May 21, 2009

A Roadmap for US Robotics
From Internet to Robotics

Organized by

Georgia Institute of Technology

University of Southern California

Johns Hopkins University

University of Pennsylvania

University of California, Berkeley

Rensselaer Polytechnic Institute

University of Massachusetts, Amherst

University of Utah

Carnegie Mellon University

Tech Collaborative

Sponsored by

Computing Community Consortium (CCC)

Computing Research Association (CRA)

Table of Contents

Overview

Robotics as a Key Economic Enabler

Over the past 50 years, robots have been primarily used to provide increased accuracy and throughput for particular, repetitive tasks, such as welding, painting, and machining, in hazardous, high volume manufacturing environments. Automating such dirty, dull, and dangerous functions has mostly involved implementing customized solutions with relatively specific, near term value. Although a sizeable "industrial" robotics industry has developed as a result, the applications for such first generation robotics solutions have proven to be relatively narrow and largely restricted to static, indoor environments, due to limitations in the enabling technology.

Within the past five years, however, tremendous advancements in ***robotics technology*** have enabled a new generation of applications in fields as diverse as agile manufacturing, logistics, medicine, healthcare, and other commercial and consumer market segments. Further, it is becoming increasingly evident that these early, next generation products are a harbinger of numerous, large scale, global, robotics technology markets likely to develop in the coming decade. Owing to the inexorable aging of our population, the emergence of such a next generation, "robotech" industry will eventually affect the lives of every American and have enormous economic, social, and political impact on the future of our nation.

Unfortunately, the United States lags behind other countries in recognizing the importance of robotics technology. While the European Union, Japan, Korea, and the rest of the world have made significant R&D investments in robotics technology, the U.S. investment, outside unmanned systems for defense purposes, remains practically non-existent. Unless this situation can be addressed in the near future, the United States runs the risk of abdicating our ability to globally compete in these emerging markets and putting the nation at risk of having to rely on the rest of the world to provide a critical technology that our population will become increasingly dependent upon. *Robotech clearly represents one of the few technologies capable in the near term of building new companies and creating new jobs and in the long run of addressing an issue of critical national importance.*

To articulate the need for the United States to establish a national robotech initiative, over 140 individuals from companies, laboratories, and universities from across the country joined forces to produce a definitive report that (1) identifies the future impact of robotics technology on the economic, social, and security needs of the nation, (2) outlines the various scientific and technological challenges, and (3) documents a technological roadmap to address those challenges. This effort was sponsored by the Computing Community Consortium (CCC) and led by 12 world class researchers from the leading robotics academic institutions in the United States. The project included three application oriented workshops that focused on efforts across the manufacturing, healthcare/medical, and services robotics markets; plus one on blue-sky research that addressed a number of enabling technologies that must be the focus of sustained research and application development in order for the U.S. to remain a leader in robotics technology and commercial development.

What follows is a summary of the major findings across all of the workshops, the opportunities and challenges specific to each of the three targeted markets, and recommended actions that must be taken if the United States is to remain globally competitive in robotics technology. Detailed reports from each of the four workshops are also available.

Roadmap Results: Summary of Major Findings

- Robotics technology holds the potential to transform the future of the country and is likely to become as ubiquitous over the next few decades as computing technology is today.
- The key driver effecting the long term future of robotics technology is our aging population both in terms of its potential to address the gap created by an aging work force as well as the opportunity to meet the healthcare needs of this aging population.
- Led by Japan, Korea, and the European Union, the rest of the world has recognized the irrefutable need to advance robotics technology and have made research investment commitments totaling over $1 billion; the U.S. investment in robotics technology, outside unmanned systems for defense purposes, remains practically non-existing.
- Robotics technology has sufficiently advanced, however, to enable an increasing number of "human augmentation" solutions and applications in a wide range of areas that are pragmatic, affordable, and provide real value.
- As such, robotics technology offers a rare opportunity to invest in an area providing the very real potential to create new jobs, increase productivity, and increase worker safety in the short run, and to address the fundamental issues associated with economic growth in an era significant aging of the general population and securing services for such a population.
- Each workshop identified both near and long term applications of robotics technology, established 5, 10, and 15 year goals for the critical capabilities required to enable such applications, and identified the underlying technologies needed to enable these critical capabilities.
- While certain critical capabilities and underlying technologies were domain-specific, the synthesis effort identified certain critical capabilities that were common across the board, including robust 3D perception, planning and navigation, human like dexterous manipulation, intuitive human-robot interaction, and safe robot behavior.

Market Specific Conclusions

Manufacturing

The manufacturing sector represents 14% of the U.S. GDP and about 11% of the total employment. Up to 75% of the net export of the U.S. is related to manufacturing. This sector represents an area of significant importance to the general economic health of the country.

In manufacturing much of the progress and the processes involving robotics technology historically have been defined by the automotive sector and have been very much driven by price and the need to automate specific tasks particular to large volume manufacturing. The new economy is much less focused on mass manufacturing, however, and more concentrated on producing customized products. The model company is no longer a large entity such as GM, Chrysler, or Ford, but small and medium sized enterprises as for example seen in the Fox Valley or in the suburbs of Chicago. The need in such an economy is far more dependent on higher degrees of adaptation, ease of use, and other factors that enable small runs of made to order products. Although the United States has continued to lead the world over the last decade in increasing manufacturing productivity, it is becoming increasingly difficult for us to compete with companies in low-salary countries producing the same products using the same tools and processes. Through the development and adoption of next generation robotics technology and the advancement of a more highly trained workforce, however, it is possible for the United States to continue to lead the world in manufacturing productivity, especially for small and medium sized companies. Doing so will enable the nation to maintain a strong, globally competitive manufacturing base, ensure our continued economic growth, and help safeguard our national security.

Logistics

The efficiency of logistics processes is essential to most aspects of our daily lives from mail delivery to the availability of food in grocery stores. The United States currently imports in excess of 100,000 containers daily, the contents of which must be processed, distributed and made available to customers. Robotics technology is already being used to automate the handling of containers at ports in Australia and elsewhere and has the potential to improve the inspection process as well. Once they leave the port or point of origin, the movement of goods usually entails multiple steps. The distribution of food from farmers to grocery stores, for example, involves several phases of transportation and handling. Although a significant portion of food prices is directly related to these transportation/logistics costs, less than 15% of the end to end distribution process has been considered for automation. Next generation robotics technology has the potential to enable greater optimization of such logistics processes and reduce the price of food and other goods by several percent. In order to realize this potential, however, there is a need to provide new methods for grasping and handling of packages and new methods for sensing and manipulation of objects.

Medical Robots

Over the last decade significant progress has been made in medical robotics. Today several thousand prostate operations are performed using minimally invasive robots, and the number of cardiac procedures is also increasing significantly. There are significant advantages associated with robotics enabled minimally invasive surgery, including smaller incisions, less time spent in the hospital, less risk of infection, faster recovery, and fewer side effects. Overall the quality of care is improved and due to shorter periods away from work there are significant economic benefits. Although the number of medical procedures for which robots are used is still relatively small, their use is expected to broadly expand as advances in next generation robotics technology provide improved facilities for imaging, feedback to the surgeon and more flexible integration into the overall process. As such, medical robotics holds the potential to have an enormous impact, economic and otherwise, as our population ages.

Healthcare

The number of people suffering strokes and other injuries attributable to aging will continue to increase and become even more pronounced. When people suffer an injury or a stroke it is essential to have them undergo regularly scheduled physical therapy sessions as soon as possible to ensure that they achieve as full a recovery as possible. Often, however, the rehabilitation/training occurs away from home and due to shortage of therapists there are often serious constraints on scheduling. Next generation robotics technology will increasingly enable earlier and more frequent sessions, a higher degree of adaptation in the training, and make it possible to perform a certain percentage of these training sessions at home. By facilitating more consistent and personalized treatment regimens in this fashion, robotics enabled rehabilitation offers the potential for faster and more complete patient recovery. Robotics technology is also beginning to be used in healthcare for the early diagnosis of autism, memory training for people with dementia, and other disorders where personalized care is essential and there is an opportunity to realize significant economic benefits. Today early products are on the market, but the full potential is still to be explored.

Services

The use of robotics technology in the service industry spans professional and domestic applications. In professional services, emerging applications include improved mining, automated harvesters for agriculture and forestry, and cleaning of large scale facilities. Domestic services applications include cleaning, surveillance, and home assistance. Today more than 4 million automated vacuum cleaners have already been deployed and the market is still growing. So far only the simplest of applications have been pursued, but an increasingly services-based U.S. economy offers significant potential for the automation of services to improve quality and time of delivery without increasing costs. As people work longer hours, there is a need to provide them with assistance in their homes to provide time for leisure activities. A big challenge in service robotics will be the design of high performance systems in markets that are price sensitive.

International Context

The promise of a thriving, next generation robotech industry has of course not gone unnoticed. The European Commission recently launched a program through which 600 mill Euros are invested in robotics and cognitive systems with a view to strengthen the industry, particularly in manufacturing and services. Korea has launched a comparable program as part of their 21st century frontier initiative, committing to invest $1B in robotics technology over a period of 10 years. Similar, but smaller programs are also in place in Australia, Singapore, and China. In the United States, funding has been committed for unmanned systems within the defense industry, but very few programs have been established in the commercial, healthcare, and industrial sectors. Although the industrial robotics industry was born in the United States, global leadership in this area now resides in Japan and Europe. In areas such as medical, healthcare and services, the United States has similarly established an early leadership position, but there are fast followers and it is not clear that we will be able to sustain our leadership position for long without a national commitment to advance the necessary robotics technology.

Further information

http://www.us-robotics.us

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

Chapter 1

Robotics and Automation Research Priorities for U.S. 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. On 19 December 2008, the U.S. government announced $13.4 billion in emergency federal loans to General Motors and Chrysler to facilitate restructuring and encourage new research and development – a clear example the U.S. of playing catch-up rather than taking technological leadership.

Federal Investments in research in manufacturing can revitalize American manufacturing. Investing a small portion of our national resources into a science of cost-effective, resource-efficient manufacturing would benefit American consumers and support millions of workers in this vital sector of the U.S. economy. It would allow our economy to flourish even as the ratio of workers to pensioners continuously decreases. Such a research and development program would also benefit the health care, agriculture, and transportation industries, and strengthen our national resources in defense, energy, and security. The resulting flurry of research activity would greatly improve the quality of "Made in the U.S.A." and invigorate productivity of U.S. manufacturing for the next fifty years.

Robotics is a key transformative technology that can revolutionize manufacturing.

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 go off-shore without it; 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.

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

This document summarizes the activities and results of a workshop on manufacturing and automation robotics that was supported by a grant from the Computing Community Consortium of the Computing Research Association. This workshop was the first of four organized on various areas of robotics, with the overall objective being the creation of a compelling vision for robotics research and development, and roadmaps for advancement of robotics technologies to maximize economic impact. 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."

Above: Robots are now commonplace in automotive manufacturing. (Source: ABB Robotics)

Below: Lightweight robots are entering the market for high speed material handling, for example in food processing and electronics packaging. (Source: Adept)

The Manufacturing and Automation Robotics Workshop was held on June 17, 2008 in Washington DC (http://www.us-robotics.us/?page_id=9). 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. White papers on perceived industrial needs were solicited from academic researchers. Authors of accepted whitepapers (available at http://www.us-robotics.us/?page_id=14) were invited to attend the workshop, where authors from industry were also invited to give short presentations on present and future uses of robotics in their companies.

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 [E07]. Fully 75% of the net export of the U.S. is related to manufacturing [State04], 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 manufacturing industry that 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 11% 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 [DoC04]. Our per capita productivity continues to increase with over a 100% increase over the last three decades. Indeed it is this rising productivity that keeps U.S. manufacturing competitive in the midst of recession and recovery and in the face of the amazing growth in China, India, and other emerging economies. Much of this productivity increase and efficiency can be attributed to innovations in technology and the use of technology in product design and manufacturing processes.

However, this dynamic is also changing. Ambitious foreign competitors are investing in fundamental research and education that will improve their manufacturing processes. On the other hand, the fraction of the U.S. manufacturing output that is being invested in research and development has essentially remained constant over this period. The U.S. share of total research and development funding the world has dropped significantly to only 30%. Our foreign competitors are using the same innovations in technology with, in some cases, significantly lower labor costs to undercut U.S. dominance, so U.S. manufacturing industry is facing increasing pressure. Our balance of trade in manufactured goods is dropping at an alarming $50 billion per decade. Additionally, with our aging population, the number of workers is also decreasing rapidly and optimistic projections point to two workers per pensioner in 2050 [E07]. 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. Finally, dramatic advances in robotics and automation technologies are even more critical with the next generation of high-value products that rely on embedded computers, advanced sensors and microelectronics requiring micro- and nano-scale assembly, for which labor-intensive manufacturing with human workers is no longer a viable option.

In contrast to the U.S., China, South Korea, Japan, and India are investing heavily in higher education and research [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 is investing $100M per year for 10 years (2002-2012) into robotics research and education as part of their 21 century frontier program. The European Commission is investing $600M into robotics and cognitive systems as part of the 7th Framework Programme. 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.

