

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 21-10-2009		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 5-Jul-2005 - 31-Mar-2009	
4. TITLE AND SUBTITLE Exploiting Discrete Structure for Learning On-Line in Distributed Robot Systems			5a. CONTRACT NUMBER W911NF-05-1-0396		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Roderic A. Grupen			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Massachusetts - Amherst Office of Grant & Contract Admin. 408 Goodell Building Amherst, MA 01003 -3285			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 45587-CS.1		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Over the course of this award, we have introduced techniques for adaptive control that exploit discrete structure in the organization of distributed and embedded systems. This general approach was evaluated using a variety of physical platforms including a distributed sensor array, a team of multiple mobile platforms, a bimanual humanoid and a new concept in mobile manipulators. The findings suggest that this approach has some advantages with respect to learning performance, fault tolerance, generalization, transfer, and programming by demonstration. In					
15. SUBJECT TERMS discrete-event structure, distributed robotics, on-line learning					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Roderic Grupen
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 413-545-3280

Report Title

Exploiting Discrete Structure for Learning On-Line in Distributed Robot Systems

ABSTRACT

Over the course of this award, we have introduced techniques for adaptive control that exploit discrete structure in the organization of distributed and embedded systems. This general approach was evaluated using a variety of physical platforms including a distributed sensor array, a team of multiple mobile platforms, a bimanual humanoid and a new concept in mobile manipulators. The findings suggest that this approach has some advantages with respect to learning performance, fault tolerance, generalization, transfer, and programming by demonstration. In particular, our cumulative publications over the award period contribute (1) a representation for creating a comprehensive array of closed-loop control circuits, (2) a discrete-event framework for characterizing the state of the dynamical system, (3) and model checking scheme for guaranteeing performance and safety while learning that significantly improves the performance of machine learning techniques in robotics, (4) an intrinsically motivated reinforcement learning technique based on the discrete-event representation, (5) mechanisms abstraction in discrete event systems that enhances transfer and provides a means of learning from demonstrations and imitation.

This project has contributed to training for one PhD student (Robert Platt - NASA and MIT), one Masters student (Edward Hannigan - digitROBOTICS and Barrett Technology), and has provided research experiences for two NSF REU candidates (Rachel Lee and Brian Russell). Finally, many of the original contributions of this work are represented in an upcoming book entitled "The Developmental Organization of Dexterous Robot Behavior," that will be published by MIT Press.

List of papers submitted or published that acknowledge ARO support during this reporting period. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Number of Papers published in peer-reviewed journals: 0.00

(b) Papers published in non-peer-reviewed journals or in conference proceedings (N/A for none)

Number of Papers published in non peer-reviewed journals: 0.00

(c) Presentations

RSS09 Workshop on Bridging the Gap between High-Level Discrete Representations and Low-Level Continuous Behavior, July 2009

The Developmental Organization of Robot Behavior - Johns Hopkins University, Department of Computer Science, 2009.

Personal Robotics - IEEE Robotics Society, Boston Chapter, January, 2009.

Whole-Body Mobile Manipulation - Brief, ARL Cooperative Technology Alliance, Reston VA, May 8, 2008.

Personal Robotics for Healthcare Applications - Invited Participant, Massachusetts Technology Leadership Council Robotics Visioning Meeting, exclusive meeting for CEOs, research directors, industry leaders, 2008, Boston.

Development and Distributed Robot Behavior, invited speaker, Microsoft Research Seminar Series, 2008.

Dexterous Mobile Manipulators - Control, Programming, Application - iRobot Corporation, 2007.

The Developmental Organization of Dexterous Robot Behavior, invited speaker, Georgia Tech Robotics Seminar Series, January 2007.

Mobile Manipulation - invited speaker, Vanderbilt University, 2005.

Number of Presentations: 9.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts): 0

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Ou, S., and Grupen, R. Learning Prospective Robot Behavior, in AAAI Spring Symposium 2009. Stanford University, Palo Alto. March, 2009.

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts): 1

(d) Manuscripts

Number of Manuscripts: 0.00

Number of Inventions:

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Shichao Ou	0.50
Shiraj Sen	0.25
FTE Equivalent:	0.75
Total Number:	2

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Roderic A. Grupen	0.08	No
FTE Equivalent:	0.08	
Total Number:	1	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Rachel Lee	0.00
FTE Equivalent:	0.00
Total Number:	1

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

- The number of undergraduates funded by this agreement who graduated during this period: 0.00
- The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00
- The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00
- Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00
- Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00
- The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00
- The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Total Number:

Names of personnel receiving PHDs

<u>NAME</u>
Total Number:

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

FINAL Report for Exploiting Discrete Structure for Learning On-Line in Distributed Robot Systems

**NSF/NASA Workshop on
Autonomous Mobile Manipulation (AMM)**

Houston, Texas, March 10/11, 2005

Final Report

August 8, 2005

prepared by:

Oliver Brock and Roderic Grupen
Laboratory for Perceptual Robotics
University of Massachusetts Amherst

Executive Summary

Research in Autonomous Mobile Manipulation (AMM) aims to create robotic agents capable of performing physical work in unstructured and open environments. Such technology will impact a variety of application areas with significant economical, societal, and scientific importance. Among them are assistive and elderly care, planetary exploration, manufacturing, and supply chain management.

