in cooperation with humans in all personal environments (homes, hotels, schools, eldercare, etc.)
  d. Basic personal manipulation systems have first versions of safety standards implemented.

- **10 years:**
  a. Inherently safe (hardware and software) **professional** mobile robots, with manipulation, operating in cooperation with untrained humans in all professional environments.
  b. Inherently safe (hardware and software) **personal** mobile robots, with manipulation, operating in cooperation with humans in all personal environments.

- **15 years:**
  a. Inherently safe mobile robots, with manipulation, operating in cooperation with untrained humans in all public, personal and professional environments.

4. Basic Research and Technologies

4.1 Architecture and Representations

Over the last 20 years, a number of established models for system organization have emerged. Characteristically, however, no agreement or overall framework for system organization has materialized. For autonomous navigation, mobility, and manipulation, there are some established methods, such as 4D/RCS and Hybrid Deliberative Architectures, but once interaction components are added such as Human-Robot Interaction (HRI), there is little agreement on a common model. Over the last few years, the area of cognitive systems has attempted to study this problem, but so far without a unified model. For wider adoption of robot systems, it will be essential to establish architectural frameworks that facilitate systems integration, component modeling, and formal design. Appropriate architectural frameworks may initially or inherently depend on the task, the application domain, the robot, or a variety of other factors. Nevertheless, a deeper understanding of the concepts underlying cognition can be expected from an incremental unification of multiple frameworks into more or less problem- or robot-specific architectures. Any of the aforementioned architectural frameworks will be intricately linked to a set of appropriate representations that capture aspects of the environment and the objects contained in it, the robot's capabilities, domain information, as well as a description of the robot's task.

4.2 Control and Planning

As service robots address real-world problems in dynamic, unstructured, and open environments, novel challenges arise in the areas of robot control algorithms and motion planning. These challenges stem from an increased need for autonomy and flexibility in robot motion and task execution. Adequate algorithms for control and motion planning will have to capture high-level motion strategies that adapt to sensor feedback. Research challenges include the consideration of sensing modalities and uncertainty in planning and control algorithms; the development of representations and motion strategies capable of incorporating feedback signals; motion subject to constraints, arising from kinematics, dynamics, and nonholonomic systems; addressing the characteristics of dynamic environments; developing control and planning algorithms for hybrid systems; and understanding the complexity of these algorithmic problems in control and motion planning.

4.3 Perception

Over the last few decades, tremendous progress has been achieved in perception and sensory processing as is seen, for example, in web-based searches such as Google images and face recognition in security applications. Mapping and localization in natural environments is also possible for engineered environments. Over the last decade, in particular, use of laser scanners and GPS has changed how navigation systems are designed and enabled a new generation of solutions. Over the last 5 years, tremendous progress has be achieved using RGB-D sensor technology and open robot software frameworks. Nonetheless, localization and planning in GPS-denied environments that are quite common remains an important research area. In addition, there has been tremendous progress on image recognition with scaling to large databases. In the future, a large number of robots will rely on sensory feedback for their operation and the application domain will go beyond prior modeled settings. There is therefore a need for reliance on multiple sensors and fusion of sensory information to provide robustness. It is expected that the use of image-based information in particular will play a major role. Vision will play a crucial role in new mapping methods, in facilitating the grasping of novel objects, in the categorization of objects and places beyond instance based recognition, and in the design of flexible user interfaces.

4.4 Robust, High-Fidelity Sensors

Advances in microelectronics and packaging have resulted in a revolution in sensory systems over the last decade. Image sensors have moved beyond broadcast quality to provide mega-pixel images. MEMS technology has enabled a new generation of inertial sensor packages, and RFID has enabled more efficient tracking of packages and people. Sensors have enabled solid progress in domains with good signal quality. As the domains of operation are widened, there will be the need for new types of sensors that allow robust operation. This requires both new methods in robust control, but more importantly, sensors that provide robust data in the presence of significant dynamic variations and a domain with poor data resolution. New methods in silicon manufacturing and MEMS open opportunities for a new generation of sensors that will be a key aspect of future progress in robotics.

4.5 Novel Mechanisms and High-Performance Actuators

There is an intricate interplay between progress in mechanical devices and actuation and the algorithmic complexity required to use them in accordance with their function. Some algorithmic problems can be solved or their solution greatly facilitated by intelligent mechanical design. Advances in mechanism design and high-performance actuators could therefore critically enable groundbreaking innovations in other basic research areas as well as enable several of the capabilities listed in the roadmap. Important research areas include the design and development of mechanisms with compliance and variable compliance, highly dexterous hands, inherently compliant hands, energy-efficient, safe, high-performance actuators, energy-efficient dynamic walkers, and many more. Of particular interest are “intelligent” mechanical designs that can subsume—through their design—a function that otherwise had to be accomplished through explicit control. Examples include self-stabilizing mechanisms or hands with special provisions to achieve form closure without explicit control.

4.6 Learning and Adaptation

Many of the basic research areas described in this section can benefit from advances in and application of learning and adaptation. Service robots occupy complex environment and live in high-dimensional state spaces. Knowledge of the environment and of the robot’s state is inherently uncertain. The robot’s actions most often are stochastic in nature and their result can best be described by a distribution. Many of the phenomena that determine the outcome of an action are difficult or even impossible to model. Techniques from machine learning provide a promising tool to address these aforementioned difficulties. These techniques can be useful for learning models of robots, tasks, or environments; learning deep hierarchies or levels of representations from sensor and motor representations to task abstractions; learning of plans and control policies by imitation and reinforcement learning; integrating learning with control architectures; methods for probabilistic inference from multimodal sensory information (e.g., proprioceptive, tactile, vision); and structured spatiotemporal representations designed for robot learning such as low-dimensional embedding of movements.

4.7 Physical Human-Robot Interaction

Gradually, the safety barriers that have been common in industrial robotics are removed and robots will to a larger degree engage with people for cooperative task execution and for programming by demonstration. As part of this, robots will have direct physical contact with the user. This requires first of all careful consideration of safety aspects. In addition, there is a need to consider how these robots can be designed to provide interaction patterns that are perceived as natural by users. This spans all aspects of interaction from physical motion of the robot to direct physical interaction with a perception of minimum inertia and fluid control. In addition, there is a need here to consider the interaction between design and control to optimize functionality.

4.8 Socially Interactive Robots

As robots engage with people, there is a need to endow the systems with facilities for cooperative interaction with humans. This interaction is needed for tasking of a system, for teaching of new skills and tasks, and for cooperative task execution. The current models for social interaction include gestures, speech/sound, body motion/pose, and physical position. There is a need here to integrate skill and task models with interpretation of human intent to enable interpretation of new and existing activities. In service robotics, there is a broad need for social interaction from encounters with novice users to cooperative tasking with an expert operator. The full span of capabilities is required to provide engaging and long-term adoption of robotics.

5. Contributors

This report documents the result of a brainstorming session that took place September 6-7, 2012 in Seattle, Washington. The report is part of the Robotics Virtual Organization (Robotics-VO) sponsored by the National Science Foundation (NSF). The present report has been authored by the workshop organizers and does not reflect the opinion of NSF. The responsibility of the report lies entirely with the authors.

Dieter Fox, University of Washington; Pieter Abbeel, University of California-Berkeley; Gaurav Sukhatme, University of Southern California; and Henrik I. Christensen, Georgia Institute of Technology; organized the workshop on service robotics. The following people attended the workshop from academia and industry:

Pieter Abbeel
University of California–Berkeley

Parag Batavia
Neya Systems

Trevor Blackwell
Anybots

Stefano Carpin
University of California–Merced

Henrik I. Christensen
Georgia Institute of Technology

Dieter Fox
Johns Hopkins University

Odest Jenkins
Brown University

Charlie Kemp
Georgia Institute of Technology

Kurt Konolige
Industrial Perception

Jana Kosecka
George Mason University

James Kuffner
Google

Rush LaSelle
Adept Technology

Stergios Roumeliotis
Minnesota University

Radu Rusu
OpenPerception

Tom Ryden
VGo

Ashutosh Saxena
Cornell University

Joshua Smith
University of Washington

Ira Snyder
Microsoft Research

Siddhartha Srinivasa
Carnegie Mellon University

Gaurav Sukhatme
University of Southern California

Larry Sweet
Symbotic

Bill Thomasmeyer
Robotics Technology Consortium

Tandy Trower
Hoaloha Robotics

Roadmap for Robot Applications in Space

1. Strategic Importance & Impact of Space Robotics

1.1 Overview

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 accomplish these missions.

> ***"Much of what we know about the Solar System (and beyond), we owe to robotic probes, orbiters, landers, and rovers."***

Looking forward, robotics, telerobotics, and autonomous systems figure heavily in NASA's strategy and are prominently mentioned in the U.S. 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 experiments 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 UAVs.

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 U.S. companies, is on the order of 700% return for every dollar invested in space exploration.

1.2 Co-Explorer Space Robotics Vignettes

Vignette 1: 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.

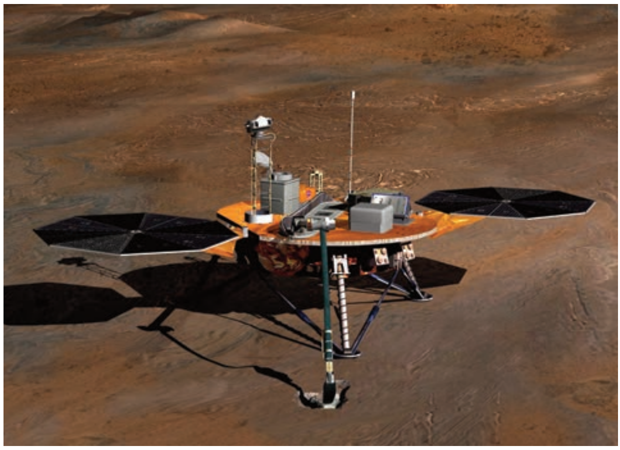
Image courtesy of NASA

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 reattaches to its dangling tether and ascends like a spider to its highline, then out of the hole and on to the next Skylight.

Vignette 2: Robot Tended Waypoint Facility

As humans prepare to venture deeper into space, consideration is being given to developing a "waypoint" facility, which would serve as a gateway to multiple destinations including cis-lunar space, the Moon, Near-Earth Asteroids (NEA), and Mars. This facility would enable assembly and servicing of satellites, telescopes, and deep-space exploration vehicles. This facility could also be used as a platform for astrophysics, heliophysics, and distant Earth observation. One candidate location for such a facility is the Earth-Moon "L2" Lagrange point where the combined gravity of the Earth and Moon allows a spacecraft to be relatively stationary over the lunar farside with little fuel expenditure.

In contrast to the ISS, which is continuously manned, a waypoint facility is expected to only be intermittently occupied. Consequently, there is a significant need for the facility to be robotically tended, in order to maintain and repair systems in the absence of human crew. These robots will perform both IVA and EVA work, remotely operated and supervised from Earth. Telerobotic tending would focus on inspection, monitoring, routine maintenance, and contingency handling of the facility (and possibly attached structures, vehicles, etc.). In particular, experience with the ISS has shown that power (generation, switching, and storage), life support (air, water, and thermal), data networking, and instruments all need to be maintained. To do this, advances will need to be made in non-contact mobile sensing, mobile dexterous manipulation, supervisory control, diagnostics and prognostics, time-delay mitigation, and safeguarded navigation.

Vignette 3: 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 with 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 (e.g., 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.

Image courtesy of NASA

Vignette 4: Crew-centric Surface Telerobots

In planning for future human exploration missions, numerous study teams have proposed having astronauts remotely operate surface robots from an orbiting spacecraft using a low-latency, high-bandwidth communications link. This concept of operations is seen as an effective method for performing surface activities that require real-time human involvement without incurring the risk and cost associated with surface EVA. In addition, this configuration would allow high-performance spacecraft computing to be used for high-level robot autonomy (perception, navigation, etc.), thus simplifying the processing and avionics required for the robot. Crew-centric surface telerobotics is considered an option for several possible missions:

- **Lunar Farside**—Astronauts orbiting the Moon (or station-keeping at the Earth-Moon "L2" Lagrange point) remotely operate a surface robot exploring the lunar farside. Astronauts would take advantage of low-latency (less than 250 ms) and high-availability communications to maximize robot utilization during a short-duration mission.

- **Near-Earth Object (NEO)**—Astronauts approaching, in orbit, or departing a NEO (e.g., asteroid) remotely operate a robot landed on surface. Astronauts would control the robot from the flight vehicle because the NEO environment (high rotation rate, rapidly varying illumination, etc.) rules out remote operations from Earth.