Novel Mobile robots are enabling new paradigms in logistics and warehouse management with improved productivity, speed, accuracy, and flexibility. (Source: KIVA Systems)

2.2. Growth Areas

The Department of Commerce and the Council on Competitiveness [CoC08, DoC04] have analyzed a broad set of 280 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%	0-120%
IP Companies	21%	15-26%
Healthcare/eldercare	62%	6-542%
Entertainment/toys	6%	4-17%
Media / Games	14%	-2-36%
Home appliances	1%	-4-7%
Capital equipment	8%	-4-20%
Automotive	0%	-11-13%
Logistics	21%	4-96%
Automation	4%	2-8%

Consolidated annual growth rates over a set of 280 U.S. companies for the period 2004-2007.

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

2.3. A Vision for Manufacturing

U.S. manufacturing today is where database technology was in the early 1960's, a patchwork of ad hoc solutions that lacked the rigorous methodology that leads to scientific innovation. In 1970 when Ted Codd, an IBM mathematician, invented relational algebra, an elegant mathematical database model that galvanized federally funded research and education leading to today's $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].

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., human are better at dealing with unexpected events to keep production lines running, while robots have better precision and repeatability, and can lift heavy parts), (e) reduce expensive medical problems, e.g., carpal tunnel syndrome, back injuries, burns, and inhalation of noxious gases and vapors, and (f) reduce time in pipeline for finished goods, allowing systems to be more responsive to changes in retail demand.

Investments in research and education in manufacturing can revitalize American manufacturing. Investing a small portion of our national resources into a science of cost-effective, resource-efficient manufacturing would benefit American consumers and support millions of workers in this vital sector of the U.S. economy. Such investments would benefit health care, agriculture, and transportation, and strengthen our national resources in defense, energy, and security. The resulting flurry of research activity would invigorate the quality and productivity of "Made in the U.S.A." for the next fifty years.

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 to create the desired revitalization of manufacturing in the U.S.

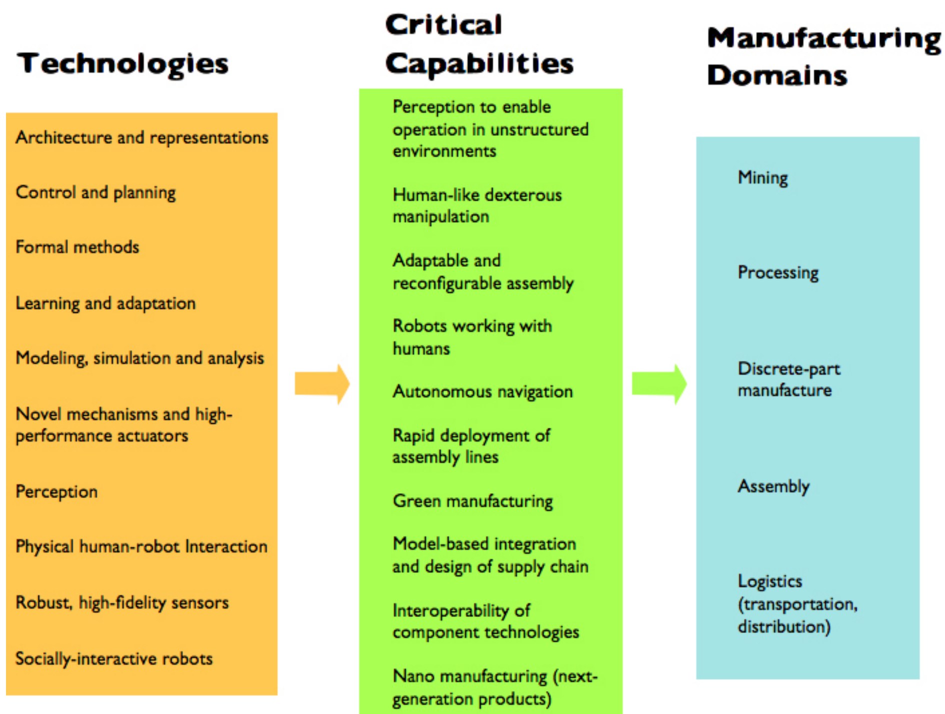

Figure 1: The roadmap process: Research and development is needed in technology areas that arise from the critical capabilities required to impact manufacturing application domains.

3.2. Robotics and Manufacturing Vignettes

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

Vignette 1: Assembly line assistant robots

> An automotive manufacturer experiences a surge in orders for its new electric car design and needs to quickly merge its production capability with other earlier models already in production. Assembly tasks are rapidly reallocated to accommodate the new more efficient car model. A set of assembly line assistant robots are brought in and quickly configured to work alongside the retrained human workers on the new tasks. One practice-shift is arranged for the robot's sensor

systems and robot learning algorithms to fine-tune parameters, and then the second shift is put into operation, doubling plant output in four days. Then, a change by a key supplier requires that the assembly sequence be modified to accommodate a new tolerance in the battery pack assembly. Engineers use computational tools to quickly modify the assembly sequence: then they print new instructions for workers and upload modified assembly programs to the assistant robots.

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

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

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

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

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

3.3. Critical Capabilities for Manufacturing

In this section, we briefly discuss the critical capabilities and give examples of possible 5, 10, and 15 year milestones. After this, in Section 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, automated guided vehicles for material handling in assembly lines, and logistics support operations like warehousing and distribution. Enabling safe autonomous navigation in unstructured environments with static obstacles, human-driven vehicles, pedestrians and animals will require significant investments in component technologies. The roadmap for autonomous navigation consists of the following milestones.

5 year: Autonomous vehicles will be capable of driving in any modern town or city with clearly lit and marked roads and demonstrate safe driving comparable to a human driver. Performance of autonomous vehicles will be superior to that exhibited by human drivers in such tasks as navigating through an industrial mining area or construction zone, backing into a loading dock, parallel parking, and emergency braking and stopping.

10 years: Autonomous vehicles will be capable of driving in any city and on unpaved roads, and exhibit limited capability for off-road environment that humans can drive in, and will be as safe as the average human driven car.

15 years: Autonomous vehicles will be capable of driving in any environment in which humans can drive. Their driving skill will be indistinguishable from humans except that robot drivers will be safer and more predictable than a human driver with less than one year's driving experience.

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 to production and distribution of finished products. 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. Human-like Dexterous Manipulation

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

5 years: Low-complexity hands with small numbers of independent joints will be capable of robust whole-hand grasp acquisition.

10 years: Medium-complexity hands with tens of independent joints and novel mechanisms and actuators will be capable of whole-hand grasp acquisition and limited dexterous manipulation.

15 years: High-complexity hands with tactile array densities approaching that of humans and with superior dynamic performance will be capable of robust whole-hand grasp acquisition and dexterous manipulation of objects found in manufacturing environments used by human workers.

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 ManufacTuring". If achieved, as with databases and computers, would enable interoperability of components and subsystems and higher quality, more reliable products, reduced costs, and faster delivery. Accordingly our roadmap should include achievements that demonstrate the following milestones.

5 years: Safe, provably-correct designs for discrete part manufacturing and assembly so bugs are not created during the construction of the manufacturing facility.

10 years: Safe, provably-correct designs for the complete manufacturing supply chain across multiple time and length scales so bugs are not created during the design of the manufacturing supply chain.

15 years: Manufacturing for Next Generation Products: With advances in micro and nano-scale science and technology, and new processes for fabrication, we will be able to develop safe, provably-correct designs for any product line.

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 necessitate robots to be smarter, more flexible, and able to operate safely in less structured environments shared with human workers. In product flow layouts for example, robots and other machines go to various operation sites on the product (e.g., an airplane or a ship) to perform their tasks, whereas in a functional layout, the product travels to various machines. The challenges of one-of-a-kind manufacturing exacerbate these difficulties. The roadmap for perception includes the following milestones.

5 years: 3-D perception enabling automation even in unstructured typical of a job shop engaged in batch manufacturing operations

10 years: Perception in support of automation of small lot sizes, for example, specialized medical aids, frames for wheelchairs, and wearable aids.

15 years: Perception for truly one-of-a-kind manufacturing including customized assistive devices, personalized furniture, specialized surface and underwater vessels, and spacecrafts for planetary exploration and colonization.

3.3.8. Intrinsically Safe Robots Working with Humans

Robotics has made significant progress toward enabling full autonomy and shared autonomy in tasks such as driving vehicles, human physical therapy, and carrying heavy parts (using cobots). Leveraging these advances to enable autonomy and shared autonomy in other tasks such as assembly and manipulation poses a significant challenge. Automotive industry experts recognize the benefits of automation support for human workers either in the form of humanoid assistants or smart machines that safely interact with human workers. To define research milestones we propose three levels of assembly line ability:

1. Level I Ability: humans require no special skills and < 1 hour of training. examples: pick and place, insertion, packing. A canonical benchmark that can be used for testing and comparison between groups might be generic tasks such as threading and unthreading a standard 1" nut and bolt.
2. Level II Ability: humans require minor skills and 1-10 hours of training. examples: cutting / shaping, soldering, riveting. A canonical benchmark might be disassembling and reassembling a specific standard flashlight.
3. Level III Ability: humans require skill and > 10 hours of training. examples: specified standard welding, machining, inspecting benchmarks.

The roadmap for robots working with humans is as follows.

5 years: Demonstrate a prototype assembly-line robot with sensors that can detect and respond to human gestures and movement into its workspace while consistently performing at Level I ability (see above) alongside a human for 8 hours without requiring any intervention from the people nearby.

10 years: Demonstrate a prototype assembly-line robot with sensors that can detect and respond to human gestures and movement into its workspace while consistently performing at Level II ability alongside a human for 40 hours without requiring any intervention from the people nearby.

15 years: Demonstrate a commercially available assembly-line robot with sensors that can detect and respond to human gestures and movement into its workspace while consistently performing at Level III ability alongside a human for 80 hours without requiring any intervention from the people nearby.

3.3.9. Education and Training

The U.S. can only take advantage of new research results and technology if there is 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-yr college and university offers concentrations in robotics to augment many Bachelors, Masters, and PhD degrees.

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

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 “Human-like 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.

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 this task is being done partly or wholly by a human, 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. 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. 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

Robots in future factories will be in physical contact with humans and other robots, if not directly, then through an object being grasped by both simultaneously. Inadvertent contact may also occur. When robots are collaborating with humans, they must be able to recognize the human activities to maintain proper task synchrony. Finally, robots must be able to communicate with humans in multiple ways; verbally and non-verbally, and must be easy to train. 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 intrinsically safe robots with fail-safe operating systems and tools to verify the safety and correctness of robot programs.

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.

5. References

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

[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.

6. Contributors

This report has its origins in presentations and discussions at a workshop on manufacturing and automation robotics that took place June 17, 2008 in Washington, DC. The report is part of the CCC study on Robotics. The Computing Community Consortium (CCC) is a project managed by the Computing Research Association (CRA) and is sponsored by the National Science Foundation (NSF). The present report has been authored by the workshop organizers and does not necessarily reflect the opinion of CRA, CCC or NSF. The responsibility of the report lies entirely with the authors.

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

Tom Batzinger
GE

Kostas Bekris
UNReno

Karl Bohringer
UW

Joe Bordogna
UPenn

Gary Bradski
Willow Garage

Oliver Brock
UMass

Jeff Burnstein
RIA

Thomas Fuhlbrigge
ABB, Inc

Henrik I Christensen
Georgia Tech

Roger Eastman
Loyola

Aaron Edsinger
Meka

Erica Fuchs
CMU

Ken Goldberg
UC Berkeley

Tom Henderson
Utah

William Joyner
Semiconductor Research Corporation

Lydia Kavaraki
URice

Clint Kelly
Georgia Tech (former SAIC)

Alonzo Kelly
CMU

Vijay Kumar
UPenn

Dinesh Manocha
UNC

Andrew McCallum
UMass

Pieter Mosterman
Mathworks

Elena Messina
NIST

Todd Murphey
Colorado

Richard Alan Peters
Vanderbilt

Stuart Shephard
KUKA

Sanjiv Singh
CMU

Larry Sweet
C&S Whole Grocers

Jeff Trinkle
RPI

Jason Tsai
FANUC

James Wells
GM

Peter Wurman
Kiva Systems

Tom Yorio
Corning

Mingjun Zhang
UTK

Chapter 2

A Research Roadmap for Medical and Healthcare Robotics

Executive Summary

Motivation and Scope

Several major societal drivers for improved health care access, affordability, quality, and personalization that can be addressed by robotic 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. Tele-medicine 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 efforts are developing affordable 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, and 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, and enabling people with disabilities, whose numbers are on the rise, to go (back) into the workforce. Today, the US 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.

Participants

The workshop contributors consisted of experts in surgical robotics, prosthetics, implants, rehabilitation robotics, and socially assistive robotics, as well as representatives from industry ranging from large corporations to startups, and representatives from the health insurance provider community. All participants contributed insights from their communities and areas of expertise; many common interests and challenges were identified, informing the road mapping 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 following problem domains were identified as those of 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 technological and research challenges: intuitive human-robot interaction and interfaces; automated understanding of human behavior; automated understanding emotional and physiological state; long term adaptation to user's changing needs; quantitative diagnosis and assessment; context-appropriate guidance; image-guided intervention; high dexterity manipulation at any scale; sensor-based automated health data acquisition; 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 lack of funding for interdisciplinary integrative projects that bring together expertise in engineering, health (and business) and develop and evaluate complete systems in human subjects studies was 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 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 successfully delivering physical and occupational therapy. 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 to assist recovery after stroke, the leading cause of permanent disability in the US. The future potential for robots in convalescence and rehabilitation is even greater. Experiments have also 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 after hospitalization. Such systems also have potential as intervention and therapeutic tools for behavioral disorders including such pervasive disorders as autism spectrum disorder, ADHD, and others prevalent among children today.