Many countries have initiated well-funded and focused research programs in Autonomous Mobile Manipulation or closely related areas. In the United States, such a coordinated research initiative is still lacking. This workshop gathered some of the leading researchers in robotics, computer vision, and related fields, to devise a set of recommendations towards the initiation of such a program. These recommendations include:

- ▶ The specification of a general research program that builds on the existing strengths of research and technology in the United States and positions the academic community competitively, relative to comparable initiatives in other countries.
- ▶ The creation of center-level funding opportunities for multi-disciplinary initiatives in AMM. This is necessary to address the broad range of research problems that arise in this area in an integrated manner. Furthermore, it is consistent with initiatives in other countries.
- ▶ A proposal for three concentrated research thrusts within the NSF Robotics and Computer Vision programs. These thrusts focus on specific research topics that are viewed as elementary building blocks for an initiative in AMM. They include a program for the development of basic skills (in particular dexterous manipulation skills), a program for the development of integrated, multi-modal sensor strategies, and a third program for the development of integrative architectures to facilitate the robust and flexible operation of robotic agents in the context of AMM.

Contents

Executive Summary	1
Preface	4
1 Introduction	4
2 State of the Art Worldwide	5
3 Enabling Technologies	6
3.1 Hardware	7
3.2 Standardization of Sensors	7
3.3 Wiring	7
3.4 Embedded, Distributed Processing	8
3.5 Actuation	8
3.6 Power	8
4 Scientific Challenges	8
4.1 Mobility	9
4.2 Representing Objects and Environments	9
4.3 Grasping and Dexterous Manipulation	9
4.4 Perception	10
4.5 Architectures	11
4.6 Human-Robot Interaction	12
4.7 Safe Manipulators	12
4.8 Commodity Hands	12
5 Recommendations	13
5.1 Areas of Focus	13
5.2 Cross-Fertilization	13
5.3 Recommended Research Initiatives	14
5.4 Infrastructure and Programmatic Recommendations	15
5.5 Hardware Infrastructure	15
5.6 Software Infrastructure	15
5.7 Intermediate Goals	16
6 Acknowledgments	17
References	18
Appendix	19
A List of Participants	19
A.1 Steering Committee	19
A.2 Focus Groups	19

A.3	All Participants	20
B	Whitepaper: Integrating Manual Dexterity with Mobility for Human-Scale Service Robotics	23
B.1	Summary	23
B.2	Commercial Potential	24
B.3	Technological and Economic Risk	27
B.4	Research and Technical Challenges	27
B.5	Action Items	28
B.6	Affiliations of the Signatories	29
B.7	The US Moon-Mars Initiative	31
B.8	American Demographic Trends	31

Preface

The Workshop on Autonomous Mobile Manipulation took place in Houston between March 10 and 11, 2005. Forty-two experts in robot design, control and representation, grasping and manipulation, perception, teaching and learning were assembled to make recommendations about the technical feasibility, scientific hurdles, and commercial potential of Autonomous Mobile Manipulators (AMM).

The motivation for this workshop was the belief among participants and contributing funding agents that AMM represents an important research area with significant scientific, economical, and societal impact. This report summarizes the discussions and findings of the workshop. It attempts to capture diverse opinions and to translate them into recommendations for the academic community as well as for funding agents and policy makers.

The AMM workshop was sponsored by the National Science Foundation (NSF) and by the National Aeronautics and Space Administration (NASA). The organizers and participants gratefully acknowledge this support.

1 Introduction

Planning for the Workshop on Autonomous Mobile Manipulation began in mid-2004, when several roboticists started a discussion about the economic and commercial potential of autonomous mobile manipulators. Such devices are capable of moving about and performing mechanical work in unstructured environments without continuous intervention from human operators. It is anticipated that integrated machines of this type will impact the health-care industry, the servicing of orbiting spacecraft and satellites, security and disaster relief, exploration and interaction with hazardous environments, military applications, and supply chain support and logistics. Appendix B presents a summary of these discussions in the form of a white paper, endorsed by a significant subset of the workshop participants. This position statement eventually led to the organization of a workshop in March of 2005. In this report, we outline the technical and scientific hurdles identified by workshop participants that must be addressed before these systems can be deployed in real-world applications.

One specification of the target technologies in AMM is the creation of a family of automated machines with the visual capacity of a two-year-old, the manual dexterity of a six-year-old, and the ability to move about in human-scale environments. Different applications place different relative emphasis on these functions. Some applications will require significant collaboration with human clients, necessitating new approaches to programming and the ability to engage in social interactions. Comprehensive control knowledge regarding interactions with the world is necessary including mobility and manual interactions with objects. When these interactions go awry, they must be adapted at run-time to the perceived state of the world. The goal of the workshop was to elaborate on the state-of-the-art and to review the dimensions of a research program that could overcome the perceived technical barriers.