- **Mars Orbit**—Astronauts in aero-stationary orbit around Mars (or perhaps landed on Phobos or Deimos) remotely operate a surface robot exploring Mars. Astronauts would control the robot from the flight vehicle when circumstances (time-critical activities, contingency handling, etc.) do not permit remote operation from Earth.

2. Critical Capabilities for Co-Explorer Space Robotics

2.1 Object Recognition and Pose Estimation

Object recognition requires sensing, often fusing multiple sensing modalities, with a perception function that can associate the sensed object with an object that is understood a priori. Pose estimation seeks to locate an object relative to a sensor coordinate frame, computing the six axis pose using sensing data. Pose estimation is often preceded by object recognition, or presumes an object so that its pose can be estimated and tracked. This technology is important for object manipulation and in mobility for object following and avoidance. The ability to identify and recognize humans and track their motion and gestures is of special interest; in manipulation for safely working with human teammates; and in mobility for avoiding collisions with pedestrians.

Sensing approaches to date have combined machine vision, stereo vision, LIDAR, structured light, and RADAR. Perception approaches often start with CAD models or models created by a scan with the same sensors used to subsequently identify the object. Major challenges include the ability to work with a large "library" of known objects (>100), identifying objects that are partially occluded, sensing in poor (high, low and sharply contrasting) lighting, estimating the pose of quickly tumbling objects, and working with objects at near and far range. A motivating use case is an IVA free-flyer on-board the ISS. Object recognition and pose estimation would enable docking to a recharging station, navigating the interior of the ISS, and interacting with crew (e.g., "roving eye" automatically supporting crew activity).

> ***"The field of mobile robotics has matured with the advance of safe, fast and deterministic motion control."***

2.2 Fusing Vision, Tactile, and Force Control for Manipulation

The field of mobile robotics has matured with the advance of safe, fast and deterministic motion control. This success has come from fusing many sensors to avoid contacting hazards. Manipulation requires forming contact, so the breadth of sensing will require proximity, then tactile, and ultimately force sensing to reach, grasp and use objects like tools. Furthermore, new approaches are required for the design and control of structurally compliant manipulators that can absorb the impact forces of contact events and tool use. Vision requires sensors that are not blocked when limbs reach for objects, but that can be pointed and transported for mobile manipulation applications. Major challenges include calibration of highly dissimilar sensors, dissimilar resolution, noise, and first principles of physics in the development of new sensors and compliant manipulators.

2.3 Achieving Humanlike 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.).

2.4 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).

2.5 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.

2.6 Exceeding Humanlike 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 first principles of physics in the development of actuation and sensing. Other challenges include two-point discrimination, contact localization, extrinsic and intrinsic actuation, back-drivability versus compliance, speed/strength/power, hand/glove coverings that do not attenuate sensors/motion but are rugged when handling rough and sharp objects.

2.7 Full Immersion, Telepresence with Haptic and Multimodal Sensor Feedback

Telepresence is the condition of a human feeling he or she is 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.

2.8 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 these 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.

2.9 Verification and Validation (V&V) of Autonomous Systems

Large software projects have such complex requirements that exhaustive and manual exploration of all possible cases is not feasible. While software validation and verification techniques have been successfully applied to many unmanned spacecrafts (Curiosity, LADEE, etc.), human rated autonomous systems are particularly challenging. New verification techniques are needed to confirm that autonomous systems meet requirements, while new validation techniques are needed to demonstrate that autonomous behavior satisfies its intended use in the intended environment.

2.10 Supervised Autonomy of Force/Contact Tasks Across Time Delay

Tasks have time constants that vary greatly, with the shortest time constants involving motions that form contacts with the environment and force controlled actions. These tasks require high-speed local control loops. As time delays approach these task time constants, the ability to teleoperate the machine degrades. Supervision is the management of a robot with autonomous skills, working along a sequence of tasks. Challenges include run time simulation to predict future states (especially for compliant manipulators required for contact tasks), visualization approaches to overlay predicted, committed and commanded states, and the ability to work ahead of real-time.

2.11 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.

2.12 Safe Mobile Manipulation for Working With and Near Humans

Merging manipulative capabilities with general mobility is sought to allow robots to go to the work site, rather than require the work to be delivered to the robot. Both manipulator arms and mobility drives each pose hazards to people. Combined, they present many risks. Challenges include tracking humans in the workspace, responding deterministically to inadvertent contact, compliance, and providing redundant sensor and software systems.