Robotics technology also has a role in augmenting basic research into human health. The ability to create a robotic system that mimics biology is one 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.

The spectrum of robotic system niches 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 health care 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 enable develop new medical procedures and devices, such as micro-scale interventions and smart prostheses, which would substantially improve risk-benefit and cost-benefit ratios. More effective methods of training of medical practitioners would lower the number of medical errors. Objective approaches for accountability and certification/assessment also contribute to this goal. Ideally, all these improvements would lower costs to society by lowering impact on families, caregivers, and employers. More directly, health care costs would be lowered due to improved quality (fewer complications, shorter hospital stays, and increased efficiency).

Population factors related to economics 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 health care, 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 health care is most directly related to its affordability. Access to physically interactive therapy robots promise to reduce the cost of clinical rehabilitative care and are the focus of an ongoing Veteran's Administration study of their cost-effectiveness. Socially assistive robotics efforts are working toward methods that could provide affordable in-home technologies for motivating and coaching exercise for both prevention and rehabilitation. It is also a promising domain for technologies for care taking for the elderly, toward promoting 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 health care is also related to location. When disasters strike and result in human injury, distance and unstructured environments are obstacles to providing on-site care and removing the injured from the scene. This has been repeatedly demonstrated in both natural disasters (such as earthquakes and hurricanes) and man-made disasters (such as terrorist attacks). Similar problems occur in the battlefield; point-of-injury care is needed to save the lives of many military personnel. Some environments, such as space, undersea, and underground (for mining) are inherently far from medical personnel. Finally, rural populations can live prohibitively far from medical centers that provide specialized health care. Telemedicine and assistive robotics can provide access to treatment for people outside populated areas and in disaster scenarios.

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

Associated with the aging population is increased prevalence of injuries, disorders and diseases. Furthermore, across the age spectrum, health trends indicate significant increases in life-long conditions including diabetes, autism, obesity, and cancer. The American Cancer Society estimates that 1,437,180 new cancer cases (excluding the most common forms of skin cancer) will be identified in the US in 2008. Furthermore, the probability of developing invasive cancers increases significantly with age [ACS Cancer Facts and Figures 2008].

These trends are producing a growing need for personalized health care. For example, the current rate of new strokes is 750,000 per year, and that number is expected to double in the next two decades. 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 most prevalent among older patients, Cerebral Palsy (CP) is most prevalent among children. About 8,000 infants are diagnosed with CP each year and there are over 760,000 persons in the US manifest symptoms of CP. Further, the number of neurodevelopmental and cognitive disorders is on the rise, including autism spectrum disorder, attention deficit and hyperactivity disorder, and others. Autism rates alone have quadrupled in the last quarter century, with one in 150 children diagnosed with the deficit today. Improved outcomes from early screening and diagnosis and transparent monitoring and continual health assessment will lead to greater cost savings, as can effective intervention and therapy. These factors will also offset the shrinking size of the healthcare workforce, while affordable and accessible technology will facilitate wellness, personalized, and home-based health care.

Increasing life-long independence thus becomes a key societal driver. It includes increasing the ability to age in place (i.e., to enable the elderly to stay 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 providing independence for caregivers, who are increasingly employed and such care is increasing informal because the economics of in-home health care are unaffordable. Lifelong health education and literacy would facilitate prevention and can be augmented by improved safety and monitoring to avoid mismedication, ensure consistency in taking medication, monitoring for falls, lack of activity, and other signs of decline.

All of the above have the effect of maintaining and improving productivity of the workforce and increasing 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 workforce (and contribute to social security) would also offset the current reduction in available labor/workforce.

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

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, and
- transcend human physical limitations in performing surgery and other interventional procedures, while still affording human control over the procedure.

Two decades after the first reported robotic surgical procedure, surgical robots 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 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 surgeon's field of view, and recognition of surgical motions and patient state to evaluate performance and predict health outcomes.

If the right information is available, many medical procedures can be planned ahead of time and executed in a reasonably predictable manner, with the human exercising mainly supervisory control over the robot. By analogy to industrial manufacturing systems, this model is often referred to as "Surgical CAD/CAM" (Computer-Aided Design and Computer-Aided Manufacturing). Examples include preparation of bone for joint reconstructions in 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 load on their backs while running) and rehabilitation in the clinic. However, 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

A patient suffering from neuromuscular injuries or diseases, such as occur in the aftereffects of stroke, often benefits 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 patient makes upper extremity or lower extremity movements physically assisted (or resisted) by a human therapist and/or robot, helps people re-learn how to move. This process is time-consuming and labor-intensive, but pays large dividends in terms of patient health care costs and return to productive labor. As an alternative to human-only therapy, a robot has several key advantages for intervention:

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

A 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 subject's real-time response. For example, the commercially available MIT-Manus rehabilitation robot showed improved recovery of both acute and chronic stroke patients. Another exciting implication of sensory-motor therapy with robots is that they can help neuroscientists improve their general understanding brain function. Through knowledge of robot-based perturbations to the patient and quantification of the response of patients with damage to particular areas of the brain, robots can make unprecedented stimulus-response recordings. In order to optimize automated rehabilitation therapies, robots and experiments must be developed to elucidate the relationship between external mechanical forces and neural plasticity. The understanding of these relationships also give neuroscientists and neurologists insight into brain function, which can contribute 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 and motivate 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 regiments are prescribed. For example, an early system was demonstrated in the cardiac ward, encouraging and coaching patients to perform spirometry exercises ten times per hour. Such systems can serve both as force multipliers in heath care delivery, providing more care to more patients, but also as a means of delivering personalized medicine and care, providing more 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, TBI, etc., as discussed in Section 1.2), the available health care needed to provide supervision and coaching for such behavior therapy is already lacking and on a recognized steady decline.

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 the 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 robots 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 tool 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 occur at any stage of life, there is a clearly growing need for personalized care for individuals with special needs. 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 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.

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 health care 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 health care. Better, more effective and accessible, as well as personalized ways of encouraging people eat right, exercise, 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 seek causes and 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 trained athletes. Wearable devices that monitor physiologic responses and interact with robotic and computer-based systems also have the potential to promote personalized wellness regiments and facilitate early detection and continuous assessment of disorders. In this context, robotics is providing enabling technologies that inter-operate with existing systems (e.g., laptop and desk-top computers, wearable devices, in-home sensors, etc.) in order to leverage advances across fields and produce a broad span of usable technologies toward improving quality of life (QoL).

3. Key Challenges and Capabilities

3.1. Motivating Exemplar Scenarios

3.1.1. Surgery and Intervention

A pre-operative image or blood test indicates that a patient may have cancer in an internal organ. The patient receives a Magnetic Resonance Imaging (MRI) scan, from which the existence of cancerous tissue is confirmed. Based spatial extent of the cancer identified through image processing and tissue models, an optimal surgical plan is determined. A surgeon uses a very minimally invasive, MRI-compatible teleoperated robot to remove the cancerous tissues. The robot is sufficiently dexterous that the surgery can be performed through a natural orifice, so no external cuts are made in the patient. During the procedure, the surgeon sees real-time images, is guided by the surgical plan, and receives haptic feedback to enable palpation and appropriate application of forces to tissue. The cancerous tissue is removed with very little margin and the patient recovers quickly with little pain and no scarring.

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 artificial hand, and receives multi-modal sensory feedback (e.g., force, texture, temperature), allowing her 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 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, even sometimes resisting the patient's motion in order to help him learn to make corrective motions. Data is recorded throughout therapy, which allows both the therapist and the robotic system 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 and project authority and competence but retain autonomy for the user while motivating continued exercises. This shortens convalescence 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 others) 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. It then serves as a social catalyst for play with other children, first in the home and then in the school lunchroom and eventually playground. Throughout, the robot collects quantitative data on user/patient behavior that can be analyzed both automatically and by healthcare providers for continuous assessment and personalized therapy/treatment/intervention delivery.

3.1.5. Personalized Care for Special-Needs Populations

Personalized robots are given to the elderly and physically and/or cognitively disabled (e.g., 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, to minimize isolation and facilitate activity and integration in society. Robots can 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 (decreasing sense of balance, lessened social interaction, diminishing vocalizations, lack of physical activity, increased isolation from family/friends, etc.) that trigger the need for early intervention.

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, etc. These robotic systems are purchased as easily and readily as current personal computers, and easily configured for the user and made inter-operable with other computing and sensory resources of the user environment. For example, robots that monitor the amount of physical activity of a overweight diabetic user to promote increased physical activity, and require 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 health care challenges noted in Sections 1 and 2 and achieve the exciting scenarios described immediately above in Section 3.1, we have developed a list of major capabilities that robotic system must have for ideal integration into medicine and health care. These capabilities, in turn, motivate research into the technologies described in Section 4.

3.2.1. Intuitive Physical Human-Robot Interaction and Interfaces

The use of robotics in medicine inherently involves physical interaction between caregivers, patients, and robots – in all combinations. Developing intuitive physical interfaces between humans and robots requires all the classic elements of a robotic system: sensing, perception, and action. A great variety of sensing and perception tasks are required, including recording the motions and forces of a surgeon to infer their intent, determining the mechanical parameters of human tissue, and estimating the forces between a rehabilitation robot and a moving stroke patient. 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 caregiver or a patient. We need to consider systems that involve many human senses, the most common of which are vision, haptics (force and tactile), and sound.

We need to consider systems that involve many human senses.

A major reason why systems involving physical collaboration between humans and robots are so difficult to design well is that, from the perspective of a robot, humans are extremely uncertain. Unlike a passive, static environment, humans change their motion, strength, and immediate purpose on a regular basis. This 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 surgery. 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, and thus cannot simply be thought of as an external system input. In addition, the loop is often closed with both human force and visual feedback, each with its own errors and delays – this 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 approaches to solving these problems, which can be used in parallel: modeling the human with as much detail as possible, 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. 5-, 10-, and 15-year goals for this capability focus on increasing complexity and uncertainty of the task at hand.

- In 5 years, robots should be able to have sophisticated understanding of desired human motion based on external sensors and brain-machine interfaces. This is especially essential for prosthesis design, and requires an appropriate mapping between human thoughts and the actions of a robotic prosthetic limb.

- In 10 years, by sensing a human's motions and inferring intent, robots should 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 stroke. By sensing the human's motions and inferring intent, the robot should limit applied force or motion to levels that are useful and intuitive for the user.

- In 15 years, robotic systems should be able to provide the full suite of physical feedback to a human operator, in particular appropriate haptic feedback. A surgeon or caregiver should be able to feel the forces, detailed surface textures, and other physical properties of a remote patient. The environment should be completely immersive, and function at any scale.

3.2.2. Automated Understanding of Human Behavior

Understanding the user's activity and intent are necessary components of human-machine and thus human-robot interaction, in order to respond appropriately and in a timely and safe fashion. Effective health systems must be able to perceive their environment and user. Because human activity is complex and unpredictable, and because vision-based perception is an ongoing challenge in robotics, automated perception and understanding of human behavior requires the integration of data from a multitude of sensors, including those on the robot, in the environment, and worn by the user. Research into algorithms for real-time on-line multi-modal sensor integration us under development, including the application of statistical methods for user modeling based on multi-modal data. Recognition and classification of human activity and intent is of particular interest, in order to enable real-time user interaction and assistance. HRI systems will only be accepted if they are responsive to the user on a time-scale the user finds reasonable (i.e., the system cannot take too long to respond nor can it respond incorrectly too often). Current methods for multi-modal perception have used various means of simplifying the hard problems of real-world object and person recognition and activity recognition and classification. For example, efforts have used color and reflective markers, bar codes, and radio frequency identification tags, all of which require some level of instrumentation of the environment. Minimizing such instrumentation and making it non-intrusive is a necessary aspect of making the technology acceptable.

Key areas of progress and promise include: (1) the use of physiologic sensing as a counterpart to standard on-robot and in-environment sensing the field has focused on to date; (2) leveraging, processing, and utilizing multi-modal sensing on-board, in the environment, and on the user for real-time HRI; and (3) understanding of user affect/emotion.

- In 5 years, robots should be able to have the ability to capture instrumented human behavior (aided with wearable markers) in controlled environments (e.g., physical therapy sessions, doctor's offices) with known structure and expected nature of interactions. Algorithms should be able to use uncertain and noisy data from such sessions to develop models of the user and the interaction.

- In 10 years, robots should be able to automatically classify human behavior from lightly instrumented users (light-weight sensors), in less structured settings (e.g., doctor's offices and homes with less-known structure), visualize those data for the user and the health care provider, and classify the activity into proscribed exercises and other activities for assessment performance. On-line modeling techniques should be able to classify observed activity and predict user performance and upcoming actions with reasonable levels of accuracy.

- In 15 years, robotic systems should be able to detect, classify, predict, and provide coaching for human activity within a known broad context (e.g., exercise, office work, dressing, etc.). The system should be able to provide intuitively visualized data for each user, which will differ based on the user's needs (e.g., the doctor will need a detailed assessment of the motor activity, the caretaker the consistency and accuracy of the exercises, the user a "score" of the activity and some helpful hints for improvement, etc.).

