

A Research Roadmap for Medical and Healthcare Robotics

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.

Table of Contents

1	Introduction.....	3
1.1	Definition of the Field/Domain.....	3
1.2	Societal Drivers.....	4
2	Strategic Findings	6
2.1	Surgical and Interventional Robotics.....	6
2.2	Robotic Replacement of Diminished/Lost Function	6
2.3	Robot-Assisted Recovery and Rehabilitation	7
2.4	Behavioral Therapy.....	8
2.5	Personalized Care for Special Needs Populations	9
2.6	Wellness/health promotion	10
3	Key Challenges and Capabilities	11
3.1	Motivating Exemplar Scenarios.....	11
3.1.1	Surgery and Intervention	11
3.1.2	Replacement of Diminished/Lost Function	11
3.1.3	Recovery and Rehabilitation.....	11
3.1.4	Behavioral Therapy.....	11
3.1.5	Personalized Care for Special Needs Populations	12
3.1.6	Wellness and Health Promotion	12
3.2	Capabilities Roadmap	12
3.2.1	Intuitive physical human-robot interaction and interfaces	12
3.2.2	Automated understanding of human behavior.....	13
3.2.3	Automated understanding of emotional and physiological state	14
3.2.4	Long term adaptation to user's changing needs	15
3.2.5	Quantitative diagnosis and assessment	16
3.2.6	Context-appropriate guidance.....	18
3.2.7	Image-guided intervention.....	19
3.2.8	High dexterity manipulation at any scale	20
3.2.9	Sensor-based automated health data acquisition	21
3.2.10	Safe robot behavior.....	21
3.3	Deployment Issues.....	22
4	Basic Research/Technologies	23
4.1	Architecture and Representations	23
4.2	Formal Methods.....	24
4.3	Control and Planning	24
4.4	Perception	24
4.5	Robust, High-Fidelity Sensors.....	25
4.6	Novel Mechanisms and High-Performance Actuators	26
4.7	Learning and Adaptation	26
4.8	Physical Human-Robot Interaction.....	27
4.9	Socially Interactive Robots.....	27
4.10	Modeling, Simulation, and Analysis.....	27
5	Contributors	28

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 in worsening of subsequent health problems. Access to health care is most directly related to its affordability. Access to physically interactive therapy robots promise to reduce the cost of clinical rehabilitative care and are the focus of an ongoing Veteran's Administration study of their cost-effectiveness. Socially assistive robotics efforts are working toward methods that could provide affordable in-home technologies for motivating and coaching exercise for both prevention and rehabilitation. It is also a promising domain for technologies for care taking for the elderly, toward promoting 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 mis-medication, 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 envelops 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.

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 regimens 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 health 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 regimens 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.

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 is 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.

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.

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.

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 knob, 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 physiology.

Aggregating over the entire population, we will have a database vastly more detailed and broader in scope than anything we have seen in the past. Such a database enables a new level of medical research based entirely on historical data. At present, medical studies are targeted to address specific issues or hypotheses, and the cost of these studies restricts 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.

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 health 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 involve 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.

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