2.13 Envisioned Five-, Ten-, and Fifteen-Year Goals and Milestones

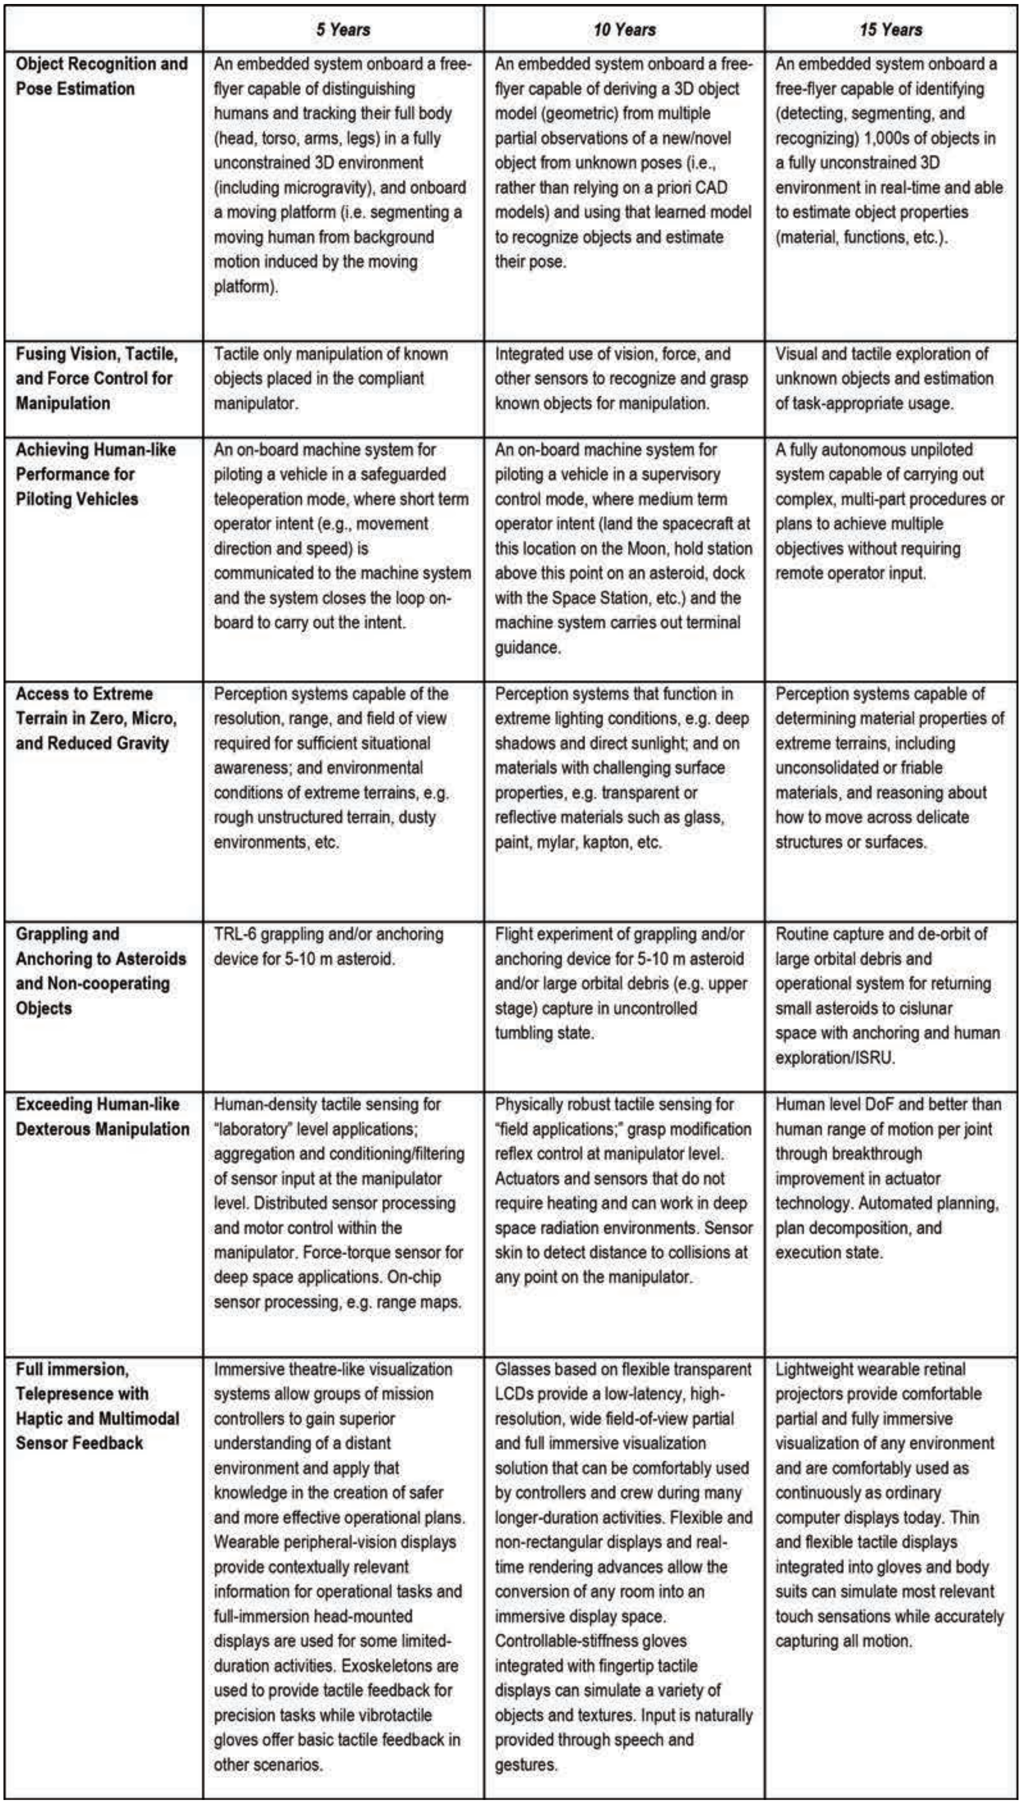

	5 Years	*10 Years*	*15 Years*
Object Recognition and Pose Estimation	An embedded system onboard a free-flyer capable of distinguishing humans and tracking their full body (head, torso, arms, legs) in a fully unconstrained 3D environment (including microgravity), and onboard a moving platform (i.e. segmenting a moving human from background motion induced by the moving platform).	An embedded system onboard a free-flyer capable of deriving a 3D object model (geometric) from multiple partial observations of a new/novel object from unknown poses (i.e., rather than relying on a priori CAD models) and using that learned model to recognize objects and estimate their pose.	An embedded system onboard a free-flyer capable of identifying (detecting, segmenting, and recognizing) 1,000s of objects in a fully unconstrained 3D environment in real-time and able to estimate object properties (material, functions, etc.).
Fusing Vision, Tactile, and Force Control for Manipulation	Tactile only manipulation of known objects placed in the compliant manipulator.	Integrated use of vision, force, and other sensors to recognize and grasp known objects for manipulation.	Visual and tactile exploration of unknown objects and estimation of task-appropriate usage.
Achieving Human-like Performance for Piloting Vehicles	An on-board machine system for piloting a vehicle in a safeguarded teleoperation mode, where short term operator intent (e.g., movement direction and speed) is communicated to the machine system and the system closes the loop on-board to carry out the intent.	An on-board machine system for piloting a vehicle in a supervisory control mode, where medium term operator intent (land the spacecraft at this location on the Moon, hold station above this point on an asteroid, dock with the Space Station, etc.) and the machine system carries out terminal guidance.	A fully autonomous unpiloted system capable of carrying out complex, multi-part procedures or plans to achieve multiple objectives without requiring remote operator input.
Access to Extreme Terrain in Zero, Micro, and Reduced Gravity	Perception systems capable of the resolution, range, and field of view required for sufficient situational awareness; and environmental conditions of extreme terrains, e.g. rough unstructured terrain, dusty environments, etc.	Perception systems that function in extreme lighting conditions, e.g. deep shadows and direct sunlight; and on materials with challenging surface properties, e.g. transparent or reflective materials such as glass, paint, mylar, kapton, etc.	Perception systems capable of determining material properties of extreme terrains, including unconsolidated or friable materials, and reasoning about how to move across delicate structures or surfaces.
Grappling and Anchoring to Asteroids and Non-cooperating Objects	TRL-6 grappling and/or anchoring device for 5-10 m asteroid.	Flight experiment of grappling and/or anchoring device for 5-10 m asteroid and/or large orbital debris (e.g. upper stage) capture in uncontrolled tumbling state.	Routine capture and de-orbit of large orbital debris and operational system for returning small asteroids to cislunar space with anchoring and human exploration/ISRU.
Exceeding Human-like Dexterous Manipulation	Human-density tactile sensing for "laboratory" level applications; aggregation and conditioning/filtering of sensor input at the manipulator level. Distributed sensor processing and motor control within the manipulator. Force-torque sensor for deep space applications. On-chip sensor processing, e.g. range maps.	Physically robust tactile sensing for "field applications;" grasp modification reflex control at manipulator level. Actuators and sensors that do not require heating and can work in deep space radiation environments. Sensor skin to detect distance to collisions at any point on the manipulator.	Human level DoF and better than human range of motion per joint through breakthrough improvement in actuator technology. Automated planning, plan decomposition, and execution state.
Full immersion, Telepresence with Haptic and Multimodal Sensor Feedback	Immersive theatre-like visualization systems allow groups of mission controllers to gain superior understanding of a distant environment and apply that knowledge in the creation of safer and more effective operational plans. Wearable peripheral-vision displays provide contextually relevant information for operational tasks and full-immersion head-mounted displays are used for some limited-duration activities. Exoskeletons are used to provide tactile feedback for precision tasks while vibrotactile gloves offer basic tactile feedback in other scenarios.	Glasses based on flexible transparent LCDs provide a low-latency, high-resolution, wide field-of-view partial and full immersive visualization solution that can be comfortably used by controllers and crew during many longer-duration activities. Flexible and non-rectangular displays and real-time rendering advances allow the conversion of any room into an immersive display space. Controllable-stiffness gloves integrated with fingertip tactile displays can simulate a variety of objects and textures. Input is naturally provided through speech and gestures.	Lightweight wearable retinal projectors provide comfortable partial and fully immersive visualization of any environment and are comfortably used as continuously as ordinary computer displays today. Thin and flexible tactile displays integrated into gloves and body suits can simulate most relevant touch sensations while accurately capturing all motion.

	5 Years	***10 Years***	***15 Years***
Understanding and Expressing Intent Between Humans and Robots	Software systems that enable robots to succinctly express their current state (health, perception of environment, task progress, etc.) and history (summarization of activities, problems, etc.) to humans via different modes (graphical displays, messaging, gesturing, etc.) depending on proximity and time delay.	Software systems that enable robots to identify and track planned human activity (steps in a pre-defined process, including contingency handling branches) using active and/or passive sensors.	Software systems that enable robots to identify human activity (pre-defined and novel), predict next steps, and offer task support. This includes recognition of user-specific characteristics and needs.
Verification & Validation of Autonomous Systems	V&V of an on-board machine system for piloting a vehicle in a safeguarded teleoperation mode. Achieve code quality and defect rate required for human rated systems.	V&V of an on-board machine system for piloting a vehicle in a supervisory control mode. Achieve code quality and defect rate required for human rated systems.	V&V of a fully autonomous system capable of carrying out complex, multi-part procedures or plans to achieve multiple objectives without requiring remote operator input. Achieve code quality and defect rate required for human rated systems.
Supervised Autonomy of Force/Contact Tasks Across Time Delay	Telerobots are capable of performing most tasks using intuitive "shop-foreman-style" instructions similar to what would be given on a shop floor to a trainee. Each command would be formed from macros of "mechanical primitive" behaviors that provide for stable and reliable transitions between N-constraints and N-1 or N+1 constraints on 6-DOF motions. For example, a move-to-contact in free space to put a tool into contact with a planar surface is a transition from 6 to 5 unconstrained DOFs. Sliding that tool-tip to drop into a hole in the surface would transition from 5 to 3 unconstrained DOFs.	Expanded dictionary of well-tuned and ultra-reliable shop-foreman-style commands such as might be given to a journeyman allow almost any task to be performed over time delay.	Dialog between operator and telerobot will be akin to that between a Factory Manager and Shop Foreman.
Rendezvous, Proximity Operations and Docking in Extreme Conditions	Autonomous control of spacecraft with proximity down to 3 meters about a small body (option to explore without landing, e.g. Touch-And-Go (TAG)); precision of TAG footprint <5m; repeated TAG capability (as opposed to TAG-sampling only once on a fly-by basis); autonomous hazard detection and avoidance; ground-based surveillance and characterization of the small body prior to close proximity ops.	Autonomous control of spacecraft with proximity down to 2 meters about a small body (option to explore without landing, e.g. TAG); precision of TAG footprint <5m; repeated TAG capability (as opposed to TAG-sampling only once on a fly-by basis); autonomous hazard detection and avoidance; autonomous surveillance and characterization of the small body prior to close proximity ops.	Agile exploration about small bodies, i.e. as opposed to above approach of extensive surveillance, characterization, approach and proximity ops, arrive, approach and explore with fast reaction to terrain and environment (e.g. comet ejecta), in an adroit manner agile closed sensing and control.
Mobile Manipulation that Is Safe for Working with and Near Humans	Structurally compliant mobile manipulator that dynamically alters its stiffness.	Demonstration of a "safe" compliant collision with a human while moving.	Demonstration of a "safe" compliant collision with a human while transporting a heavy payload.

3. Research & Development Areas for Space Robotics

The NASA robotics and autonomous systems technology roadmap uses a combined push and pull approach to identify individual technologies that are enabling or strongly enhancing for planned NASA missions over the next few decades. The individual technologies are broken down into several major areas, each of which is broken down into subareas of research, as described below.

3.1 Sensing & Perception

This research area includes sensors and algorithms needed to convert sensor data into representations suitable for decision making. Traditional spacecraft sensing and perception included position, attitude, and velocity estimation in reference frames centered on solar system bodies, plus sensing spacecraft internal degrees of freedom, such as scan-platform angles. Current and future development will expand this to include position, attitude, and velocity estimation relative to local terrain, plus rich perception of characteristics of local terrain—where "terrain" may include the structure of other spacecraft in the vicinity and dynamic events, such as atmospheric phenomena. Enhanced sensing and perception will broadly impact three areas of capability: autonomous navigation, sampling and manipulation, and interpretation of science data. Perception tends to be very computationally intensive, so progress in this area will be closely linked to progress in high-performance onboard computing. Metrics for measuring progress in sensing and perception technology include resolution, range, accuracy, tolerance of environmental conditions, and power. Subareas for research include:

- Perception
- Relative Position & Velocity Estimation
- Terrain Mapping, Classification, and Characterization
- Natural and Man-Made Object Recognition
- Sensor Fusion for Sampling and Manipulation
- Onboard Science Data Analysis

3.2 Mobility

Mobility is defined as the ability to move between places in the environment, as distinct from intentionally modifying that environment. Examples include moving between places on a planetary surface or in a planetary atmosphere, or to reach a point in the subsurface. Space robotics involves the need to reach sites of scientific interest (e.g., on the sides of cliffs) which requires a unique focus on certain aspects of extreme-terrain surface mobility, free-space mobility and landing/attachment. Space robotics also has environmental constraints such as thermal extremes that require rad-hard computing. While high-speed operations are of only limited use, mission success will often depend on reliable, sustained operations, including the ability to move through the environment long distances without consuming too much of the mission timeline. Mass, and to some degree power, generally needs to have a much greater degree of emphasis in the design process for space missions than others.

In the next few decades, robotic vehicles designed for planetary surfaces will approach or even exceed the performance of the best piloted human vehicles on Earth in traversing extreme terrain and reaching sites of interest, despite severe terrain challenges. The human ability to quickly assess subtle terrain geometric and non-geometric properties (e.g., visually estimating the properties of soft soil) at long-range fast enough to pilot vehicles at speeds near the limits set by physical law is lacking in today's best obstacle detection and hazard avoidance systems. For free-flying vehicles, in a microgravity environment or flying through an atmosphere, we can expect that robotic vehicles will become capable of utilizing essentially all available vehicle performance, in terms of acceleration, turn rate, stopping distance, etc., without being limited by the onboard sensors, computational throughput, or appropriate algorithms in making timely decisions. Coordination of multiple robotic systems is another active area of research. Combinations of heterogeneous systems, such as flying and roving systems is potentially useful for surface missions, pairing long-range sensing on the flyer with higher resolution surface-sensing on the rover. Metrics for measuring progress in mobility technology include range, payload, speed, life, and mass.

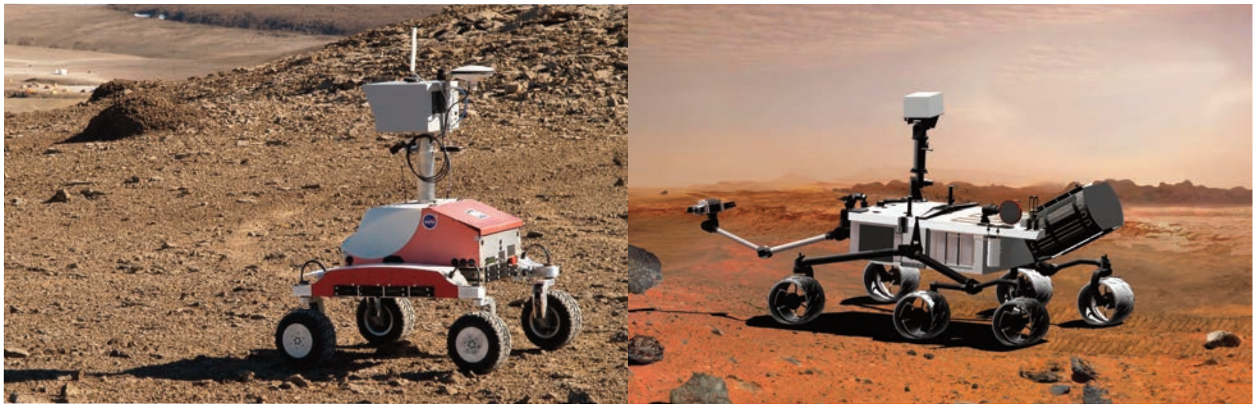

Images courtesy of NASA

Subareas for research include:

- Extreme Terrain Mobility
- Below-surface Mobility
- Above-surface Mobility
- Small Body/Microgravity Mobility

3.3 Manipulation Technology

Manipulation is defined as making an intentional change in the environment. Positioning sensors, handling objects, digging, assembling, grappling, berthing, deploying, sampling, bending, and even positioning the crew on the end of long arms are tasks considered to be forms of manipulation. Arms, cables, fingers, scoops, and combinations of multiple limbs are embodiments of manipulators. Here we look ahead to missions' requirements and chart the evolution of these capabilities that will be needed for space missions. Metrics for measuring progress in manipulation technology include strength, reach, mass, power, resolution, minimum force/position, and number of interfaces handled.

Research subareas include:

- Robot Arms
- Dexterous Manipulators
- Modeling of Contact Dynamics
- Mobile Manipulation
- Collaborative Manipulation
- Robotic Drilling & Sample Processing

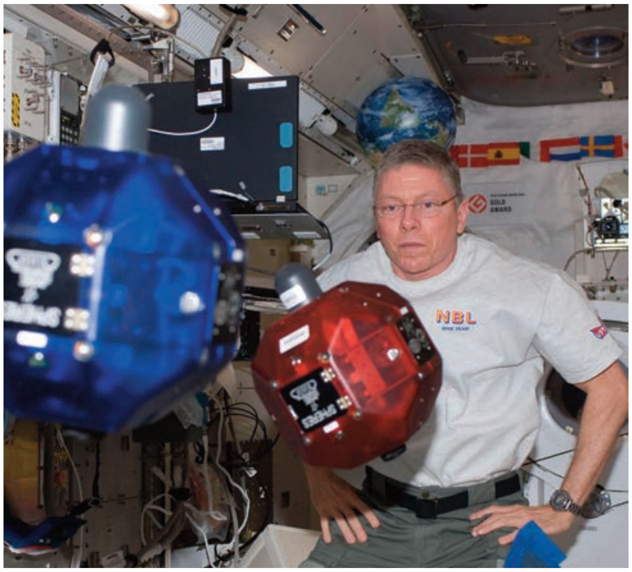
Images courtesy of NASA

3.4 Human-Systems Interaction

The ultimate efficacy of robotic systems depends greatly upon the interfaces that humans use to operate them. As robots and the tasks assigned to them grow more complex, the demands placed on the interfaces used to control them also increase. An excellent human-system interface enables a human to rapidly understand the state of the system under control and effectively direct its actions towards a new desired state. This research area explores advanced technologies for improving situational awareness of a human operator, capturing the operator's intent, and enabling the safe operation of robots in the vicinity of humans and critical systems.