3.2.3. Automated Understanding of Emotional and Physiological State

The ability to automatically recognize emotional states of users in support of appropriate, personalized robot behavior is critical for making personalized robotics effective, especially for health-related applications that involve vulnerable users. Emotion recognition has been studied in voice and speech signals, facial data, and physiologic data. Given the complexity of the problem, emotion understanding, modeling, and classification will directly benefit from strides in all of the areas listed above: activity recognition, physiologic data processing, and multi-modal perception. Emotion understanding requires processing multi-channel data from the user, and reconciling inconsistencies (e.g., between verbal v. facial signals). Incongruence in such signals can confuse the recipient; analogously, human perception of synthetic multi-channel expressions of emotion (e.g., on embodied robots equipped with articulated faces, voices, and bodies) is not yet well understood and merits in-depth research in order to inform principled system design. The power of empathy is well recognized in health care: doctors who are perceived as empathetic are judged as most competent and have the least lawsuits. Creating empathy in synthetic systems is just one of the challenges of perceiving and expressing emotion. Furthermore, 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. Since personality is known to have impact on health outcomes, the ability to perceive, model, and express it and the associated emotions is an important aspect of human-machine interaction aimed at improving human health and quality of life.

Creating empathy in synthetic systems is just one of the challenges.

Physiologic data, such as measures of frustration, fatigue, and interest, are invaluable in understanding the state of the user and enabling robots, and machines in general, to enable them to assist the user and optimize performance. Physiologic data sensors are typically wearable sensors and devices that provide real-time physiologic signals (e.g., heart rate, galvanic skin response, body temperature, etc.). These signals are highly individualized and typically complex to intuitively visualize and usefully analyze. Active research in the field is addressing methods for extracting metrics, such as frustration, and saliency relative to external activity, from physiologic data. Research is also focusing on connecting and accessing bioelectrical signals with wearable or implantable devices. With the exception of some implantable devices, lightweight wearable sensors with wireless capabilities for data transmission and low-weight batteries are not yet readily available. The promise of wearable sensory technologies has been recognized widely and developments toward addressing these issues are in progress. The ability to capture physiologic data in an un-encumbering way and transmit that data to a computer, robot, or caregiver, has great potential for improving health assessment, diagnosis, treatment, and personalized medicine. Such data complement standard robotics sensors (vision, laser, infra red, sonar) and provide invaluable user data for modeling and intelligent human-machine interaction.

- In 5 years, a variety of wearable devices should interface wirelessly with assistive robots to inform the development of user models and state and activity classification algorithms. Multi-modal algorithms should be developed that can take highly uncertain visual data and combine it with other sensory data toward emotion state classification.

- In 10 years, smaller-scale and lighter-weight wireless wearable sensors providing a range of physiologic data should be available as real-time input into algorithms that use population and individual models of the user to detect and classify as well as to some degree predict user physiologic state. Multi-modal algorithms should take inputs from vision and wearable sensors to seamlessly integrate toward reliable real-time physiologic and emotional state recognition.

- In 15 years, off-the-shelf wireless physiologic sensing devices should inter-operate with computer- and robot-based coaching systems that can use the data to develop and apply user models in real-time to facilitate bio-feedback and other forms of feedback to the user and classification of user physiologic and emotional state for facilitating sophisticated human-robot and more generally human-machine interaction.

3.2.4. Long Term Adaptation to User's Changing Needs

The need for system adaptation and learning is especially evident in human-robot interaction domains. Each user has specific characteristics, needs, and preferences to which the system must be attuned. Furthermore, those very characteristics, needs, and preferences can change over time as the user gets accustomed to the system and as the health state of the user changes, both over the short term (convalescence), medium term (rehabilitation) and life-long (life-style changes, aging). To be accepted, usable and effective, robot systems interacting with human users must be able to adapt and learn in new contexts and at extended time-scales, in a variety of environments and contexts.

Challenges in long-term learning include the integration of multi-modal information about the user over time, in light of inconsistencies and changes in behavior, and unexpected experiences. Machine learning, including robot learning, has been adopting increasingly principled statistical methods. However, the work has not addressed the complexities of real-world uncertain data (noisy, incomplete, and inconsistent), multi-modal data about a user (ranging from signal-level information from tests, probes, electrodes, and wearable devices, to symbolic information from charts, questionnaires, and patient interviews), and long-term data (over months and years of treatment).

The ability to interact with the user through intuitive interfaces (gestures, wands, speech) and learn from demonstration and imitation have been topics of active research for some time. They present a novel challenge for in-home long-term interactions where the system is subject to user learning and habituation, as well as diminishing novelty and patience effects. Robotics learning systems have not yet been tested on truly long-term studies (over weeks and months) and life-long learning is not yet more than a concept.

Robot systems interacting with human users must be able to adapt and learn.

Finally, because learning systems are typically difficult to assess and analyze, it is particularly important that such personalized, adaptive technologies be equipped with intuitive visualization tools of their system state as well as the health-state of the user.

Taking these challenges into account, an ideal adaptive, learning health-care robot system would be able to predict changes in the health state of the user/patient and adjust the delivery of its services accordingly; it would adjust its methods for motivating, encouraging, and coaching the user continually, to retain its appeal and effectiveness by sustaining user engagement over the long term. Such a system would have quantitative metrics to show positive health outcomes based on health professional-prescribed convalescence/intervention/therapy/prevention methods.

- In 5 years, adaptive and learning systems should use increasing amounts of real-world health data and be shown to operate on such data in spite of its noisy, dynamically changing and complex nature. User models should enable the system to adapt its interaction style with the user to improve user task performance within a particular context (e.g., specific exercise).
- In 10 years, adaptive and learning systems should be extended to operate on long-term data (months and more) and multi-modal patient data toward more general-purpose comprehensive user modeling beyond a particular context (e.g., from a specific exercises to overall daily activity).
- In 15 years, adaptive and learning systems should be available as software on standard computers, facilitating in-home health-care monitoring and wellness promotion. Taking user-provided data over time and from multiple modalities as well as healthcare provider information (as part of checkout procedure, for example), to continue to update comprehensive models of user health state, and visualize and report those to the user, family, and healthcare providers, and use those to continue to optimize human-machine interaction for improved health practices.

3.2.5. Quantitative Diagnosis and Assessment

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 health such as strength, interact with the patient in social ways to acquire behavioral data (e.g., eye gaze, gesture, joint attention, etc.) 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. Quantitative diagnosis and assessment requires sensing of the patient, 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 replace them. As robots facilitate aging in place (e.g., in the home), automated assessment becomes more important as a means to alert a caregiver, who may 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 the field of robotics or have been used by robots in diagnosis and assessment. Others, such as using intelligent socially assistive robots to quantify behavioral data, are entirely novel and present new ways of treating data that had, to date, been only qualitative. The myriad steps in diagnosis/assessment need to each be improved and then combined into a seamless process. These steps include: apply stimulus (if necessary), acquire data, make a diagnosis or assessment of patient health, relay the information in a useful form with appropriate level of detail to a caregiver, integrate caregiver input to revise diagnosis/assessment, and perform actions what will allow collection of more or different data (if needed) to make a better informed diagnosis/assessment. In some settings, this process is self-contained (i.e., administered within a controlled session) while in others it may be a more open-ended procedure (i.e., administered in a natural environment, such as the home). Achieving this sophisticated process requires reaching several major milestones.

- In 5 years, a robot should be able to extract relevant metrics, such as arousal, heart rate, movement capability, eye gaze direction, social gestures, etc. in the real world. Off-line analysis of bioelectrical and behavioral signals would be conducted and optimal ways of relaying the information to the robot system and caregiver developed. Integration of multi-modal physiological sensing and visualization of data is essential.

- In 10 years, we should be able to access bioelectrical signals using external hardware instrumentation and have direct analysis of both bio-electrical and movement behaviors to provide detailed diagnosis and/or assessment. Robotic devices are used to stimulate the patient as needed to acquire appropriate data, from the motor to the social. Algorithms for automatically extracting salient behaviors from multi-modal data should enable for data segmentation and analysis, for aiding quantitative diagnosis.

- In 15 years, we can accomplish connecting and easily accessing bioelectrical signals with wearable or implantable devices. This is linked to integrated unencumbered multi-modal sensing and intuitive data visualization environment for the user and caregiver. Real-time algorithms enable not only off-line but also on-line quantitative analysis of such data to inform in situ diagnosis as well as long-term patient tracking. Systems are developed for in-home use and detection of early symptoms of pervasive disorders, such as autism spectrum disorder, from behavioral data.

3.2.6. Context-Appropriate Guidance

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 *human-machine collaborative systems,* in which the operator works "in-the-loop" with the robot during the 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 tasks such as grasping an object with a prosthetic hand, performing a delicate surgical procedure, or assisting an elderly patient to get out of bed.

Many types of assistance, or guidance, can be provided. In prosthesis control, it may be decades before we have sufficient understanding of the human nervous system in order to provide sensory feedback that allows humans to easily control an artificial hand with as many joints as a real hand. Thus, low-level robotic controllers are needed to help automatically control the joints that are not directly controlled by the human. The motion of the automatically controlled joints should be complementary to the human-controlled joints, and the resulting behavior so intuitive that the human operator does not even notice that some autonomy is taking place. Another example is the use of "virtual fixtures" in surgery. The term "virtual fixture" refers to a general class of guidance modes, implemented in software and executed by a robotic device, that help a human-machine collaborative system perform a task by limiting movement into restricted regions and/or influencing movement along desired paths. Virtual fixtures can enhance robot-assisted minimally invasive surgery by ensuring that the manipulator inside the patient does not enter forbidden areas of the workspace, such as organ surfaces that should not be cut and delicate tissue structures. At the same time, the surgeon should be able to override the virtual fixture if desired. A final example of such guidance includes coaching of physical, cognitive and/or social exercises toward rehabilitation of a variety of conditions. Implementing such guidance modes requires that the robot understands the task the human operator or user is trying to do, the current state of the human (both physically and the human's intent), and have the physical and/or social means for providing assistance. The milestones below are based on increasing uncertainty of the task, human operator, and environment.

- In 5 years, a robot should be able to track, record, and suggest optimal procedure performance for set of well-defined procedures or behaviors, with clear steps. Recognition of human behavior/state and corresponding robotic assistance should be achievable in laboratory environment.
- In 10 years, a robot should be able to recognize and classify human behavior and intent achievable in a modified environment in which the environment and/or people are augmented to make perception easier. Novel devices should be used to provide augmentation in an unobtrusive manner.
- In 15 years, the robot should be able to achieve the 10-year performance in an unmodified environment. A robotic system should be able to assemble relevant historical data and consultations with expert caregivers for tricky situations, even bringing them into the control loop if necessary.

3.2.7. Image-Guided Intervention

We now consider robotic image-guided intervention, which concentrates on visualization of the internal structures of a patient in order to guide a robotic device and/or its human operator. This is usually associated with surgery and interventional radiology, although the concepts described here could more broadly apply to any health care needs in which the patient cannot be naturally visualized. No matter the application, such interventions require advances in image acquisition and analysis, development of robots that are compatible with imaging environments, and methods for the robots and their human operators to use the image data.

Novel materials, actuation mechanisms, and sensors are required to create robots that can be seamlessly integrated into the interventional suite.

Sensor data are essential for building models and acquiring real-time information during surgery and interventional radiology. Real-time medical imaging techniques such as magnetic resonance imaging (MRI), ultrasound, spectroscopy, and optical coherence tomography (OCT) can provide significant benefits when they enable the physician to see subsurface structures and/or tissue properties. In addition, images acquired pre-operatively can be used for planning and simulation. New techniques such as elastography, which non-invasively quantifies tissue compliance, are needed in order to provide images that provide useful, quantitative physical information. For robot control, the necessary speed and resolution of imagers is not yet understood. We must determine how to integrate these with robotic systems to provide useful information to the surgeon and the robot to react to patient health in real time.

One of the most useful forms of imaging is magnetic resonance imaging (MRI). The design of MRI-compatible robots is especially challenging because MRI relies on a strong magnetic field and radio frequency (RF) pulses, and so it is not possible to use components that can interfere with, or be susceptible to, these physical effects. This rules out most components used for typical robots, such as electric motors and ferromagnetic materials. In addition, surgery or interventional radiology inside an imager places severe constraints on robot size and geometry, as well as the nature of the clinician-robot interaction. Novel materials, actuation mechanisms, and sensors are required to create robots that can be seamlessly integrated into the interventional suite.

With the abundance of different types of interventions, it is useful to consider milestones that address the different types of surgery that could be accomplished with robotic assistance. Each of these milestones involves the same concepts for semi- and fully automated robot behaviors, only at different levels of complexity.

- In 5 years, we should be able to use images to perform ultra-minimally invasive diagnosis and therapy, using needles that can reach desired targets while avoiding delicate structures. Robots should enable automatic transformation of image data to physical models of specific patients to guide these interventions.
- In 10 years, we should have swimming microrobots capable of local drug delivery, using automatic vessel structure model from spatial imaging. In addition, these robots will need to have imager-compatible locomotion and control design (using fluid mechanics models) and real-time automatic pathology localization from imager-compatible physiological sensing.
- In 15 years, we can achieve semi-automated and automated surgical assistants that use fully real-time image-to-model generation (including geometry, mechanics, and physiological state). The image data should be used to generate on-line planners and control for organ retraction and resection in dexterous minimally invasive surgical procedures.