Metrics for measuring progress in human-systems interaction technology include efficiency indices, such as mean-time-to-intervene and the mean-time-between-interventions.

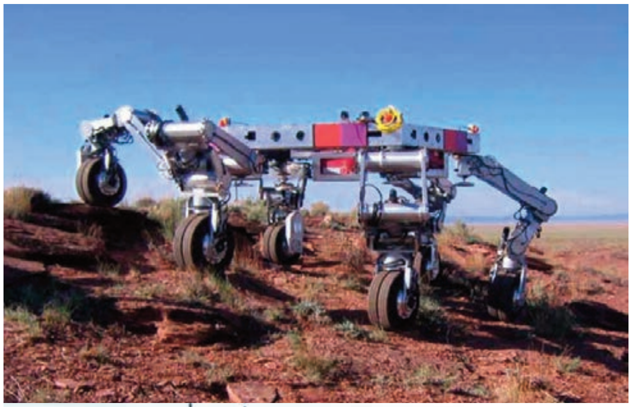

Images courtesy of NASA

Research subareas include:

- Multimodal Human-Systems Interaction
- Supervisory Control
- Robot-to-Suit Interfaces
- Intent Recognition and Reaction
- Distributed Collaboration
- Common Human-Systems Interfaces
- Safety, Trust, & Interfacing of Robotic/Human Proximity Operations

3.5 Autonomy

"Autonomy" is the ability for a system to perform a task, or function, without external support. Autonomous systems are capable of operating independent of external communication, commands, or control. One manifestation of this is a computerized system operating without human support; another manifestation of this is a crewed vehicle operating without ground support. Locus of control describes the entity (or entities) that determines what commands to issue, when to issue them, and then performs the actual command transmission. The locus of control determines whether, and to what degree, a system is autonomous from another system.

Image courtesy of NASA

Autonomy is distinguished from "automation," which is the ability of (computerized) systems to perform a function without input. Automatic control of a system takes place without human intervention or commanding. The function can be performed via ground and/or onboard software interaction. This does not exclude the possibility of operator input, but such input is explicitly not required for an automated function. Automation is a property that can be enabled by software or hardware, and need not be an all-or-nothing proposition, but can be subject to operator discretion.

> ***"Autonomy enables functional improvements with, and without, humans in the loop during missions."***

Autonomy is a critical area for space robotics technology investment and development as it enables functional improvements with, and without, humans in the loop during missions. For space missions, there is a spectrum of autonomy in a system from basic automation (mechanistic execution of action or response to stimuli) to fully autonomous systems that are able to act independently in dynamic and uncertain environments. Autonomy's fundamental benefits are: increasing a system's operations capability, cost savings via increased human labor efficiencies and reduced needs, as well as increased mission assurance or robustness to uncertain environments.

Autonomy can also be applied to aid with data interpretation and decision making for earth-observing satellites and deep-space probes as well as planetary surface rovers. Such onboard autonomous science data analysis improves the capabilities of existing sensors, relieves deep space communication bottlenecks, and enables transformative new operational modes to address novel science issues. Automatic onboard data analysis and understanding can opportunistically apply full-resolution targeting to detected changes, areas of interest or compositional anomalies and improve the science return of long-range (over-the-horizon) traverses.

Subareas for research include the following:

- Integrated Systems Health Management
- Dynamic Planning & Sequencing Tools
- Autonomous Guidance & Control
- Adjustable Autonomy
- Terrain Relative Navigation
- Path & Motion Planning with Uncertainty
- Autonomous Onboard Science Data Analysis

4. Contributors

This chapter is based on the content of NASA's *Robotics, Tele-Robotics, and Autonomous Systems Roadmap* published in April, 2012; input and feedback from a teleworkshop held on November 26, 2012; and contributions from the individuals listed below.

Rob Ambrose
NASA

Sarah Bergbreiter
University of Maryland

Terry Fong
NASA

Jeremy Frank
NASA

Vijay Kumar
University of Pennsylvania

Roland Menassa
General Motors Corporation

Rob Platt
University at Buffalo–SUNY

Luis Sentis
University of Texas–Austin

Mark Schwabacher
NASA

Vytas SunSpriral
SGT, Inc.

William Thomasmeyer
Robotics Technology Consortium

Richard Voyles
National Science Foundation

Ian Walker
Clemson University

David Wettergren
Carnegie Mellon University

Red Whittaker
Carnegie Mellon University

Brian Wilcox
NASA

Roadmap for Robot Applications in Defense

1. Strategic Importance & Impact of Unmanned Systems

Robotic systems developed for military applications are generally referred to by the Department of Defense (DoD) and those in the defense robotics industry as **unmanned systems**. 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. They provide the U.S. military and its allies with an increasingly valuable means for conducting a wide range of operations in the modern combat environment. Military operations such as environmental sensing, precision targeting, and precision strike are all conducive to using unmanned systems as are missions and applications shared with homeland security, such as chemical, biological, radiological, and nuclear (CBRN) detection, counter-improvised explosive device (C-IED) actions, port security, and humanitarian assistance. It should come as no surprise that the military has been developing and successfully fielding these systems in rapidly increasing numbers across the air, ground, and maritime domains and that border security, law enforcement, and fire rescue are eager to adopt and adapt unmanned systems for civilian use as well.

Today's deployed forces have seen firsthand how effective unmanned systems can be in combat operations. Their 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 future combat scenarios. Accordingly, every two years, the Department of Defense publishes an *Unmanned Systems Integrated Roadmap* that describes the Department's vision for the continued integration of unmanned systems into the DoD joint force structure and to identify steps that need to be taken to affordably execute this integration. This section of the national robotics roadmap is largely an extract of the findings contained in the last published version (FY11) of the DoD roadmap and supplemented by input obtained from participants in a defense robotics workshop held on December 4, 2012.

> ***"Unmanned systems offer versatility, functionality, and the capacity to reduce the risk to human life."***

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. The United States faces a complex and uncertain security landscape. The rise of new 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.

Over the next two decades, U.S. 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. Future adversaries will rely less on conventional force-on-force conflicts to thwart U.S. actions and more on tactics that allow them to frustrate U.S. intentions without direct confrontation. Moreover, as technological innovation and global information flows accelerate, non-state actors will continue to gain influence and capabilities that, during the past century, remained largely the purview of states.

> *The next quarter century will challenge U.S. Joint Forces with threats and opportunities ranging from regular and irregular wars in remote lands, to relief and reconstruction in crisis zones, to cooperative engagement in the global commons.... There will continue to be opponents who will try to disrupt the political stability and deny free access to the global commons that is crucial to the world's economy.... In this environment, the presence, reach, and capability of U.S. military forces, working with like-minded partners, will continue to be called upon to protect our national interests.*
>
> *—The Joint Operating Environment 2010: Ready For Today, Preparing For Tomorrow*

With significant budget and force structure reductions in all four services and across the law enforcement agencies, the U.S. no longer has the resources or manpower to physically cover the scope and the geographic separation that this future landscape portends. Unmanned systems can help fill gaps in presence and capability in order to address these future threats in a timely and efficient manner.

1.2 The Role of Unmanned Systems

Unmanned systems enhance combat capability in most of the joint operational functional areas including engagement, sustainment, mobility, and survivability/force protection. Unmanned systems help reduce the load on defenders and help mitigate the risks to the forces responsible for these areas by providing early warning and information and increasing stand-off from hazardous areas. The Department of Defense's vision for unmanned systems is the seamless integration of diverse unmanned capabilities that provide flexible capabilities for joint warfighters. 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 envisions unmanned systems seamlessly operating with manned systems as "**Co-Defenders**" able to aid in human decision making and reduce the degree of required human control.

Image courtesy of NASA

UAVs offer intelligence and reconnaissance-gathering abilities, as well as offensive operations.

The DoD understands the effect that innovation and technology in unmanned systems can have on the future of warfare and the ability of the United States to adapt to an ever-changing global environment. Accordingly, 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. The Joint Staff will continue to support unmanned systems when they fulfill joint requirements and are effective and affordable. Unmanned systems must:

- Provide capabilities more efficiently through modularity, commonality and interoperability.
- Be more effective through greater autonomy and greater performance.
- Be more survivable with improved and resilient communications, development for anti-permissive environments, and more security from tampering.
- Be trustworthy and reliable.
- Take the "man" out of unmanned. Unmanned systems must strive to reduce the number of personnel required to operate and maintain the systems.

The DoD is working to advance operational concepts that leverage the synergies between manned and unmanned 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 environment.

DS1-MA Manipulator Arms, developed by RE² have been integrated onto many of the Armadillo Micro Unmanned Ground Vehicles (MUGVs).

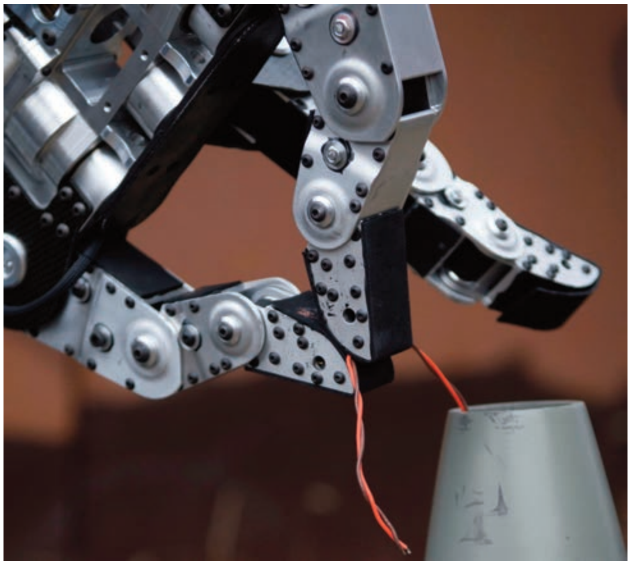

Conformal End-Effector (CE2) technology, developed by RE², Inc., provides an end-effector for the manipulation systems onboard the next generation of unmanned ground vehicles (UGVs).

Images courtesy of RE², Inc.

Robotics technology developed for unmanned systems to satisfy military needs is oftentimes well-suited for "dual use" in commercial 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 the DoD in developing and maturing such technologies will continue to serve the greater good in the development of robotics products and applications for various commercial markets.

Unmanned systems can also contribute to the public safety sector through what is referred to as safety, security, and rescue robotics. Federal, state, and local agencies already incorporate ground robots into bomb squads and aerial vehicles for border security. This is the tip of the iceberg as law enforcement, fire rescue, disaster management, bridge inspection, and port security officials are beginning to adopt these technologies. Safety, security, and rescue (SSR) robotics rely on advances in military robots to provide platforms and software that can be easily modified to their unique requirements, while SSR applications complete a feedback loop by encouraging the development of tactics, procedures, and sensors for DoD functions such as peacekeeping missions and humanitarian response.

1.3 Types of Unmanned Systems

1.3.1 Unmanned Aircraft Systems (UAS)

The air domain has been the most visible to the general public as the DoD has fully embraced UAS capability in tactical, operational, and strategic roles. The air domain has also consumed by far the largest share of the overall DoD investment in unmanned systems. These efforts have fielded a large number of UAS capable of executing a wide range of missions. Originally, UAS missions focused primarily on tactical reconnaissance; however, this scope has been expanded to include a much wider range of capabilities. 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.

In 2009, the DoD flew roughly 19,000 sorties and completed almost 500,000 UAS flight hours just in support of Operation Enduring Freedom and Operation Iraqi Freedom. In May 2010, unmanned systems surpassed one million flight hours and in November 2010, achieved one million combat hours. As the number of fielded systems and applications continue to expand, the number of flight hours is expected to dramatically increase.