3.2.8. High-Dexterity Manipulation at Any Scale

Device design and control is key to the operation of all medical and health robotics, since they interact physically with their environment. Accordingly, one of the most important technical challenges is in the area of mechanisms. For example, in surgical applications, the smaller a robot is, the less invasive the procedure is for the patient. And in most procedures, increased dexterity results in more efficient and accurate surgeries. We also consider the possibility of cellular-scale surgery; proofs-of-concept of this have already been implemented in the laboratory. Another example is rehabilitation; current rehabilitation robots are large and relegated to the clinic. Similarly, human physical therapists have limited availability. Yet for many patients, effective long-term therapy clearly calls for longer and more frequent training sessions than is affordable or practical in the clinic. Human-scale wearable devices, or at least ones that can be easily carried home, would allow rehabilitative therapies to be applied in unprecedented ways. Finally, consider a dexterous prosthetic hand. To fully replicate the joints of a real hand, using current mechanisms, actuator designs, and power sources would require the hand to be too heavy or large for a human to naturally use. Small, dexterous mechanisms would make great strides toward more life-like prosthetic limbs.

Miniaturization is challenging in large part because current electromechanical actuators (the standard because of their desirable controllability and power to weight ratio) are relatively large. Biological analogs (e.g., human muscles) are far superior to engineered systems in terms of compactness, energy efficiency, low impedance, and high force output. Interestingly, these biological systems often combine "mechanisms" and "actuation" into an integrated, inseparable system. Novel mechanism design will go hand-in-hand with actuator development. In addition, every actuator/mechanism combination will need to be controlled for it to achieve its full potential behavior, especially when dexterity is required. Models need to be developed in order to optimize control strategies; this may even motivate the design of mechanisms that are especially straightforward to model.

Goals for systems that achieve high dexterity at any scale will naturally differ greatly depending on the medical application (e.g. the surgery, rehabilitation, and prosthetics examples given above). Thus, a natural set of milestones for mechanism design is to consider capabilities linked to each of these applications in order of increasing complexity.

- In 5 years, robotic hands for prostheses should have sufficient degrees of freedom and dexterity with lightweight structure so as to achieve natural manipulation. Mobile manipulators should be available to deal with structured environments (e.g., pick up and deliver specific objects).
- In 10 years, robotic manipulators for surgery should be able to perform snake-like maneuvers at great depth – such as that required for natural orifice surgery. Manipulators for everyday objects should be expanded to handle more general objects and tasks (pick up, deliver, turn know, open door, push button, move slider, etc.).
- In 15 years, micro-scale robots should be able to assist in dexterous microsurgery in small structures such as the eye, as well as cellular-scale surgery. Mobile manipulation with on-board power and computation should manipulate objects in everyday environments safely.

3.2.9. 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. Other sensors are also becoming more effective and more common. Our cell phones include accelerometers, cameras, and GPS, which provide considerable information. 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 physAggregating 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 the 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.iology.

- In 5 years, begin a concerted data collection effort. Begin aggregating existing health data (in appropriately anonymous format) toward facilitating analysis. Work with the various health communities and interested data collection parties to facilitate access to anonymous data. Learn from successful models (e.g., Iceland genetic database).
- In 10 years, apply data mining algorithms to growing body of data. Deploy sophisticated data sharing techniques to facilitate access not only to the research community but also to heath professionals and patients.
- In 15 years, make 15 years worth of health data for a nation and beyond available in anonymous form to all interested researchers, health care professionals, and lay users through a suitable web interface, while continuing to collect data long term and make it available.

3.2.10. Safe Robot Behavior

The challenge of safe robot action and reaction is as old as the field of robotics itself. However, safety takes on a new dimension when directly close-up interactions with human users, often vulnerable ones, constitute the core of the robot's purpose. Providing appropriate response to human behavior (e.g., knowing difference between inadvertent human behavior and specific intent) represents a new technical challenge.

The robot must be able to anticipate dangerous behavior or conditions (i.e., create virtual constraints) and respond to any urgent conditions in home environments under all conditions. Such operation is much more readily achieved in non-contact systems, i.e., HRI that does not involved physical touch and application of force between the user and the robot. When contact is involved, research is focusing on inherently safe mechanisms at the mechanical and hardware level to facilitate safety well before the software level.

Safety of behavior has more profound implications than merely physical interaction. While socially assistive robotics does 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 possibilities must be taken into account in the context of safe system design.

- In 5 years, continue development on inherently safe actuation, low-weight/strength and affordable robot bodies for service and socially assistive robotics for in-clinic and in-home testing for specific tasks.
- In 10 years, create affordable prototypes for in-clinic and in-home robot systems for extensive evaluation with heterogeneous users (health care providers, family, patient). Collect longitudinal data on safety and usability.
- In 15 years, 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.

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 health care delivery is inherently qualitative in nature, having to do with therapy, motivation, and social interaction; while such methods are standard in the social sciences, they are not recognized or accepted by the medical community. Because medical and health robotics must work with both trained specialists and lay users, it is necessary to gain acceptance from both communities. This necessitates reproducibility of experiments, standards, code re-use, hardware platform re-use/sharing, clinical trials, sufficient data for claims of efficacy, and moving robots from lab to real world. As systems become increasingly intelligent and autonomous, it is necessary to develop methods for measuring and evaluating adaptive technologies that change along with the interaction with the user.

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

Resources must be directed toward funding collaborative ventures that bring together the necessary expertise in engineering, health, and business.

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

A thriving industry requires a training in research, implementation, evaluation, and deployment of health care 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 care giving.

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 multi-modal 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 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, and 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 and a plan for how a robot can reach a tumor in an organ. 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

Robot perception, which uses sensor data and models to develop an understanding of a task or 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 health care 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 not distracting 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. Such tasks 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, below.

4.5. Robust, High-Fidelity Sensors

We focus here on two types of sensing especially important for medicine and health care: biocompatible/implantable sensors and force/tactile sensing. 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.

4.6. Novel Mechanisms and High-Performance Actuators

For systems ranging from ultra-minimally invasive surgery robots to human-size prosthetic fingers, 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 humans 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, 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. We need to continue to explore and develop these and other potential robot actuators.

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 statistical machine learning applied to real-world uncertain and multi-modal 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 data for grounding in relevance. In general, learning and adaptation to users, to environments, and to tasks 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 use 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. Socially Interactive Robots

Effective social interaction with a user (or a set of users) is critically important for enabling medical and health robotics to become useful for 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 bi-directional and thus involves both multi-modal 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.

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.

We now provide some specific examples of models needed for medicine and health care. 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 his 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 on-line real-time human state and activity detection, classification, and prediction.

5. Contributors

This document is based on the workshop entitled "A Research Roadmap for Medical and Healthcare Robotics", held June 19-20, 2008 in Arlington, VA. The workshop was sponsored by the Computing Community Consortium (CCC), part of the Computing Research Association (CRA), through a grant from the U.S. National Science Foundation (NSF).

Listed below are 37 researchers and industrial representatives who attended the workshop or otherwise contributed to this document. The workshop and the development of this document were led by Maja Mataric', Allison M. Okamura, and Henrik Christensen.

Ron Alterovitz
UC Berkeley, UCSF

David Brown
Kinea Design

M. Cenk Cavusoglu
Case Western Reserve University

Howie Choset
Carnegie Mellon University

Henrik Christensen
Georgia Tech

Mark Cutkosky
Stanford University

Hari Das Nayar
NASA/JPL

Jaydev Desai
University of Maryland

Aaron Dollar
Harvard/MIT

Aaron Edsinger
Meka Robotics

Miguel Encarnação
Humana

Brian Gerkey
Willow Garage

Neville Hogan
Massachusetts Institute of Technology

Ayanna Howard
Georgia Tech

Robert Howe
Harvard University

Chad Jenkins
Brown University

Dan Jones
Intuitive Surgical

Timothy Judkins
Intelligent Automation

James Koeneman
Kinetic Muscles, Inc.

Venkat Krovi
SUNY Buffalo

Corinna Latham
AnthroTronix, Inc.

Ming Lin
UNC Chapel Hill

Jay Martin
Orthocare

Maja Mataric'
University of Southern California

Allison Okamura
Johns Hopkins University

Marcia O'Malley
Rice University

Charlie Ortiz
SRI

Brian Scassellati
Yale University

Reid Simmons
Carnegie Mellon University

Bill Smart
Washington University in St. Louis

John Spletzer
Lehigh University

Thomas Sugar
Arizona State University

Stewart Tansley
Microsoft Research

Russell Taylor
Johns Hopkins University

Frank Tendick

Chris Ullrich
Immersion

Holly Yanco
UMass Lowell

Chapter 3

A 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 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 38,000 professional robots are in use today and the market is growing rapidly every year. Several typical professional robots are shown in figure 1.

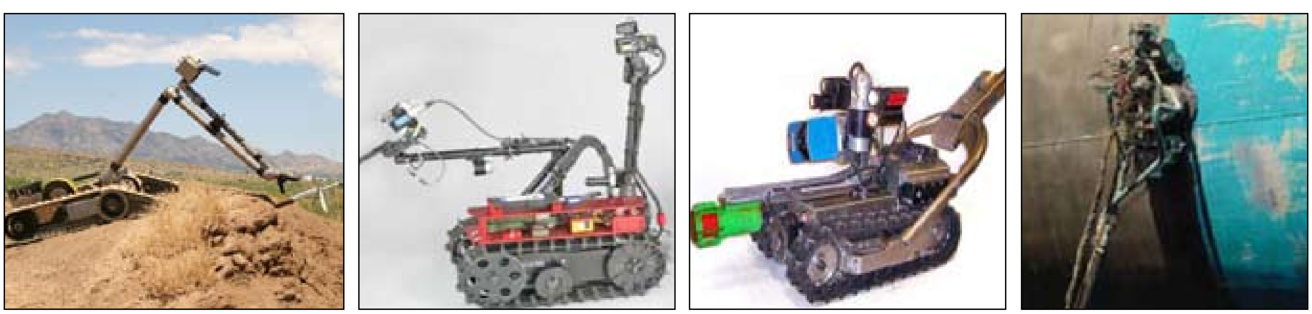

PackBot™ Tactical Robot | TALON® Hazmat Robot | Responder™ Pipeline Robot | Envirobot™ Paint Stripping Robot

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. The by far largest group of personal service robots consists of domestic vacuum cleaners; over 3 million iRobot Roomba's alone have been sold worldwide and the market is growing 60%+ / year. In addition, a large number of robots have been deployed for leisure applications such as artificial pets (AIBO), dolls, etc. With more than 2 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.


Roomba™ Vacuum Cleaning Robot


ATRSTM Robotic Wheelchair System


Verro™ Pool Cleaning Robot


LEGO® Mindstorms™ Educational

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 have the greatest impact. Among the areas identified are healthcare, national infrastructure and resource management, energy and the environment, security, transportation and logistics, and education and entertainment.

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 work force 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 retiree 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.

The United States is on the threshold of a 20-year trend that will see a near doubling of the number of retiree workers as a percentage of the current workforce.

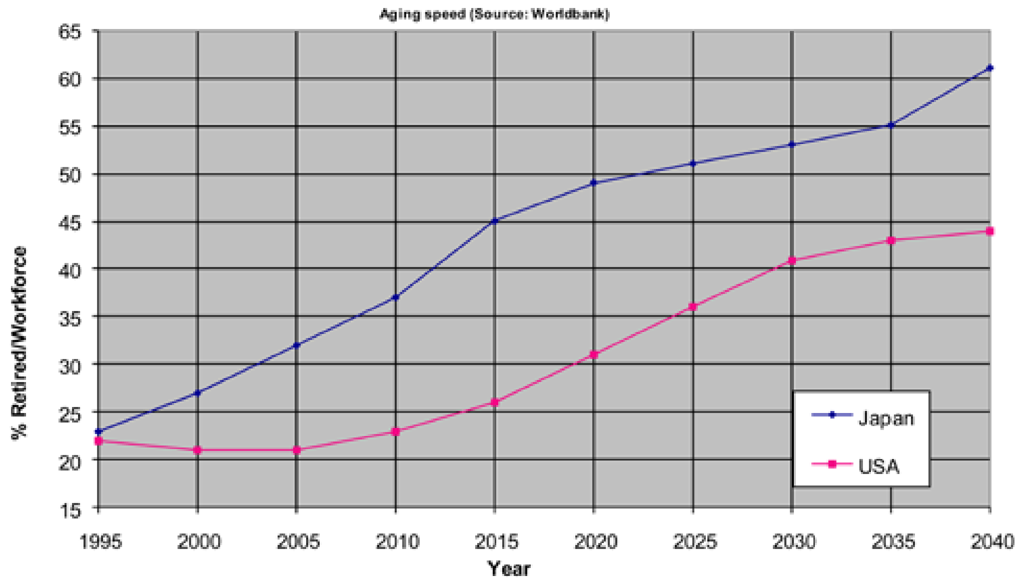

Figure 3. The changes in demographics in USA and Japan respectively.

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 daVinci 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 1999 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 and the engineering and system integration principles required to produce intelligent machines to accomplish certain missions.

2.2. Near-Term Opportunities and Factors Affecting 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 over-ridden by human operators.

Examples of such robotics solutions today include tele-operated systems such as the daVinci 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 three break-out groups to identify technical and scientific challenges pertinent to the applications and business drivers described in the previous section. The first break-out group focused on **application and systems design**; the second group discussed **action, cognition, planning, and other elements of robotic intelligence**; and the final group identified challenges in **human robot interaction**. This section summarizes their findings. Because the challenges identified by the three groups span the boundaries between the respective topic areas, 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 open problems remain. 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 every day, 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. If 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 very high precision laser-based measurement systems 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 develop a new 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 which emphasize planning and control—**perception** becomes an important component of the research agenda towards autonomous manipulation.

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

Substantial progress in manipulation is needed for almost all of the service robotics applications.