1.3.2 Unmanned Ground Vehicles (UGVs)

Unmanned Ground Vehicles (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 operations in Iraq and Afghanistan began, the DoD has acquired and deployed thousands of UGVs. Approximately 8,000 systems of various types have seen action in Operation Enduring Freedom and Operation Iraqi Freedom. As of September 2010, these deployed UGVs have been used in over 125,000 missions, including suspected object identification and route clearance, as well as 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 UGVs have helped with many missions, but resulted in many challenges, not the least of which are configuration, sustainment and maintenance costs. UGVs continue to provide tremendous benefit to the ground commander. In order to meet the challenges anticipated in future conflicts, however, there will need to be improvements in user interfaces, and UGV mobility, reliability, endurance and survivability. These improvements need to keep pace with advances in 360° sensing, recording fidelity, and CBRN and explosive detection.

1.3.3 Unmanned Maritime Systems (UMS)

Over 90% of the information, people, goods, and services that sustain and create opportunities for regional economic prosperity flow across the maritime domain. In response to emerging threats such as piracy, natural resource disputes, drug trafficking, and weapons proliferation taking place in all maritime regions, the DoD continues to expand the range of missions supported by Unmanned Maritime Systems (UMS). UMS are defined as unmanned vehicles that displace water at rest and can be categorized into two subcategories: unmanned underwater vehicles (UUVs) and unmanned surface vehicles (USVs). Like UAS and UGVs, UMS have the potential to save lives, reduce human risk, provide persistent surveillance, and reduce operating costs. The use of UMS is not new. After World War II, USVs were used to conduct minesweeping missions and test the radioactivity of water after each atomic bomb test. More recently, UUVs conducted mine-clearing activities during Operation Iraqi Freedom in 2003. Building on the experience and contribution from the first generation of fielded UMS and advances in autonomy, the shift is underway from UMS merely serving as an extension of the sensor systems of manned ships and submarines into an integrated force to provide full mission capabilities.

> ***"Unmanned systems can help fill gaps in defense presence and capability in order to address future threats in a timely and efficient manner."***

1.4 Vignettes

The following vignettes offer examples of the increased capability and flexibility inherent in unmanned systems as the DoD continues to field unmanned technologies and integrate resulting systems into its existing force structure.

1.4.1 Nuclear Contamination Threat

A weak 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 wash down and 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 that was formed to manage the disaster quickly determines that the reactor fuel core is exposed to the sea and that the reactor was not 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, limiting exposure of personnel to the radioactive contamination.

> ***"Inherently safe robots would enable modes of human-robot interaction that can increase acceptance of robotic technology in everyday life."***

The UGVs continue monitoring contamination levels and collecting samples, but now also start conducting decontamination of the oil terminal control station and the local power and water facilities. Highly contaminated soil is placed into steel drums, and larger UGVs are used to dig pits and bury contaminated building and pipeline materials. Advanced sensor technology and control logic allows the UGVs to operate around the clock with human operators serving solely in a monitoring function. 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.

1.4.2 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 onboard commander of 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.

Image courtesy of U.S. Army via Flickr

Meanwhile, the JMOC analysts recognize the pipeline welder in the UUV data snapshot as one recently stolen and acquired by rebel antigovernment forces. The JMOC then 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 could 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 then 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.

1.4.3 Homeland Critical Infrastructure Protection and Inspection

The Port Authority of New Jersey and New York and the Port of Miami receive notice that a terrorist event may occur within the next two weeks along either the Manhattan or 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 in the channel 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.

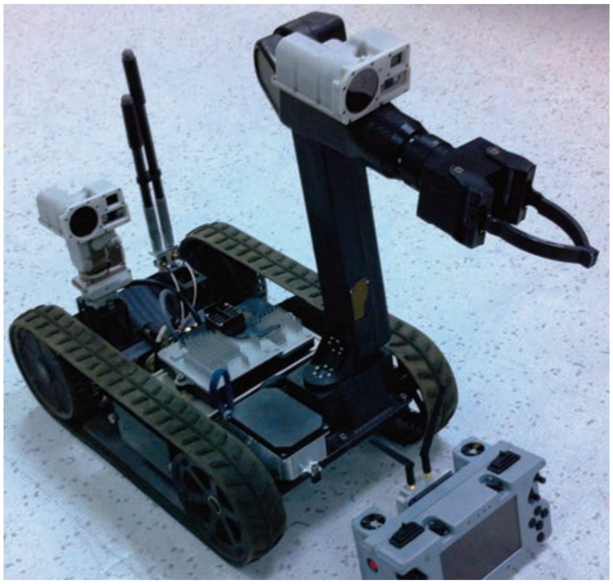

Image courtesy of Naval EOD Technology Division

AEODRS Increment 1 preproduction representative model, one of the first examples of a DoD open architecture unmanned system.

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. A swarm of brightly colored small UAS are deployed 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 USVs 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 USVs continue to contain and monitor the suspect fishing vessel, and one identifies what could be explosives strapped to the hull, possibly for a suicide mission if captured. The Coast Guard vessels, while staying 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 then 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-equipped USV 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 be returned to their routine inspection tasks.

2. Critical Capabilities

Current technology and future advancements in unmanned systems technology can and will enable single platforms to perform an even greater number 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 Department and across the various portfolios. There are nine Tier One JCAs with each representing 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 it currently and could, in the future, contribute to the missions of the Department. Unmanned systems are key contributors in the Battlespace Awareness, Force Application, Protection, and Logistics JCAs.

2.1 Battlespace Awareness

Battlespace Awareness is a JCA in which unmanned systems in all domains have the ability to significantly contribute well into the future. This capability area is an underpinning function that applies across virtually all other JCAs and one that lends itself to tasks and missions being conducted by UMS in all domains. Applications in this JCA include tasks such as ground or aerial urban reconnaissance, which is performed today by UAS such as Predators, Reapers, and Global Hawks and by UGVs such as PackBots and Talons; as well as tasks such as Expeditionary Runway Evaluation, Nuclear Forensics, and Special Reconnaissance. To achieve these missions, unmanned systems development and fielding must include the technology and processes required to:

- Translate vast quantities of sensor data into a shared and relevant understanding of the environment.
- Enable greater onboard processing of data that facilitates improved change detection, Aided Target Recognition (AiTR) and Automatic Target Recognition (ATR).
- Enable mission endurance to extend from hours to days to weeks so that unmanned systems can conduct long-endurance persistent reconnaissance and surveillance in all domains.
- Provide the systems with their own organic perception from onboard sensors so that they can autonomously contribute to Battle Space Awareness regardless of their intended primary mission.
- Enable greater cognitive functions and collaborative awareness individually and among UMS in a cohort.
- Have cognitive functions that result in actions (e.g., move to physical proximity of the target and drop sensor to further inspect an area of interest).

2.2 Force Application

Force Application is a JCA involving maneuver and engagement and includes elements such as target detection and identification, ballistic or non-ballistic firing solution, selection of the firing platform, and battle damage assessment (BDA) of the results of the engagement. 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. In the air domain, projected mission areas for UAS include air-ground, air-to-air combat, and suppression and defeat of enemy air defense. On the ground, UGVs are projected to conduct missions such as non-lethal crowd control, dismounted offensive and defensive operations, and to some degree, mounted operations such as armed reconnaissance. And, in the maritime domain, UUVs and USVs are projected to be particularly suited for mine laying and mine neutralization missions. Currently, the nature of munitions on UMS are driven by a number of factors, including the capability of the platform to carry the munition. As UMS get larger, so can their payloads.

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, has 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.

2.3 Protection

Protection is a joint capability area that includes force susceptibility, vulnerability, and safety. Protection not only includes those measures a unit takes to harden its positions, systems and personnel, but it includes all 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. This includes early warning associated with flank security of ground air or maritime forces. It includes the wide array of threats—whether the threat is forces or systems or chemical agents. As the future enables greater automation with respect to both navigation and manipulation, unmanned systems will be able to perform tasks such as firefighting, decontamination, forward-operating base security, installation security, obstacle construction and breaching, vehicle and personnel search and inspection, mine clearance and neutralization, sophisticated explosive ordnance disposal, casualty extraction and evacuation, and maritime interdiction. In the Protection JCA teaming within domains and collaboration across domains will likely prevail.

2.4 Logistics

The Logistics JCA is ideally suited for employing unmanned systems in all domains to deploy, distribute, and supply forces. Transportation of supplies is an application particularly suited for unmanned systems in widely varying ground terrain. Maintenance related tasks such as inspection, decontamination, and refueling can be performed by unmanned systems. Munitions and materials 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 rescue can also be performed by unmanned systems. Unmanned systems will perform logistics tasks on home stations as well as on forward-deployed locations.

2.5 Homeland Safety, Security, and Inspection

One example of how robotic co-defenders can be used for the prevention of terrorism and for civilian applications is infrastructure protection and inspection. Unmanned maritime and aerial vehicles can cooperatively 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 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, or to respond to an event.

2.6 Envisioned JCA Unmanned System Goals

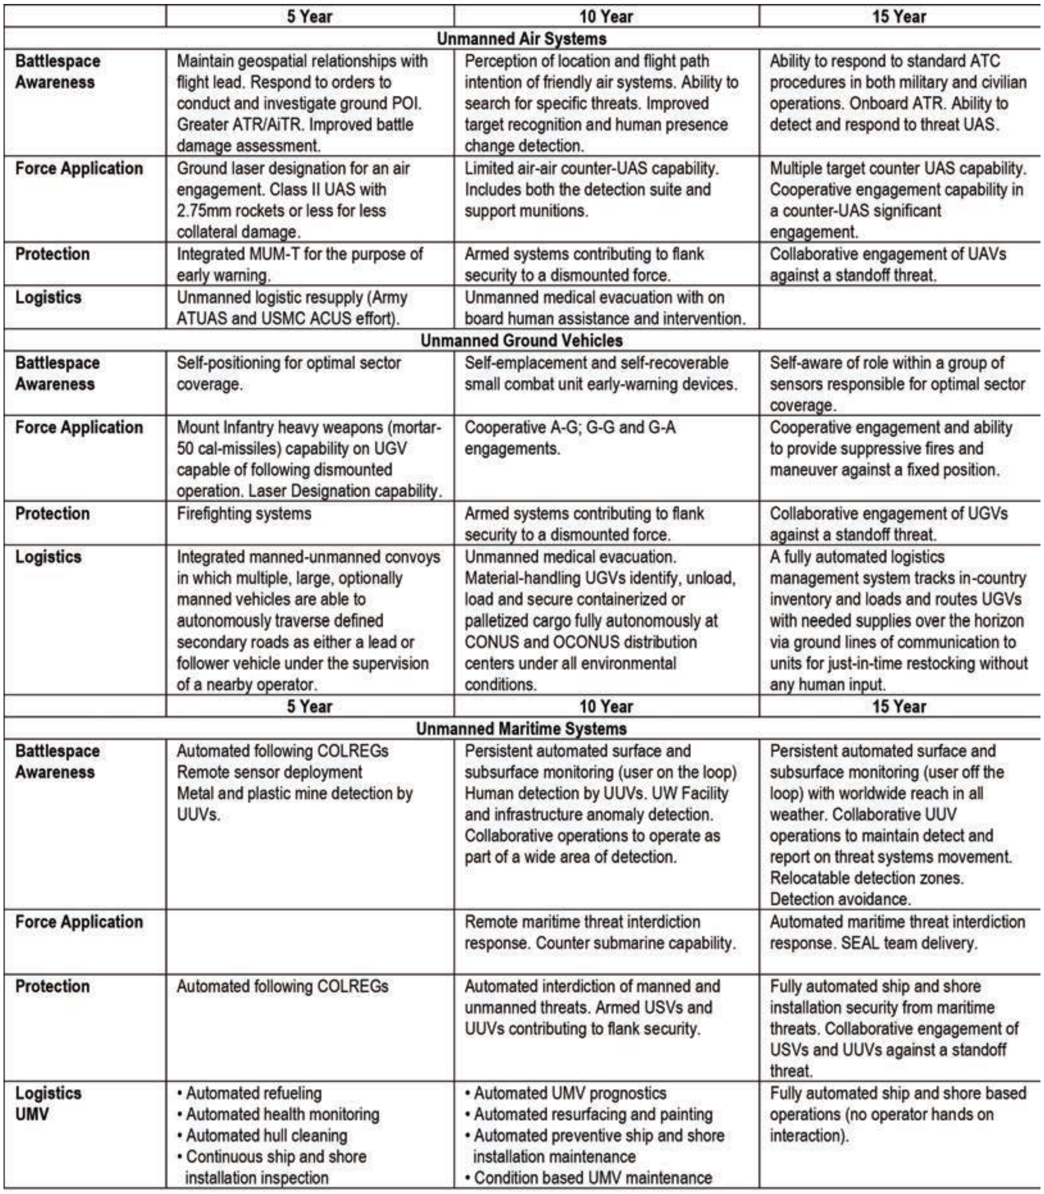

	5 Year	**10 Year**	**15 Year**
	Unmanned Air Systems		
Battlespace Awareness	Maintain geospatial relationships with flight lead. Respond to orders to conduct and investigate ground POI. Greater ATR/AiTR. Improved battle damage assessment.	Perception of location and flight path intention of friendly air systems. Ability to search for specific threats. Improved target recognition and human presence change detection.	Ability to respond to standard ATC procedures in both military and civilian operations. Onboard ATR. Ability to detect and respond to threat UAS.
Force Application	Ground laser designation for an air engagement. Class II UAS with 2.75mm rockets or less for less collateral damage.	Limited air-air counter-UAS capability. Includes both the detection suite and support munitions.	Multiple target counter UAS capability. Cooperative engagement capability in a counter-UAS significant engagement.
Protection	Integrated MUM-T for the purpose of early warning.	Armed systems contributing to flank security to a dismounted force.	Collaborative engagement of UAVs against a standoff threat.
Logistics	Unmanned logistic resupply (Army ATUAS and USMC ACUS effort).	Unmanned medical evacuation with on board human assistance and intervention.	
	Unmanned Ground Vehicles		
Battlespace Awareness	Self-positioning for optimal sector coverage.	Self-emplacement and self-recoverable small combat unit early-warning devices.	Self-aware of role within a group of sensors responsible for optimal sector coverage.
Force Application	Mount Infantry heavy weapons (mortar-50 cal-missiles) capability on UGV capable of following dismounted operation. Laser Designation capability.	Cooperative A-G; G-G and G-A engagements.	Cooperative engagement and ability to provide suppressive fires and maneuver against a fixed position.
Protection	Firefighting systems	Armed systems contributing to flank security to a dismounted force.	Collaborative engagement of UGVs against a standoff threat.
Logistics	Integrated manned-unmanned convoys in which multiple, large, optionally manned vehicles are able to autonomously traverse defined secondary roads as either a lead or follower vehicle under the supervision of a nearby operator.	Unmanned medical evacuation. Material-handling UGVs identify, unload, load and secure containerized or palletized cargo fully autonomously at CONUS and OCONUS distribution centers under all environmental conditions.	A fully automated logistics management system tracks in-country inventory and loads and routes UGVs with needed supplies over the horizon via ground lines of communication to units for just-in-time restocking without any human input.
	5 Year	**10 Year**	**15 Year**
	Unmanned Maritime Systems		
Battlespace Awareness	Automated following COLREGs Remote sensor deployment Metal and plastic mine detection by UUVs.	Persistent automated surface and subsurface monitoring (user on the loop) Human detection by UUVs. UW Facility and infrastructure anomaly detection. Collaborative operations to operate as part of a wide area of detection.	Persistent automated surface and subsurface monitoring (user off the loop) with worldwide reach in all weather. Collaborative UUV operations to maintain detect and report on threat systems movement. Relocatable detection zones. Detection avoidance.
Force Application		Remote maritime threat interdiction response. Counter submarine capability.	Automated maritime threat interdiction response. SEAL team delivery.
Protection	Automated following COLREGs	Automated interdiction of manned and unmanned threats. Armed USVs and UUVs contributing to flank security.	Fully automated ship and shore installation security from maritime threats. Collaborative engagement of USVs and UUVs against a standoff threat.
Logistics UMV	• Automated refueling • Automated health monitoring • Automated hull cleaning • Continuous ship and shore installation inspection	• Automated UMV prognostics • Automated resurfacing and painting • Automated preventive ship and shore installation maintenance • Condition based UMV maintenance	Fully automated ship and shore based operations (no operator hands on interaction).

3. Technological Challenges

Over the last decade, the DoD has achieved great successes from the use of unmanned systems for both peacetime and wartime operations. These successes have led to a significant increase in the number of unmanned systems planned and procured, and the future looks to see an exponential increase in the quantity and diversity of unmanned systems applications. As the DoD steers a path toward the vision described herein, certain technological challenges must be overcome in order to realize the full potential offered by unmanned systems. Many of these technological challenges, in areas such as high-dexterity manipulation, high-fidelity sensors, 3D navigation, and intuitive human-robot interaction, are "**dual-use**," meaning they can be used to satisfy both military and commercial goals and objectives. The challenges with respect to these enabling, dual-use robotics technologies are well articulated in other sections of this document. Rather than repeat them here, the following sub-sections instead summarize five major technological challenges: Interoperability, Autonomy, Communications, Power/Propulsion, and Manned-Unmanned Teaming, all of which are of unique and/or particular importance to the field of defense robotics.

3.1 Interoperability

3.1.1 Need for Increased Interoperability

Traditionally, unmanned systems have been procured as vertically integrated, vendor-proprietary solutions, consisting of the vehicle system, control station, communications channels, and encryption technologies, designed to accomplish a specific mission or capability. These single-system variants have typically been "closed" systems utilizing proprietary interfaces with development of the entire system conducted in parallel with close interdependencies between components and procured as a whole through the platform prime contractor. Although optimal for a single system and for addressing urgent operational needs on short notice, this approach has unfortunately produced a collection of discrete, disjointed solutions with significant functional overlap and no means for minimizing life cycle costs through the utilization of common components for each system.

To maximize the potential of unmanned systems, warfighters must be able to seamlessly command, control, communicate with, exploit and share sensor information from unmanned systems across multiple domains. Properly implemented, interoperability can serve as a force multiplier, improve joint warfighting capabilities, decrease integration timelines, simplify logistics, and reduce total ownership costs (TOC). System interoperability is critical in achieving these objectives and requires the implementation of mandated standards and the adoption of the open systems architecture concept. The DoD's goal is to move to fewer Service/Agency-unique, stand-alone capabilities and toward substantially improved interoperability standards that lead to an improved collaborative operational environment and greater interoperability. The DoD's unmanned systems will need to demonstrate interoperability in a number of areas:

- Among similar components of the same or different systems—the plug-and-play use of different sensors on an unmanned vehicle.

- Among different systems of the same modality—an open common ground control station (GCS) architecture for multiple, heterogeneous unmanned vehicles.

- Among systems of different modalities—the ability of air, ground, and maritime vehicles to work cooperatively.

- Among systems operated by different Military Departments under various CONOPS and TTPs, i.e., in joint operations—joint service systems working in concert to execute a common task or mission.

- Among systems operated and employed by coalition and allied militaries under the governance of various concepts of employment, TTPs (e.g., in multinational combined operations or NATO STANAGs)—the ability of coalition and allied systems to work in concert to execute a common task or mission based on predefined roles and responsibilities.

- Among military systems and systems operated by other entities in a common environment—the ability of military UAS to share the NAS and international airspace with commercial airliners and general aviation aircraft.

- Among systems operated by non-DoD organizations, allies, and coalition partners (i.e. in combined operations)—the ability of assets from organizations such as Customs and Border Protection (CBP) and Department of Homeland Security (DHS) to coordinate, interoperate, and exchange information with DoD assets of the same modality and same model.

3.1.2 Open Architecture

Open architectures (OAs) facilitate interoperability between systems by providing a framework for developing joint interoperable systems that adapt and exploit open-system design principles and architectures in order to:

- Provide more opportunities for competition and innovation

- Field affordable, interoperable systems

- Minimize total ownership cost

- Yield systems that are more rapidly developed and more easily upgraded

- Achieve component software reuse

- Facilitate development test and operational test evaluations

OAs seek to establish commonality in standards, services, transport mechanisms, and applications, as well as facilitate development by utilizing a common set of interfaces and services, associated data models, robust standard data busses, and methods for sharing information. OAs involve the use of COTS components with published, standard interfaces, where feasible, at all levels of system design. This approach avoids proprietary, stovepipe solutions that are vendor-specific and enables innovation to be better captured and integrated into systems design. The OA approach allows for expanded market opportunities, simplified testing and integration, and enhanced reusability throughout the program life cycle. The OA process encourages innovation, allows information sharing among competitors, and rewards Government and industry for this collaboration. It allows programs to include small businesses in systems acquisition activities as a valuable, affordable, and innovative source of technologies and capabilities. The result is a better product.

At a minimum, a common set of interfaces and messaging standards is required for interoperability with respect to exchanging information. Without a common semantic understanding of what data represents, however, there is significant opportunity for lack of interoperability, even if messages are correctly parsed and interfaces are followed. Therefore, a key aspect is the recognition that data modeling is a separate, core aspect for defining interoperable systems. This aspect includes specifying definitions, taxonomies, and other semantic information to ensure there is a common understanding about what information a specific data item imparts. Service-oriented architectures (SOAs) provide a framework for facilitating the design of software in a standardized way to produce interchangeable and interoperable software components called *services*. SOAs increase functionality by incorporating new services, which are developed separately but integrated within the system's common framework as a new capability. Their interfaces are independent of application behavior and business logic, and this independence makes the interfaces agile in supporting application changes and enables operations across heterogeneous software and hardware environments.

In recognition of the rapidly changing technology, unmanned systems architectures would also benefit strongly from being defined at a platform-independent model (PIM) level, which is devoid of technology dependence. The PIM level allows for definition of domains, software components, interfaces, interaction patterns, and data elements without flattening them to a specific set of computing, communications, and middleware technologies. Aside from enabling technology-independent design, this approach, as formalized in model-driven engineering principles, fosters interoperability.

3.2 Autonomy

Today's unmanned systems require significant human interaction and intervention to operate. As these systems continue to demonstrate their military utility, produce greater quantities of data, and are fielded in greater numbers, the demand for manpower will continue to grow and create an increased burden on the military. Moreover, this increasing manpower requirement is occurring at a time when constrained budgets are limiting growth in manpower authorizations. With limited manpower resources to draw upon, the mility is seeking ways to improve the efficiency of unmanned systems operations. The appropriate application of autonomy is a key element in reducing this burden.

3.2.1 Need for Increased Autonomy

Autonomy, as defined in the Defense Science Board (DSB) Task Force Report on the *Role of Autonomy in DoD Systems* published in July 2012, is a capability or set of capabilities that enables a particular action of a system to be automatic or, within programmed boundaries, self-governing, under the supervision of a human operator. The DSB report makes the case that "instead of viewing autonomy as an intrinsic property of an unmanned system in isolation, the design and operation of autonomous systems needs to be considered in terms of human-system collaboration." As such, autonomy enables unmanned systems to function as truer co-defenders, requiring guidance and direction as opposed to the constant attention required to operate today's unmanned systems. Autonomy reduces the human workload required to operate unmanned systems, enables the optimization of the human role in the system, and allows human decision making to focus on points where it is most needed. These benefits can increase operational capability and result in manpower efficiencies, cost savings, and greater speed in decision making. Greater degrees of autonomy, for example, better enable a single operator to control more than one unmanned system while greatly reducing the need for high bandwidth communication. Autonomy can enable operations beyond the reach of external control or where such control is extremely limited (such as in caves, under water, or in areas with enemy jamming or degraded communications). Autonomy can also help extend vehicle endurance by intelligently responding to the surrounding environmental conditions (e.g. exploit/avoid currents) and appropriately managing onboard sensors and processing (e.g. turn off sensors when not needed). Similar efficiencies can be gained by automating the tasking, processing, exploitation, and distribution of data collected by unmanned systems.

The Automated Mine Detection System (AMDS), developed by Carnegie Robotics, LLC, searches for landmines and IEDs.

Unclassified. Distribution A. Approved for Public Release. U.S. Army RDECOM CERDEC NVESD. Release Date: March 6, 2012

The flexibility required of autonomous systems in dynamic, unstructured environments must not adversely affect safety, reliability, or the ability to collaborate with the operator or other autonomous systems, nor overly complicate the predictability needed for U.S. commanders to "trust" the autonomy. "Trust" will be established through robust operational Test & Evaluation (T&E) along with safeties and safeguards to ensure appropriate behavior. Complex autonomous systems must be subject to rigorous "red team" analysis in order to evaluate the full range of behaviors that might emerge in environments that simulate real-world conditions. Safeties and safeguards are also required to mitigate the consequences of failures. Robust safeties and control measures will be required for commanders to trust that autonomous systems will not behave in a manner other than what is intended on the battlefield. For unmanned systems to fully realize their potential, they must be able to achieve a highly autonomous state of behavior and be able to interact with their surroundings. This advancement will require an ability to understand and adapt to their environment, and an ability to collaborate with other autonomous systems, along with the development of new verification and validation (V&V) techniques to prove the new technology does what it should.