Participant 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**, possibly by LIDAR. Advances in dexterous manipulation are likely to require **skin-like tactile sensors for robotic hands**. But also specialized sensors, for example for safety, termed **safety sensors**, 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 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. 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 activities and intent**. Novel algorithms for **affordance recognition** are required to support the type of planning described in the previous subsection. Participants also discussed the need for **accurate sensor models** in support of perceptual 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. Robot 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 change.

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.

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 work space (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.

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 were 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

ATRSTM 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 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 ability to re-install on subsequent vehicles.

3.1.2. Agriculture

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 night time 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.

Autonomous Tractor

3.1.3. Infrastructure

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, and 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 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.

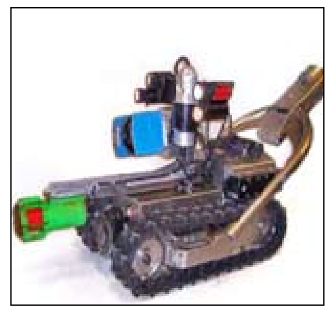

Responder™ Pipeline Robot

3.1.4. Mining

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 2010. 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.

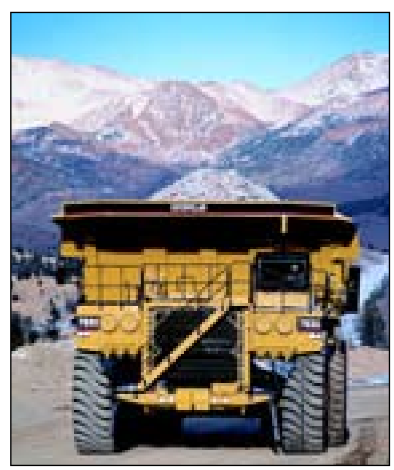

Autonomous Haul Truck

Top three finishers in the 2008 DARPA Urban Grand Challenge

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.

Although American drivers log nearly twice as many miles (1.33 trillion per year) as they did 25 years ago, 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 and 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 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 1999 to inspire young people to be science and technology leaders. As a measure of the growing popularity of robotics competitions, FIRST is expecting over 195,000 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. While other "man-packable" robots equipped with instruments including infrared cameras, night vision sensors and millimeter-wave radar have been used at disaster sites, including the World Trade Center, to search for victims.

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 above. 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 on the next page shows these research areas and technologies in the left column; they are described in more detail in Section 4.

3.2.1. Human-like Dexterous Manipulation

Even simple tasks, such as picking up unknown objects, still represent major research challenges. The level of dexterity and capabilities in physical reasoning required for autonomous manipulation in the context of professional and domestic service robotics seems far out of reach. Pressing problems in this area include adequate sensors and associated perceptual capabilities, dexterous hands and safe manipulators, planning under uncertainty, advanced control, skill learning and transfer, and modeling and simulation.

Some participants believed that the required competency in manipulation can only be achieved when these different areas are advanced in a coordinated fashion rather than in isolation. For example, novel, skin-like tactile sensors hold great promise for dexterous in-hand manipulation. However, we lack the algorithms to process the data from such sensors. It is conceivable that techniques from computer vision could interpret the tactile information as an image and therefore are able to compute useful

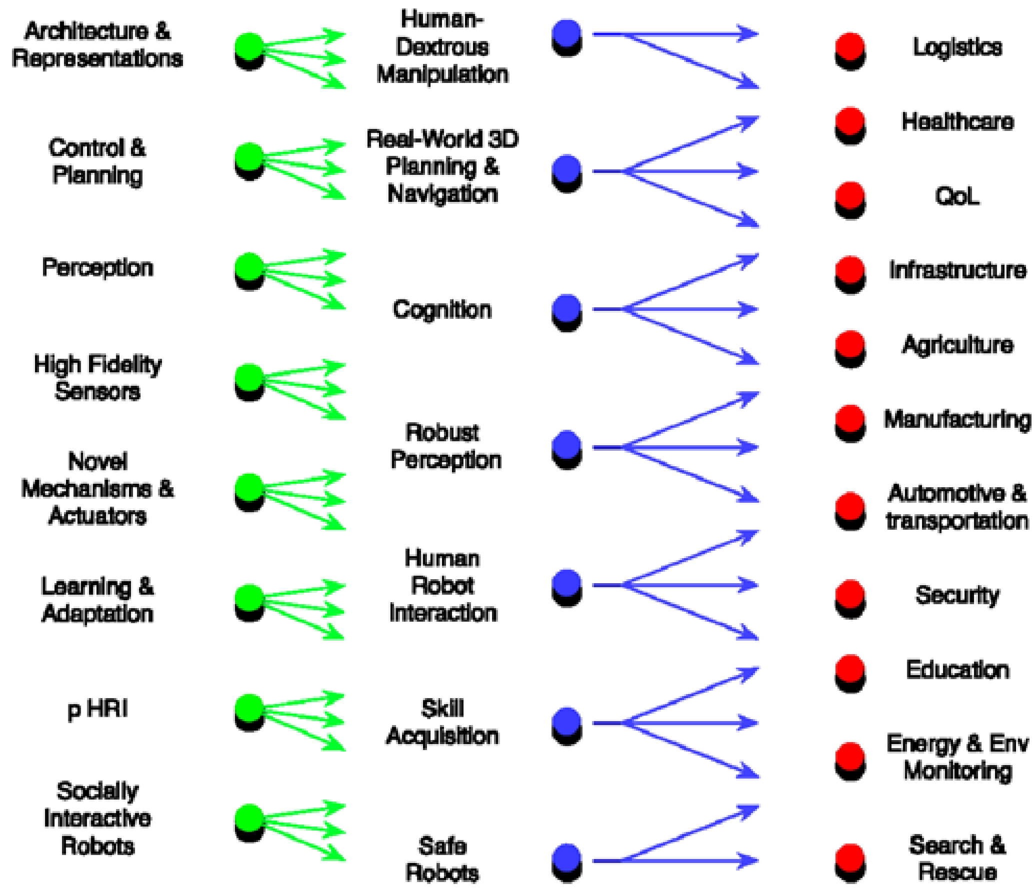

Figure 4. Overview of the roadmap for domestic and industrial service robotics: Sustained research and development in the basic research areas in the rightmost column of the figure will enable a number of elementary capabilities, shown in the middle column of the figure. These capabilities in turn enable progress in the application areas on the right.

abstractions of the high-dimensional tactile data. At the same time, inspiration from computer vision algorithms may enable the design of simpler tactile sensors that contain simple local pre-processing tailored to the specific algorithms they support.

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

5 years: Robots perform limited pick and place task in the home and in industrial settings; robots are able to reliably open doors and cabinets. These manipulation tasks are accomplished partially by engineering the environment, partially by equipping robots with specialized (or at least not very general purpose) end-effectors, and by making simplifying assumptions regarding the environment.

10 years: Robots robustly manipulate large, graspable, rigid, possibly articulated objects and tools without possessing a priori models. Robots improve the robustness and applicability of manipulation and grasping skills with experience. Robots acquire generalized manipulation knowledge to give them information about the use of objects and tools, even if they have not encountered them before.

15 years: Robots possess hands with nearly human levels of mechanical dexterity. Hands are covered with high-resolution tactile skin. Robots are able to perform robust, sensor-based, prehensile and non-prehensile manipulation of objects. They possess rudimentary capabilities of manipulating flexible objects.

3.2.2. Real-World 3D Planning and Navigation

Autonomous service robots accomplish tasks by moving about their environment and by interacting with their environment. These motions and 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, clear obstacles out of its path, remove obstructions, or use tools. To achieve this level of competency, substantial advances at the intersection of motion planning, task planning, and control have to be made. Historically, these areas have progressed in isolation. The problems posed by service robotics, however, can only be addressed through a tight integration of these techniques.

Consider the task of picking up a cup to which access is obstructed by a box. To reason about pushing the box to the side to pick up the cup, the robot has to reason about its own capabilities, the geometry of the scene, constraints imposed by actuation and joint limits, the contact dynamics and friction that arise when pushing the box, etc.

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 and map its environment. "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 and for object manipulation.

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

<u>5 years:</u> Robots in research laboratories can 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 learn semantic maps about their environment through exploration and 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.

<u>10 years:</u> Given an approximate and possibly incomplete model of the static part of the environment (possibly given a priori or obtained from data bases vie 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, interaction, and instruction. The robot modifies its environment to increase the chances of achieving its task (remove obstructions, clear obstacles, turn on lights), and it can detect and recover from some failures.

<u>15 years:</u> Service robots can 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 to dynamic changes in the environment in a way that is consistent with the global objective. 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 between objects and the environments (sliding, pushing, throwing, etc.) and the use of tools and other objects.

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 is 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:

<u>5 years:</u> 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 in-door environment.

<u>10 years:</u> 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.

<u>15 years:</u> 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 relative 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 typically used for recognizing and interacting with single, known objects. To enable scalability there is a need to have facilities for categorization of percepts and generalization across scenes, event and activities. Already today there are methods for mapping and interpretation of scenes and activities and the main challenge is in scalability and robustness for operation in unconstrained environments.

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

<u>5 years:</u> Demonstration of a robot system that can categorize spaces and automatically associate semantics with particular places. The sensing will be integrated over time for robust operation in large scale scales such as mall or a building structure. The robot will be able to recognize hundreds of objects.

<u>10 years:</u> Demonstration of a robot system that can perceive event and activities in the environment to enable it to operate over extended periods of time.

<u>15 years:</u> Demonstration of a robot that integrates multiple sensory modalities such as GPS, vision and inertial to acquire models of the environment and use the models for navigation and interaction with novel objects and events.

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 multi-modal 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. Adequate capabilities in skill learning will be enabled by advances in perception, representation, machine learning, cognition, planning, control, activity recognition, and other related areas.

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 difference 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. There is a need to consider standards for safety to enable certification. 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.

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

5 years: A safety standard for service robotics has been defined and accepted worldwide, which specifies allow impacts and energy transfers. Basic manipulation systems have first versions of safety standard implemented.

10 years: An inherently safe robot for operation in proximity of humans is demonstrated for industrial application scenarios.

15 years: A robot system that does mobile manipulation in cooperation with humans is demonstrated and the safety is demonstrated both for hardware and software components.

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 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. Nonetheless, localization and planning in GPS-denied environments which are quite common remains a very 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 ground-breaking 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, task 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 multi-modal sensory information (e.g., proprioceptive, tactile, vision); structured spatio-temporal 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 here a need 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 brainstorming session that took place 7-8 August 2008 in San Francisco, CA. The report is part of the CCC study on Robotics. The Computing Community Consortium (CCC) is a project managed by the Computing Research Association (CRA) and is sponsored by the National Science Foundation (NSF). The present report has been authored by the workshop organizers and does not reflect the option of CRA, CCC or NSF. The responsibility of the report lies entirely with the authors.

The CCC workshop on service robotics was organized by Oliver Brock, University of Massachusetts, Bill Thomasmeyer, The Technology Collaborative, Inc, and Henrik I Christensen, Georgia Institute of Technology. The workshop was attended by the following people from academia and industry:

Chad Jenkins
Brown

Nicholas Roy
MIT

Aaron Dollar
MIT

Stefano Carpin
UC Merced

Jana Kosecka
George Mason

Andrew Ng
Stanford

Andrea Thomaz
Georgia Tech

Jing Xiao
UNC Charlotte

Charles Rich
WPI

Candace Sidner
WPI

Stewart Tansley
Microsoft Research

Joshua Smith
Intel

Eric Berger
Willow Garage

Martin Buehler
iRobot

Paolo Pirjanian
Evolution Robotics

Bill Townsend
Barrett Technology

Scott Thayer
Red Zone

Chris Urmson
CMU / GM

Cynthia Breazeal
MIT

Michael O'Connor
Novariant

Paul James
Adept

Eric Whinnem
Boeing - Mfr Tech

Charlie Kemp
Georgia Tech

Trevor Blackwell
Anybots

Dan Miller
Anybots

Brian Carlisle
Precise Automation

Parag Batavia
Foster-Miller

Andreas Hoffman
Vecna

James Kuffner
CMU

Alex Foessel
Deere

Oliver Brock
UMass Amherst

Bill Thomasmeyer
Tech Collaborative

Henrik Christensen
Georgia Tech

Jake Huckaby
Georgia Tech

Chapter 4

Robotics: Emerging Technologies and Trends

1. Introduction

Robotics integrates many different component disciplines and technologies, such as computing technology, sensors, materials, actuators, control, and artificial intelligence. Advances in these component technologies have often driven robotics forward, and in return robotics has often provided the applications to motivate advances in the component technologies.

The Workshop on Emerging Technologies and Trends examined continuing and likely future advances in technology, to explore their impact on robotics, and to map the implications for future robotics research directions and funding policy. The Workshop identified 63 different technological advances which promise to impact robotics, and 35 different new applications which will be enabled by various advances. Projections were given for achievable developments in 5, 10, and 15 year time frames. This report presents the workshop's findings, grouped into the following categories:

- actuation systems
- energy and power systems
- fabrication and materials technology
- micro-and nano-technology
- human-robot interfaces
- communications and networking
- planning and control
- robustness and reliability
- machine learning
- perception

Progress in such technologies is essential for capabilities and applications envisioned for robotics, and therefore serves as a foundation for the application-oriented workshops on Manufacturing and Automation, Healthcare and Medical Robotics, and Service Robotics.

2. Strategic Findings

2.1. Actuation Systems

Advances in actuation systems are critical for progress in many applications. For many mobile robot and human-robot interaction applications we need fast, safe, low-cost, and reliable robots. This in turn requires new actuators and transmissions that have high torque-to-weight, power-to-weight, are safe interacting with people, are robust to impact, have appropriate impedance for interactive tasks, and have reasonable speed and efficiency. Current actuators can be made to achieve good performance in a few of these metrics, but not good overall performance. This deficit has led to heavy, slow, fragile robots which are dangerous for people to work with. Just as jet engine technology is a critical strategic advantage for high performance aircraft, actuation technology will provide a competitive advantage in robotics.

Improvements in actuator technology in the coming years won't be dramatic, but steady advances and particularization to robotics will expand the capabilities of robots. The following segments of the world robot market are expected to grow rapidly if necessary actuator technologies are developed.

- *Physical assistance of humans by robots,* such as carrying a person upstairs, requires light weight actuation with appropriate impedance.
- *Micro mobile sensor nodes for security* requires miniature high performance actuators, such as for microflight.
- *Power suits and prosthetics* require high power, high torque, high efficiency actuators.
- *Domestic robots* require safe, low cost, low inertia actuators for mobility plus manipulation.

2.2. Energy and Power Systems

Two key issues, energy storage and power delivery, are both vital to many robotics domains, particularly mobile and autonomous robots. They determine payload, mission duration, and service interval. They are also vital for transportation and consumer electronics, which are much larger markets than robotics and which possess intense research programs. In fact, most progress will come from these areas.

Energy storage involves constraints arising from the fundamental chemistry, as well as numerous economic and safety-related issues. In those applications where it is practical, liquid hydrocarbons will have an energy density advantage for at least the next twenty years. In other applications batteries dominate, and are likely to improve two to three times in the same time period, but still will have ten times lower energy density than gasoline.

Other energy storage technologies have significant issues for practical usage in many applications. Hydrogen requires significant storage overhead. Monopropellants (e.g., hydrogen peroxide) have potential, but they have major safety and handling concerns. Hydrogen peroxide autodenonates at high concentrations (around 80%), but high concentrations are required to get enough energy.

Power delivery is dominated more by design than by chemistry. It is relatively unimportant in batteries for long mission durations. Improved batteries are in the commercial pipeline. There are several companies in the short-term process of commercializing new lithium battery electrode materials that can provide increases in power density and/or cycle life. However, the improved batteries in this category are unlikely to exceed the energy density of current lithium-polymer batteries with cobalt oxide electrodes. There are a group of companies working on batteries that can give 2x-3x improvement in energy density, largely in the area of metal-air batteries, such as rechargeable zinc-air batteries.

Power delivery is vital in fuel-based systems. Engines are already heavily optimized, with clear efficiency trade-offs. Fuel cells should see major improvement in power density, but it is not clear that they beat engines.

In the face of such modest prospects for power and energy, the main option is to explore related technologies with strong applications to robotics.

- *Harvesting.* Requirements for energy storage can be relaxed if energy can be acquired in the field.
- *Efficiency.* Requirements for energy storage can be relaxed if less energy is consumed to perform the task.
- *Miniaturization.* Most robotic systems operate on much smaller scales than transportation vehicles. Some are much smaller than consumer electronics (micro air vehicles, smart dust, etc.) New power and energy systems need to be scaled to robot-relevant sizes. Robotics can also exploit advances of hobbyists. Miniature engines and batteries are well established challenges in remote-controlled aircraft.

2.3. Fabrication and Materials Technology

Advances in fabrication and materials technology are also critical to produce the next generations of robots, by enabling novel lightweight, safe, low cost, compliant, and durable structures. *Integrated fabrication technologies* will result in compact, lightweight subsystems that are rugged and have high performance. *Miniaturization technologies* at the meso/micro/nano scale will lead to proliferation of low-cost components that are easily integrated. *Smart material technologies* will lead to compliant wearable sensors that can either form robot skin or that can be placed on humans for measurement. *New materials technologies* will yield light-weight, soft and safe robot structures. Appplications that would be enabled are:

- *Wearable robots,* e.g., prosthetics, strength assistance for the infirm, and rehabilitation
- *Servant/domestic robots* for elder and other assistance
- *Miniature robots* for safety and security, e.g., search and rescue swarms

2.4. Micro and Nano Technology

Micro and nano robotics will provide unprecedented capabilities to observe and interact with the microworld, from molecules to cells and organs. By engineering robotic systems with characteristic dimensions in the range from nanometers to millimeters, these robots will be able to interact with the environment in ways, and using physical properties, that are not possible with conventional macro-scale systems.

There are two aspects of micro and nano technology: (1) micro and nano manipulation by large robots, and (2) robots that are themselves very small. Micro and nano manipulation by large robots will improve our understanding of materials and biological structures, as well as assist in micro fabrication. The construction of very small robots will yield meso-scale untethered devices for flight, ambulation and swimming. Artificial bacteria and cell systems are also examples of microbots. Small-sized micro-or nano-robots can act in swarms to produce enhanced capabilities.

- The overall human environment can be better observed by using swarms, in such applications as environmental monitoring and search and rescue.
- Swarms of microrobots can perform in vivo medical therapies such as micromanipulation or drug delivery.
- Artificial bacteria can be designed with sensing, manipulation and locomotion capabilities.
- Cell systems can act as miniature factories and as embedded sensors for disease detection.

2.5. Human-Robot Interfaces

Human-robot interaction is central to many of the most exciting applications of robotics, including medical robotics, assistive robotics, prosthetics, rehabilitation, transportation, human augmentation, entertainment and education. Human-robot interfaces include: voice interaction; visual interaction including gesture and inference of intent from visual monitoring; neural interfaces including physical probes, EEG (brainwaves), and surface EMG; physical interaction including exoskeletons and haptic devices; physical manipulation of the human body for transportation or rehabilitation; intelligent prosthetics; and invasive devices such as some biomonitors, surgical instruments, and neural interfaces. A review of this list shows that human-robot interfaces are central to the applications that will affect us most profoundly.

Advances in human robot interaction are coming rapidly and have already had an impact. The Nintendo Wii illustrates that a relatively small advance in gesture recognition can have a significant impact on the human experience. Robotic systems are already using behavior recognition to automate lab science with animal tests. Likewise teleoperated surgery has proven to be feasible and commercially successful. Neural interfaces have made striking progress in the last few years, even progressing to a first commercial system. Other applications on the horizon include:

- factory robots and military logistics robots that learn assembly and warehousing tasks by observing and assisting humans without programming;
- domestic service robots that learn specific non-repetitive tasks by observing humans, inferring how to help, and modifying their assistance through speech and gesture of the human;
- a team of search and rescue robots, working with humans, that can deploy themselves and autonomously negotiate to relay disparate information from several sites to human users such that the number of human operators is less than the number of robots;
- exoskeleton “man-amplifiers” for military and commercial use that slip over portions of the body or the entire body and intuitively enhance or replace human speed/strength/dexterity by responding to nerve impulses or other forms of human “thought”.

Related technologies include speech interaction, gesture and behavior recognition, biocompatible materials, neuroscience, and bioengineering.

2.6. Communications and Networking

Communications and networking are fundamental enabling technologies for numerous high-priority applications. Communications and networking are required whenever robotic systems are distributed spatially, whenever remote data or computing resources must be accessed, or whenever human interaction is required. Advances in communications and networking will lead to more capable, more robust, and more easily deployable systems. Applications enabled by advances in communications and networking include:

- distributed mobile sensor networks, for environmental monitoring, monitoring the civil infrastructure, monitoring container shipping yards, etc.;
- autonomous passenger vehicles, which need to communicate with each other, with traffic control systems, and with offboard data and computing resources such as navigation assistance;
- domestic robotic systems, such as home security systems, home health care and assistive robotics, and home automation;
- distributed micro-or nano-systems deployed for biological instrumentation, such as intracellular scientific studies.

2.7. Planning and Control

Planning and control are the decision sciences employed to determine what actions a robot will take. They include some of the most profound challenges in robotics. Even with perfect sensing and hardware, it is clear that robotic planning and control fall well short of human performance in most tasks. With the current level of planning and control algorithms, robots often have to be employed in narrowly prescribed scenarios, following very detailed programming written laboriously by humans. In other instances, autonomy is impossible, and a human has to be integrated as part of a telerobotic system. Applications enabled by advances in planning and control will include:

- factory robots adapting to a new task without new programming;
- currently teleoperated tasks such as surgery and bomb disposal will transition to supervised autonomy;
- logistic automation will allow robots to autonomously move goods from producers through distribution network to consumers;
- autonomous control of high speed vehicles in dynamic environments will enable passenger vehicle autopilots;
- autonomous robots capable of working in novel environments without reprogramming will make domestic robotics affordable.

2.8. Robustness and Reliability

Robustness and reliability ensure that a robot will continue to do its job, even when the operating environment departs from the ideal. Failures are expensive and inconvenient, and in some applications are unacceptable. Improved robustness and reliability will enable many applications:

- When failure is unacceptable: robotic surgery, passenger vehicle autopilots, critical safety and security applications.
- When failure is expensive and inconvenient: everything else.

2.9. Perception and Machine Learning

Perception and machine learning are the techniques for using sensory information, either to recognize the state of the world around the robot (perception) or to improve its ability to do a job (machine learning). It is largely perception and machine learning that distinguish a robot from an ordinary machine. With perception and machine learning a robot can adjust its actions to the situation at hand and can improve with experience, both hallmarks of intelligent behavior. Advances in perception and machine learning enable new applications:

- Factory robots can perform a new job without re-programming, and possibly by watching a human or another robot doing the job;
- Installation of robots in all applications can be done without expensive special purpose programming, by watching humans or by being instructed by a human.
- Very challenging dynamic tasks like walking also benefit from machine learning techniques.

3. Key Challenges / Capabilities

3.1. Motivating/Exemplar Scenarios

3.1.1. Security - National, Industrial, Home

Security is a broad set of applications. In national security we include military operations such as combat, reconnaissance, surveillance, and logistics support, and homeland security such as border security, airport security, and other critical infrastructure installations. Robotics technology has already had a substantial impact on national security, through the use of unmanned aerial vehicles in military operations, and teleoperated bomb disposal robots. Autonomous and semi-autonomous ground vehicles are just beginning to be deployed. Video surveillance technology is also being deployed.

Industrial and home security will more likely be focused on surveillance. Widespread deployment in homes will be dependent on easily installed highly affordable systems

Several emerging technologies bear directly on security applications. Increased mobility and improved perception, intelligence, and communications would result in more capable military robots, especially in reconnaissance and surveillance. Sensor networks, either mobile or fixed, with improved perception and communications, would more effectively monitor security zones, including borders, airports, marine shipping terminals, industrial security zones, and civil infrastructure facilities. All security applications benefit from advances in sensor networks, perception software, and communications networks, home security especially.

3.1.2. Manufacturing and Distribution

Robotics and automation are well entrenched in manufacturing, most notably in automobile assembly and consumer electronics. Robotics technologies are also being deployed in the distribution chain, from fully automated shipping yards and warehouses to inventory monitoring and point of sale systems at retail outlets. The advent of more capable and more mature technologies could have a broad impact in both manufacturing and distribution. Improved technologies for power, actuation, and materials will lead to more capable and safer systems. Improvements in communications, perception, and human robot interaction will lead to more easily deployed, more affordable, and more adaptable systems. Combining all these technologies with new programming techniques, protocols, and standards will change the underlying economics, so that the technology could be adopted more broadly. Early in this scenario we would see increased adoption by large manufacturing companies, but later on we would see adoption by smaller businesses, ultimately including very small shops.

3.1.3. Transportation

The main focus of transportation is passenger vehicles on streets and highways. The benefits of passenger vehicle automation are enormous in cost, time, comfort, and most of all safety. Early adopters may include closed private systems such as mining or logging companies. Some of the technologies are already fairly mature, and deployed in automated shipping yards for example. Technology for urban driving in a closed environment was demonstrated in the recent DARPA Urban Challenge, making heavy use of laser ranging, radar, and GPS. Many challenges remain to produce a reliable, robust system which can safely cope with all the uncertainties of normal driving. Nonetheless, the technology will gradually shift into our vehicles. Collision warning, skid correction, navigation assistance, and adaptive cruise control are all examples of technology that is already being deployed.

Transportation is so broad that every robotics-related technology is relevant. The most obvious are perception, mobility, and navigation but advances in human robot interaction, networking and communications, actuation, power, control, and machine learning are all necessary.

3.1.4. Medical and Health Care

Medical and health care applications include surgery, rehabilitation therapy, prosthetics and orthotics, medical imaging, monitoring and therapeutic assistance. Robotics technology will impact medical and health care in many ways, providing new tools and techniques for professionals, as well as enabling individuals to monitor and cope with their own health more effectively. Robotics technology has been employed both in medical imaging and computer aided surgery for many years now, with the technology continuing to advance rapidly. New tools and techniques are being developed and demonstrated that will enable procedures that were previously impossible, or costly and dangerous. Neural prosthetics are being demonstrated in case studies, and limb prosthetics are advancing rapidly. Techniques for robotic monitoring of therapy at home are also being demonstrated. Ultimately health care technology can be embedded in our homes, and monitor our health and nutrition and behavioral anomalies. While almost every robotic technology is relevant to health care, especially relevant areas include micro and nanotechnology, power, actuation, and physical human-robot interfaces.

3.1.5. Domestic Robotics

Domestic robotics means use of robotics technology in the home, although some of the most important applications, security and healthcare, are covered in other applications. The most important remaining applications would be education, entertainment, cleaning, and communications. There are already notable successes in entertainment and cleaning, with a few million vacuum-cleaning robots in service. Affordability, ease of installation, and ease of use are key factors in domestic impact. Key technologies are actuators, power, human interaction, communications and networking, perception, safety, robustness and reliability.