3.2.2 Autonomy Technology Challenges

Autonomous capabilities have been enabled by advances in computer science (digital and analog), artificial intelligence, cognitive and behavioral sciences, decision aids, machine training and learning, and communication technologies. In order to achieve operational acceptance and trust of these autonomous capabilities in the highly dynamic unmanned system environment, improvement is essential in advanced algorithms that provide robust decision-making capabilities (such as machine reasoning and intelligence); automated integration of highly disparate information; and the computational construct to handle data sets with imprecision, incompleteness, contradiction, and uncertainty. Significant advances have been made in autonomy, but many challenges still exist. For relatively static environments and undemanding missions and objectives, systems governed by rule-based autonomous programs can be highly effective. However, most DoD environments and mission tasks dictate that unmanned systems operate in complex and uncertain environments, as well as possess the ability to interact and collaborate with human operators and human teammates, including being able to respond to external stimuli in a way that facilitates the manned-unmanned team's ability to survive. Additionally, autonomous systems need the capability to interact and work together with other autonomous systems to adapt to and learn from changes in the environment and missions, and to do so safely and reliably.

3.2.2.1 Autonomy that Dynamically Adjusts to Meet Mission Requirements

While reduced reliance on human operators and analysts is the goal of autonomy, one of the major challenges is how to maintain and facilitate unmanned systems' interactions with the operator and other human agents. An alternative statement of the goal of autonomy is to allow the human operator to "work the mission" rather than "work the system." In other words, autonomy must be developed to support natural modes of interaction with the operator. These decision-making systems must be cognitively compatible with humans in order to share information states and to allow the operator and the autonomous system to interact efficiently and effectively. Autonomous capabilities should dynamically adjust based on workload and the perceived intent of the operator. The goal is not about designing a better interface, but rather about designing the entire autonomous system to support the role of the warfighter and ensure trust in the autonomy algorithms and the system itself. Historically, the DoD has used four defined levels of autonomy: human-operated, human-delegated, human-supervised, and fully autonomous. The 2012 DSB report on autonomy, however, suggested the future use of a broader, three-part autonomous system reference framework architecture to better guide the conception, design, and test and evaluation of autonomous systems.

3.2.2.2 Sensing & Understanding Complex, Dynamic Environments

Autonomous systems must be able to sense and understand the complex and uncertain environments in which they operate. They must be able to create a model of the surrounding world by conducting multi-sensor data fusion (MDF) and converting these data into meaningful information that supports a variety of decision-making processes. The perception system must be able to perceive and infer the state of the environment from limited information and be able to assess the intent of other agents in the environment. This understanding is needed to provide future autonomous systems with the flexibility and adaptability for planning and executing missions in a complex, dynamic world. Recent advancements in computational intelligence (especially neuro-fuzzy systems), neuroscience, and cognition science may lead to the implementation of some of the most critical functionalities of heterogeneous, sensor net-based MDF systems. Needed developments to advance these types of processing capabilities include being able to: reconfigure the weighting of sensor input; dynamically adapt malfunctioning sensors and calibrate misleading data; perform intelligent adaptive data association; and intelligently manage, fully optimize, and dynamically reconfigure fusion clusters to adapt to a changing environment. While robustness in adaptability to environmental change is necessary, the future need is to be able to adapt and learn from the operational environment because every possible contingency cannot be programmed a priori. This adaptation must happen fast enough to provide benefits within the adversary's decision loop, and the autonomy should be constructed so that these lessons can be shared with other autonomous systems that have not yet encountered that situation.

3.2.2.3 Collaboration with Other Autonomous Systems

In addition to understanding the environment, unmanned systems must also possess the ability to collaborate through the sharing of information and de-confliction of tasking. Collaborative autonomy is an extension of autonomy that enables a team of unmanned systems to coordinate their activities to achieve common goals with minimal human oversight. This trend in autonomy will continue to reduce the human role in the system. Autonomous unmanned teams may be capable of faster, more synchronized fire and maneuver and a more complete and efficient manner of conducting area R&S, including the ability to respond more quickly to investigate issues within the area of operations. This trend will lead to a shift toward strategic decision making for a team of vehicles and away from direct control of any single vehicle. The ability to collaborate is one of the keys to reducing force structure requirements. The collaborative autonomy that is developed must be scalable to both larger numbers of heterogeneous systems as well as increased mission and environment complexity. Collaborative autonomy must be able to adapt to the air, ground, and maritime traffic environment and to changes in team members, operators, and the operational environment.

3.2.2.4 Autonomy & Data Tasking, Processing, Exploitation, and Distribution (TPED) Processes

Current TPED processes are manpower intensive and therefore offer huge opportunities for reducing the degree of human involvement. In today's combat environment, most full-motion video (FMV) and still imagery is monitored and used in real time, but then stored without being fully analyzed to exploit all information about the enemy. This challenge has been exacerbated by the large numbers of ISR-capable, long-endurance unmanned systems being fielded. These systems are collecting great quantities of information and overwhelming current TPED processes. Near-term steps might include implementation of change detection and automatic target recognition or aided target recognition software to enable automated cueing that identifies and calls attention to potential threats. Applications of face recognition software could enable the ability to identify individuals of interest. Increased automation in communications intelligence sensors has the potential to identify key words and even specific voices to rapidly alert operators to targets of interest. Ultimately, automated cross-cueing of different sensor types in a networked environment could enable greater autonomy in tasking systems and their sensors to identify and track threats more rapidly. Increased processing power and information storage capacities also have the potential to change how unmanned systems operate. For example, many current UAS transmit ISR data that is processed and exploited in ground stations. If more processing and exploitation processes can be accomplished onboard a UAS (like the automatic target recognition or communications intelligence examples discussed above), the system can disseminate actionable intelligence for immediate use and reduce bandwidth requirements. By accomplishing more of the TPED process onboard the unmanned system, the link bandwidth can then be focused on transmitting only what's needed, and the overall bandwidth requirements can be reduced.

3.2.2.5 New Approaches to Verification & Validation

Verification and validation (V&V) is the process of checking that a product, service, or system meets specifications and that it fulfills its intended purpose. To ensure the safety and reliability of autonomous systems and to fully realize the benefits of these systems, new approaches to software V&V are required to complement existing hardware testing and evaluation. Today's V&V processes utilize existing industry standards for software certification for manned systems and ignore the unique non-deterministic aspect of autonomy. Existing processes will be severely stressed due to the growth in the amount and complexity of software to be evaluated; for example, current regression testing methods are not adequate to ensure correct operation of autonomous capabilities. Without new V&V processes, such as the use of trust audit trails for autonomy, the result will be either extreme cost growth or limitations on fielded capabilities. Efforts leading to advancements in computational intelligence as well as the appropriate V&V processes are essential. Enhanced V&V tech-

> ***"To ensure the safety and reliability of autonomous systems and to fully realize the benefits of these systems, new approaches to software verification and validation are required."***

nologies would provide both near-term cost reduction and enhanced capabilities for current autonomous systems and would enable otherwise cost-prohibitive capabilities in the future. New autonomous system test and analysis capabilities are also required to assess intelligent single-vehicle and group behaviors. These technological enhancements would lead to more effective development, testing, and operations of current and future autonomous systems.

3.3 Communications

Current non-autonomous unmanned systems operations involve a high degree of human interaction with the systems via various means for C2 and transmission of operational data. Protection of these communication links and the information flowing through them is critical to these operations. For some ground and maritime systems, these types of exchanges of information can use a cable for the transmission path, but for highly mobile unmanned operations, the exchange is more likely to use signals sent across the electromagnetic spectrum (EMS) or by other means (e.g., acoustical or optical).

Unmanned Ground Vehicles—Until recently, most unmanned systems utilized several radios: one for data, one for video, and sometimes one for voice. Because of congestion, frequency competition, and regulatory challenges in several theaters, many of these communication systems were redesigned to operate at higher frequencies. However, use of these higher frequencies reduced the operational effectiveness in dense foliage and urban areas.

Unmanned Aircraft Systems—Small, hand-carried and/or hand-launched systems (e.g., the Raven) utilize LOS communications, including the Army-developed digital data link (DDL) system. Large UAS (e.g., the Predator, Reaper, Gray Eagle, and Global Hawk) utilize both LOS, specifically the common data link (CDL) mandated for use in ISR platforms, and BLOS communications, generally using satellite communications.

Unmanned Maritime Systems—There are unique challenges related to UMS that make intermittent communications the norm such that multispectral capabilities are used to meet communications requirements. Primary trade-offs to be considered when communicating with a USV or UUV that supports dynamic tasking, querying, and data dissemination include data rate, processing capability, range, detectability, and negotiating the maritime environment.

3.3.1 Need for Improved Communications

Future unmanned systems' sensors will capture an alarming amount of data. Left unchecked, sending all that data to local or remote sites will tax current technology and available resources, both manpower and funding. As a result, tomorrow's unmanned systems will need to utilize technical strategies which can more efficiently deal with extremely large data sets. How to best deal with that amount of data and the distribution of the needed information within that data to the right warfighters at the right time will be a major challenge.

The DoD's desire is to operate unmanned systems in theater or within the United States and its possessions so that communication constraints do not adversely affect successful mission execution. As the number of fielded systems grows, planners face challenges such as communication link security, radio frequency spectrum availability, de-confliction of frequencies and bandwidth, network infrastructure, and link ranges. The DoD needs communication technologies that overcome these limitations and that are agile, robust, redundant, efficient, and affordable. Specifically, the DoD must significantly improve communication transmission efficiencies; attain better bandwidth efficiencies; increase transmitter and receiver efficiencies; and acquire communications systems that are of less size and weight, require less power, and provide more efficient cooling to operate. Improved communication transmission technologies alone cannot achieve the necessary capacity, however. The DoD must also pursue a fundamental shift to a future state where the collected data is preprocessed to determine the critical data that must be rapidly passed on to the warfighters, with the rest stored for later retrieval as may be needed.

To support the DoD's goals, communication systems need to support multiple frequency bands, limited bandwidth, variable modulation schemes, error correction, data fusion, encryption, and compression. These strategies also need to mandate efficient use of the spectrum, reduce frequency use overhead, allow for data security and ensure improved clarity of the available frequency spectrum. All this support, of course, needs to be done so that no electromagnetic interference (EMI) is caused within those systems or within other nearby spectrum-dependent systems (SDS).

3.3.2 Communications Technology Challenges

There are numerous challenges to meeting these goals. First, operating a higher density of unmanned systems within relatively small areas creates increased local data rate demands. Second, size, weight, power, and cooling are limiting factors on many platforms, for both onboard systems and ground / surface control systems. Third, the fidelity of the communication links must be ensured. Fourth, latency associated with digital systems must be reduced, especially for takeoff and landing of large UAS. These challenges will be exacerbated by an expected decrease in available spectrum available due to an increase in the civil uses of spectrum. The challenges in attaining this goal include developing, procuring, testing, and fielding communication systems that can operate with greater effectiveness, efficiency, and flexibility even in congested and adversarial environments.

> ***"The DoD's desire is to operate unmanned systems so that communication constraints do not adversely affect successful mission execution."***

The ability to update and reconfigure parts of a communication system by software changes has been available for several years. Such communication systems need to conform to a standards-based architecture (e.g., service-orientated architecture) that supports multiple networks to enable rapid and transparent configuration changes without removing the radios from operation. Such multiple-input, multiple-output (MIMO), multi-carrier, and multi-waveform capabilities, along with the software control of these functions, are needed within future subsystem developments. Ultimately, it is desired that these reconfiguration changes be done "automatically" so the systems adapt dynamically in response to sensed changes in the operational environment. The need to support operations in which there are intermittent wireless propagation links has become commonplace. This support has resulted in increased use of advanced error control coding, MIMO configurations, various path diversity techniques, integrated networking, and data diversity—all to provide improved end-to-end quality of service. Future effectiveness of unmanned communication systems is contingent on continued advancements in antennas, transmit/receive systems, underwater communications, spectrum considerations, signal processing, network systems, and optical communications.