3.1.6. Science and Technology

In some instances robotics technologies enable new or improved techniques in other fields of science and technology. There are many examples: micro and nanotechnology can enable superior instrumentation of biological systems, even down to the cellular level. Laboratory automation procedures enable large scale experimentation. Behavior recognition algorithms support previously impossible or expensive animal studies. Improved perception has the potential for broad impact across virtually all parts of scientific study. Several cases are already demonstrated involving lab mice studies for drug discovery, behavioral studies of insects, and even astronomical studies.

Another interesting trend arises from a fundamental connection between biology and robotics. Both fields address the problems faced by an agent which survives and achieves some behavioral goals in the real world. In robotics, the application of biological observations to the design of robots is called biomimetics. There is also some flow of ideas and techniques from robotics to biology.

3.2. Capabilities Roadmap

3.2.1. Actuation Systems

The impact of actuation system developments is sketched for two robotics systems. Ornithopters are flapping-wing robots, for example, bird-like (50 cm) or insect-like (5 cm). The capabilities envisioned are full roll/pitch/yaw/thrust/lift control authority, and a 30-minute mission duration.

<u>5 years:</u> 50 cm ornithopter -cargo capacity equal to actuator system mass. 5 cm ornithopter -cargo capacity equal to 1/10 actuator system mass

<u>10 years:</u> 50 cm ornithopter -cargo capacity equal to 10X actuator system mass. 5 cm ornithopter -cargo capacity equal to actuator system mass

<u>15 years:</u> 50 cm ornithopter -cargo capacity equal to actuator system mass, 30 day mission duration. 5 cm ornithopter -cargo capacity equal to 10X actuator system mass.

Powered exoskeletons for motion assistance are envisioned with an 8 hour mission duration (work shift). A key development are actuators with passive and/or variable impedance.

<u>5 years:</u> Cargo capacity equal to exoskeleton mass.

<u>10 years:</u> Cargo capacity equal to 10X exoskeleton mass.

<u>15 years:</u> Cargo capacity equal to 10X exoskeleton mass, load carried in arms.

3.2.2. Energy and Power Systems

Harvesting. There are two issues in harvesting: mechanisms for energy location and acquisition; and algorithms for power management and energy-seeking behaviors. Harvesting must be well-matched to consumption: it is pointless to harvest milliwatts of vibration energy in a robot that consumes several watts.

Efficiency. Efficiency is mostly a mechanism problem: improved actuators and valves, energy recovery (elastic and/or regenerative), and algorithms for power management. Improvements here also have application to industrial automation, a larger market.

Miniaturization. Miniaturization is often simply an issue of repackaging, but sometimes scaling laws dominate.

3.2.3. Nano and Micro Technology

5 years

- Sub-mm untethered devices capable of in vivo targeted delivery and sensing demonstrated in animal models.
- Robotic nanomanipulation of sub-100nm non-spherical objects with precise 5DOF control in fluid, including force servoed nanomanipulation of high molecular weight molecules.
- High throughput cell assays with relevant industrial applications.
- Wired instrumented cell systems, i.e., networks of nanosensors acquiring and processing in real time data in multiple modalities and multiple temporal and spatial scales, for biological investigations ("understanding biology")
- Sub-cm untethered devices capable of flight, ambulation or swimming in outdoor environments.

10 years

- Increasingly complex in vivo sensing and manipulation therapies by sub-mm micro-robots.
- Tracking and precise control of sub-100nm non-spherical objects in vivo.
- Intracellular nanorobots for biological investigation.
- Artificial bacteria: sensing and locomotion.
- Molecular assays based on robotic nanomanipulation, including investigation of Force-structurefunction relationships in proteins using manipulation.
- Instrumented cell systems, but wireless and embedded in tissue (possibly explants), for applications also in disease detection.
- Artificial bacteria-like robots powered from external fields.
- Sub-cm untethered devices capable of all-terrain mobility, and working in coordinated swarms for tasks such as disaster search and rescue.

15 years

- In vivo self-assembling microrobots.
- Standardized microrobot platform for delivering a variety of in vivo medical therapies.
- Artificial bacteria: sensing, manipulation, and locomotion.
- Experimental platform for drug development based on nanomanipulation of molecules.
- Instrumented cell systems, but having actuator capabilities (e.g to move around, to release drugs, or to mechanically interact with cells) and in vivo.
- Artificial autonomous bacteria.
- Swarms of untethered sub-cm devices capable of large area environmental monitoring and control such as crop pollination, health monitoring, and pest protection.

3.2.4. Human-Robot Interfaces

5 years

- Robust control of exoskeletons using surface electromyography embedded in links.
- Reliable gesture/voice command of service robots based on limited taxonomies with few-to-no
- training iterations.
- Shared control of surgical procedures using real-time sensory feedback and imaging.
- Commercially available, low cost (video game-ish), PHANToM-like haptic device.
- Adaptable standards and practices for intuitive telepresence interfaces for one-to-one control.

10 years

- Neurological control of exoskeletons using neuro-prosthetics and advanced surface EMG.
- Inference of complex intention from natural gesture/voice interaction with humans.
- Remote control of surgical procedures over limited networks using real-time sensory feedback and imaging with predictive correction and pacing of the procedures.
- Adaptable standards and practices for intuitive telepresence interfaces for 1:N (human:robot) control of remote robots.
- Adaptable standards and practices for safe interaction of humans with robots based on task and sensory feedback.

15 years

- Neurological control of exoskeletons using brainwaves (EEG).
- Invasive neurological control that is widely acceptable to potential users.
- Seamless cooperation interfaces using gesture, voice, and communication-through-the-task.
- Adaptable standards and practices for intuitive, self-arbitrating interfaces for many humans controlling many remote robots.

3.2.5. Communications

5 years

- Protocols for sparse, highly volatile multi-hop, ad-hoc networks with high bandwidth and low latency.
- Improvements in localization in UWB networks and better spectrum utilization. Integration of wide-area, local-area, and personal-area networks for more seamless local-toglobal coverage (heterogeneity).
- Cooperative communication clusters in ad-hoc networks for better connectivity, power utilization, etc.
- Integration of the cellular and satellite phone networks to augment other networks.

10 years

- Cognitive Radio -extension of theory and development of deployable systems.
- Improvements in meta-level methods and protocols (middleware) for heterogeneous agents for task allocation, resource sharing, etc.
- General methods and theories encompassing non-traditional communication (biological forms as opposed to RF).
- Communication through the task.
- Local collaboration strategies that better deal with poor network connectivity in volatile environments (all environments).

15 years

- Expanded use of commercial multimedia standards (video object layer, content keying, etc).
- Incorporation of cognitive radio standards with multi-agent collaboration/cooperation strategies.
- Robust middleware approaches for generic task allocation, resource distribution and data sharing across multiple network modalities with frequent outages.

4. Research/Technologies

4.1. Actuation Systems

The following are projections for actuation developments in 5, 10, and 15 year time frames.

5 years: higher torque and power to weight than muscle (> 20 Nm/kg and > 50 W/kg) with appropriate impedance in 1 gram to 1 kg size-scale.

10 years: 5X improvement in torque and power to weight with appropriate impedance, in milligram to kg size-scale, sufficient bandwidth, power plant efficiency > 40%, minimal holding power, low complexity control.

15 years: high performance actuators for light weight/safe mobile manipulation hierarchical for precision/power tradeoff and fault tolerance alternate power sources (food rather than batteries?)

4.2. Energy and Power Systems

The following are projections for energy and power developments in the 5, 10, and 15 year time frames.

5 years

- **Harvesting.** Develop a 10 kg mobile robot capable of fully autonomous recharging in an office environment.
- **Efficiency**. Construct a robust legged locomotion system that consumes 2x power compared to a similar biological system.

- **Miniaturization.** Fabricate a 100 mg battery exceeding 500 kJ/kg and 1 kW/kg energy storage and power delivery capability.

10 years

- **Harvesting.** Develop a 10 kg mobile robot capable of fully autonomous recharging and power management outdoors (> 25% duty cycle). Develop a 1 gram UAV capable of operation at a 5% duty cycle while harvesting in an indoor/office environment.
- **Efficiency.** Improve the legged locomotion system to equal power to biological model.
- **Miniaturization.** Produce a fuel-based power source with converter mass < 50 mg, power > 100 W/kg, and efficiency > 10%.

15 years

- **Harvesting.** Develop a 1 gram UAV capable of operation at a 25% duty cycle while harvesting outdoors.
- **Efficiency.** Improve the legged locomotion system to 50% power of the biological model.
- **Miniaturization.** Produce a fuel-based power source with converter mass < 50 mg, power > 1 kW/kg, and efficiency > 10%.

4.3. Fabrication and Materials Technology

Roadmaps are organized by four areas, addressing several key issues: cost; throughput; heterogeneous integration; complexity (2d or 3d); compliance; strength-to-weight; time-to-market; biodegradable/recyclable/green; and modularity.

4.3.1. Integrated Fabrication Roadmap

Enabled technologies include low cost, consumer robots, high performance (high power and low weight), high reliability, and protection of US intellectual properties.

5 years: integrated sensing, mechanics, actuation, wiring, packaging (encapsulating key components)

10 years: a unified programmable fabrication process (e.g. printing, deposition, etc.) of heterogeneous components with no manual assembly

15 years: Y10 + high throughput

4.3.2. Miniaturization Roadmap

Enabled technologies include security, healthcare, and mobile sensor networks for search and rescue.

5 years: miniaturized integrated sensing, mechanics, actuation, wiring, comm, packaging (encapsulating key components); Combined meso/micro/nano components and systems.

10 years: high throughput

15 years: Y10 + low cost + use of micro-and nano-technology based intelligent materials

4.3.3. Smart Material Roadmap

Enabled technologies include power suits and prosthetics.

5 years: compliant wearable sensor skin

10 years: compliant wearable sensor and actuator skin customized to fit the individual wearer's skeletal structure

15 years: compliant wearable sensor and actuator skin with integrated power/energy; modular, mass-produceable smart-material actuators co-fabricated with the skin structure

4.3.4. Materials Roadmap

Enabled technologies include higher performance, safer, more agile robots.

5 years: light-weight, soft and safe robot structures (e.g. humanoid with mass of 20 kg for light duty household tasks); new materials with controllable adhesion/friction properties for agile and high efficiency mobility and fragile part manipulation.

10 years: embedded sensors and actuators in soft materials for robot limbs and bodies; controllable softness; highly maneuverable flapping wing based centimeter scale flying robots.

15 years: both low-cost mass production and customizable fabrication of soft robots, from meso to nanoscale, with recyclable materials.

4.4. Planning and Control

5 Years

- A new generation of filtering algorithms that exploit task constraints to minimize sensing, computation, and estimation.
- Robust, publicly available libraries to compute optimal feedback plans in several (3 to 6) dimensions.
- Rapid replanning approaches based on precomputation and low-cost memory, providing an alternative to explicitly stored feedback plans.
- Development of mathematical criteria for predicting the quality of motion primitives or behaviors in planning algorithms.

10 Years

- Sensor-centric planning and control algorithms that operate directly in the reduced information spaces resulting from minimalist filters.
- Development of sampling-based techniques that compute feasible feedback plans or control laws in high-dimensional spaces.
- Unified the notions of real-time replanning and feedback control, including mathematical analysis of convergence properties.
- Automatic synthesis of motion primitives based on optimizing newly introduced criteria.

15 Years

- Optimal feedback motion planning for nonlinear systems with obstacles in ten or more dimensions.
- Unified theory and techniques for reducing the topological, combinatorial, and dimensional complexities of information spaces.
- Algorithms that eliminate the artificial boundaries between sensing, planning, control, and learning.
- Dramatic performance improvements in applications that use planning and control due to automatically synthesized primitives.

5. Contributors

This report resulted from a Workshop on Emerging Technologies and Trends that took place 14–15 August 2008 at Snobird, Utah. The report is part of the Computing Community Consortium (CCC) study on Robotics. The CCC is managed by the Computing Research Association (CRA) and is sponsored by the National Science Foundation (NSF). The present report has been authored by the workshop organizers and does not reflect the option of CRA, CCC or NSF. The responsibility of the report lies entirely with the authors.

The CCC Workshop on Emerging Technologies and Trends was organized by John M. Hollerbach, University of Utah, Matthew T. Mason, Carnegie Mellon University, and Henrik I. Christensen, Georgia Institute of Technology. The workshop was attended by the following people from academia and industry:

H. Harry Asada
Massachusetts Institute of Technology

Michael S. Branicky
Case Western Reserve University

Craig R. Carignan
Georgetown University

Henrik I. Christensen
Georgia Institute of Technology

Ronald S. Fearing
University of California at Berkeley

William R. Hamel
University of Tennessee

John M. Hollerbach
University of Utah

Steven M. LaValle
University of Illinois

Matthew T. Mason
Carnegie Mellon University

Bradley J. Nelson
ETH Zrich

Gill Pratt
Olin College

Aristides A. G. Requicha
University of Southern California

Bryan P. Ruddy
Massachusetts Institute of Technology

Metin Sitti
Carnegie Mellon University

Gaurav S. Sukhatme
Univerity of Southern California

RussTedrake
Massachusetts Institute of Technology

Richard M. Voyles
University of Denver

Mingjun Zhang
University of Tennessee