3.4 Propulsion & Power

3.4.1 Need for Better Propulsion & Power Technology

The dramatic increase in the development and deployment of unmanned systems across the entire spectrum of air, ground, and maritime missions has led to a concurrent increase in the demand for efficient, powerful, often portable, and logistically supportable solutions for unmanned system propulsion and power plant requirements. As these systems continue to demonstrate their value, operators want them to function longer without refueling and to do more tasks, demands that tax the internal power sources. For the purpose of this section, propulsion and power consist of the prime power to provide thrust and electrical power conversion, management, and distribution necessary for the operation of the electrically driven subsystems required to perform an unmanned vehicle's mission. Regardless of energy source, total vehicle design, from materials used to autonomous response to the physical environment, needs to be considered up front to maximize endurance.

A wide array of propulsion systems is used in unmanned systems, including combustion engines powered by heavy fuel or gasoline, jet engines, electric systems, fuel cells, solar power, and hybrid power systems. These propulsion systems can be divided into three groups according to vehicle size and mission: turbine engines, internal combustion, and electrical. The thresholds are not simple or clean cut, but are highly dependent on mission goals. Some of the parameters taken into consideration to determine the optimum propulsion system include size, weight, airflow, range, efficiency, and speed. Similarly, numerous power systems are in use, including batteries, engine-driven generators, solar power and hybrid systems. The T&E of propulsion and power is critical as we consider a world of declining energy reserves and the strategic initiatives in alternative energy being made by the DoD. Endurance is perhaps one of the most compelling aspects of unmanned systems. While power and propulsion systems are much improved over comparable manned systems, the search continues for even more efficient systems to provide greater endurance, speed and range.

3.4.2 Propulsion

These challenges are currently being addressed for UAS applications under the highly efficient embedded turbine engine (HEETE) and efficient small-scale propulsion (ESSP) products. HEETE will demonstrate engine technologies that enable fuel-efficient, subsonic propulsion that supports future extreme endurance and range requirements with embedded engines incorporating complex inlets and exhausts. HEETE has two challenges: packing a high-bypass engine internally and delivering large amounts of electrical power regardless of throttle or flight condition. ESSP will cover a full spectrum of technologies for propulsion systems for vehicles ranging from 100 to 2500 lbs. These products promise game-changing system capabilities. The S&T challenge to meet the ESSP goals is the simultaneous combination of high power density with high efficiency (low specific fuel consumption) in a design space not typically addressed by either gas turbine or piston engine systems. For smaller platform applications, fuel cells offer an attractive alternative for internal combustion engines as field power generators, ground vehicle and aircraft auxiliary power units (APUs), and primary power units for small UAS. Because these systems do not generate power via combustion processes, they offer significantly lower SFC rates relative to advanced heavy fuel engines or diesel power generators. Solid oxide fuel cell (SOFC) systems represent a compelling power system option due to their high efficiencies, fuel flexibility, and low audible signature.

3.4.3 Power

Power sources are critical enablers for all of the desired unmanned systems capabilities. Improved power sources will have to be compact, lightweight, and reliable; provide enough power for the desired mission; and satisfy a full range of environmental and safety requirements. Design of power sources must be optimized for specific platforms and use profiles. Depending on the platform and mission requirements, applicable technologies may include energy harvesting (e.g., photovoltaic), electrical energy storage devices, fuel cells, and generators. It may be attractive to hybridize two or more of these technologies depending on the expected use profile. To implement these hybrid systems, the development of the proper control schemes must also be conducted. Recently, there has been a great deal of effort invested to improve the power density of power generation systems with very good progress, but work is still needed to improve other power systems' critical metrics. Some of these needed metrics and improvements are life, reliability, efficiency, optimized performance over varying engine speed, wide temperature range, production variability, control strategy, and parameters that capture the fact that unmanned subsystems typically do not have the redundancy of manned systems. Early scrutiny of the vehicle design will lead to improved power management. Power-sharing architectures allow for tailoring the source of power generation to minimize the cost in fuel burn. Some of the key technologies needed to implement a power-sharing architecture are reliable power management control logics, high-power high-speed solid-state power controllers (SSPCs), a modulating generator control unit (GCU), and high-capacity electrical accumulator units (EAUs).

3.4.4 Batteries

Concerning battery chemistries and fuel cells, in the near term (up to 5 years), incremental power and energy performance improvements will continue to be made in the area of rechargeable lithium ion batteries. Lithium ion batteries will see broader military and commercial application, and significant cost reductions will be made as the manufacturing base matures. Near-term availability of small, JP-8 fuel-compatible engines is expected. There is mid-term (5 to 15 years) potential for significant incremental performance advances through the discovery and development of alternative lithium ion chemistries. Mid-term development of fuel cells with moderate power levels will begin to be introduced based on low-weight hydrocarbon fuels (e.g., propane). The technical feasibility of heavy hydrocarbon-fueled (e.g., JP-8) fuel cell systems will be proven at the kilowatt class. In the long term (beyond 15 years), there is the potential for revolutionary improvements through the discovery and development of completely new battery chemistries and designs.

3.5 Manned-Unmanned (MUM) Teaming

MUM teaming refers to the relationship established between manned and unmanned systems executing a common mission as an integrated team. More specifically, MUM teaming is the overarching term used to describe platform interoperability and shared asset control to achieve a common operational mission objective. This term also includes concepts of "loyal wingman" for air combat missions and segments of missions such as MUM air refueling. This capability is especially vital for missions such as target cueing and hand off between manned and unmanned systems, where the operators not only require direct voice communications between the participants, but also a high degree of geospatial fidelity to accurately depict each team member's location with regard to the object being monitored.

Other examples of MUM teaming include:

- Manned helicopter with one or more supporting UAS
- Manned ground systems controlling multiple UAS
- Manned air systems controlling (or supporting) multiple UAS and UGVs
- Manned surface system control varying combination of USVs and UUVs
- Air deployment of multiple USVs and UUVs

U.S. military forces have demonstrated early progress in integrating unmanned systems within the existing manned force structure, but much needs to be done in terms of concepts, organization, training, and integration to achieve the full potential offered by unmanned technology. Practical applications of MUM teaming continue to evolve as confidence in unmanned vehicle reliability and functionality matures. Employment concepts are limited by data links, vehicle control interfaces, and level of autonomy. Improving MUM teaming is both a technology challenge (such as connecting the systems) and a policy challenge (such as establishing the rules of engagement for synchronously operating semi-autonomous unmanned and manned systems).

3.5.1 Unmanned Aircraft Systems

The initial combat operations in Afghanistan and Iraq validated the urgent need to integrate UAS capabilities with manned aircraft, specifically the attack platforms. Commanders recognized that they could dramatically reduce sensor-to-shooter times and improve situational awareness of helicopter pilots, while drastically reducing collateral damage and the potential for fratricide. MUM teaming has been successfully demonstrated in combat operations to provide enduring surveillance of hostile activities in real/near-real time, to provide critical tactical data used to plan and support combat operations, to accurately geo-locate potential targets, to laser designate targets, and to provide battle damage assessment. Under the command of a centralized source, armed UAS, in particular, have the ability to provide force protection through early warning, investigate areas of interest, and engage targets directly or cooperatively (either as a designator, and laser locator, or weapons platform) with other air and ground systems.


Image courtesy of DARPA

A160T Hummingbird Unmanned Helicopter/Rotorcraft.

Some aspects of this has been demonstrated successfully in combat operations with attack helicopter crews. The attack helicopter crew is able to see on their cockpit display UAS sensor outputs that give them overhead views to the target and surrounding area. This capability greatly enhances the attack helicopter crew's ability to identify, classify, and verify target locations to reduce the risk of fratricide. Current MUM teaming applications are limited due to the fact the control interface currently requires a dedicated crew member to fly the UAS while another crew member flies the manned aircraft. However, some automated MUM mission segments are being developed. For example, the Navy and USAF have developed and demonstrated technology for MUM air refueling and have simulated cooperative MUM air combat missions.

3.5.2 Unmanned Ground Vehicles

MUM teaming has steadily increased as technology has improved and users have found new and innovative methods to exploit this enhanced mission capability. Current missions include reconnaissance, surveillance, and target acquisition (RSTA); transport; countermining; explosive ordnance disposal; and the use of armed unmanned tactical wheeled vehicles for checkpoint security inspections. The integration of one-system remote video terminal (OSRVT) technology and distributed UGV control into ground combat vehicles is leading to the adaptation of TTPs because all parties now receive the same picture at the same time, regardless of their location.

These developments have also been the catalyst for the creation of the common robotic controller, a joint project between the Army and USMC to develop a universal, wearable controller to operate a wide variety of unmanned systems, including UGVs, UA, and unattended ground sensors. This effort is currently aimed at smaller platforms, but could be transitioned to include limited control (i.e., payload only) for larger platforms as the technology matures.

> ***"All maritime missions will benefit from reduced timelines and improved accuracy of information."***

3.5.3 Unmanned Maritime Systems

MUM teaming is critical for the maritime environment. This is especially true for the undersea domain where physics prevent man from safely performing tasks to the same fidelity. There are many different aspects of MUM teaming for unmanned systems that have been explored and implemented in various degrees: long-endurance undersea gliders that send data ashore and receive human-initiated mission updates in near-real time; UUVs that enable efficient port security, harbor defense, and mine clearance operations through change detection and autonomous investigation of mine-like objects; and UUVs that extend the footprint of manned hydrographic and bathymetric survey platforms to gather higher volumes of data while enabling people to focus on the tasks that require human oversight. An enduring and integrated net of undersea sensors partnered with USVs or UAVs for communication and controlled from a common command center will revolutionize how undersea missions are conducted by bringing transparency to an otherwise opaque battlespace. All maritime missions will benefit from reduced timelines and improved accuracy of information from which the combat commander can make engagement decisions.

3.5.4 MUM Technology Challenges

Several challenges persist that will continue to affect the amount of time it takes this technology to transition from the invention and adaptation phase to the acceptance phase. Technical challenges range from near-term issues, such as the limited ability to integrate and deconflict various radio frequencies across a secure communications network, to longer-term issues, such as the ability of one person to supervise one or more UAS and UGVs while simultaneously flying his or her primary aircraft. This ability requires a high degree of hardware and software interoperability, scalable autonomy, human system interfaces (HSIs), new collaborative control algorithms, decision aids, and network mission tools. The platforms must do significant levels of onboard processing to not only reduce bandwidth required, but also collaborate with other unmanned vehicles with minimal operator input. Other technical challenges result from the size, weight, and power limitations of the various platforms and the desire for increased performance and capability based on the subsystems they are asked to carry. The ability to communicate from a highly maneuverable aircraft to a highly maneuverable future UAS will require significant advances in autonomy and HSI. This advancement can be compounded if LPI communication is needed for missions such as EA, SEAD, or control of long-dwell, insect-size vehicles collecting information inside buildings.

4. Contributors

This chapter is based on the content of the Department of Defense's *Unmanned Systems Integrated Roadmap, FY2011–2036*, input and feedback from a defense robotics workshop hosted by the Robotics Technology Consortium on December 4, 2012 in Washington, D.C., and contributions from the individuals listed below.

Grant Bagley
Concepts to Capabilities Consulting

Robert Bechtel
Soar Technology

Sara Blackmer
Pratt & Miller Engineering

Henrik I. Christensen
Georgia Institute of Technology

Edward Chupein
United States Air Force

Andrew Culhane
TORC Robotics

Robert Finkelstein
Robotic Technology

Matthew Fordham
Applied Research Associates

Gordon Fraser
ICS NETT

Satyandra K. Gupta
National Science Foundation

Glen Henshaw
U.S. Naval Research Lab

Rich Ivey
BAE Systems

Brian Julian
MIT Lincoln Laboratory

Vijay Kumar
University of Pennsylvania

Alberto Lacaze
Robotic Research

Chris Mailey
AUVSI

Mario Mairena
AUVSI

Jay McConville
Chandler/May Inc.

Adrian Michalicek
Lockheed Martin

Wesley Mitchell
Lockheed Martin

Robin Murphy
Texas A&M University

Don Nimblet
Lockheed Martin

John Northrop
John H. Northrop & Associates

James Overholt
TARDEC

Mike Passaretti
Honeybee Robotics

Robin Pope
SAIC

Mike Ryan
BRTRC

Adam Schuman
Chatten Associates

Jim Shinn
ARA

Bill Thomasmeyer
Robotics Technology Consortium

Eduardo Tores-Jara
Worcester Polytechnic Institute

Tom Van Doren
HDT Robotics

Rocco Wall
Integrated Microwave Technologies

Peter Wells
QinetiQ North America

Joe Zinecker
Lockheed Martin

Noah Zych
Oshkosh Corporation