Autonomous Mobile Manipulation requires the creation of technologies that are capable of

projecting mechanical work through communication networks to remote locations via robotic surrogates. An integrated **autonomous mobile manipulation** system would successfully marry technologies that are often treated independently. Mobility concerns the mapping and traversal of relatively large scale and variable terrain. It is generally agreed that many useful technologies exist already to map and navigate in certain kinds of environments. Technologies for indoor environments are relatively mature, and large scale traverses of outdoor environments are the subject of a recent DARPA Grand Challenge. In the context of AMM, *mobility* in large-scale environments introduces a variety of objects, tasks, and environmental conditions for dexterous manipulation. This variety and variability in the environment necessitates corresponding resourcefulness and variability in the robot’s behavior. The robot must observe, learn, anticipate, and reason about contingencies in order to execute manual skills in new situations.

Manipulation entails mechanical work to modify the arrangement of objects in the world. This spans a range of tasks from relatively coarse earth moving tasks to dexterous assembly tasks. Specific AMM systems will address some portion of this range. Research in this area is relatively mature in structured, industrial settings. The dexterous end of the manipulation continuum continues to be a challenge. A large body of literature exists for describing the mechanics of manipulating known geometries, however, sensor-based approaches capable of accommodating the characteristic variety in AMM remain on the horizon.

The term *autonomous* in AMM requires that many tasks can be performed without continuous human intervention. In order to achieve this goal, mechanisms for extracting information from the environment, recovering from failure, and modifying plans based on run-time feedback are the central challenging issues. Moreover, autonomy extends into the processes of learning and adaptation—AMM systems must perform sequences of activities based on a high-level task descriptions and must have the capacity to learn and model new circumstances. Autonomy does not denote “isolation,” however. To meet functional goals, AMM systems should co-exist and interact with humans in human environments.

The report begins by surveying the state-of-the-art in important related technologies (Section 2). Following this summary, Section 3 reviews enabling technologies. These pieces of the integrated AMM are technologies that rely on established engineering practice, where new approaches could significantly influence the performance and cost of AMM systems. Workshop participants also identified a number of critical research challenges, which are described in Section 4. Based on these findings, Section 5 describes programmatic recommendations to funding agents and policy makers. These recommendations aim to create a national research environment to enable the United States to build on existing strengths in order to assume scientific leadership in the area of Autonomous Mobile Manipulation.

2 State of the Art Worldwide

Experimental platforms are fundamental prerequisites to conducting research aimed at AMM. The development of robust dexterous manipulation in unstructured environments cannot be advanced using simulation alone. The number of existing platforms can therefore be viewed as a measure of the potential for short term progress in this field.

Countries in Asia and Europe are making significant investments in humanoid robotics. This area differs from AMM in the focus on anthropomorphic robots, but otherwise shares many of the scientific challenges. In this section we briefly review the initiatives underway outside of the United States. The survey demonstrates that to maintain competitiveness with other nations, a concerted research effort has to be initiated immediately. Such an effort should be focused on expanding the competitive advantages of research in the US, before other countries have caught up. For more information, the reader is referred to a presentation by Robert Ambrose, given at the workshop [1].

In Japan and Korea about a dozen anthropomorphic, humanoid robots have been developed and serve as experimental testbeds for ongoing research activities. In most cases these platforms combine legged locomotion with robot hands and extensive sensor packages. Manipulation capabilities are generally achieved by simple grippers rather than dexterous hands. Experimental platforms in the US, by way of contrast, have focused on bi-manual, humanoid upper bodies with multi-fingered hands. These domestic technology efforts include: Robonaut at the NASA Johnson Space Center (Robonaut has demonstrated untethered mobility as well); Domo at MIT; and Dexter at UMass Amherst.

In Asia and Europe governments and funding agencies are making significant financial investments in research activities associated with humanoid robotics and mobile manipulation. For example, Japan is investing \$30 billion over the next five years, exceeding the anticipated total operating budget of the National Science Foundation for the same time span. In Germany, a well-funded, multi-university research initiative on humanoid robotics has been ongoing for several years. Foreign industry, such as Toyota, Honda, and Sony, have initiated substantial research initiatives in the area of humanoid robotics. In the US, no comparable initiatives exist.

In spite of the lack of available experimental platforms and funded research initiatives, the US can still maintain a scientific lead in several relevant areas [1]. This lead is particularly visible in the areas of dexterous manipulation and perception for manipulation and navigation. This report recommends the initiation of concentrated research efforts to maintain this lead and leverage it for progress in Autonomous Mobile Manipulation.

3 Enabling Technologies

The successful deployment of autonomous robots in the context of AMM will require progress in a variety of areas. At the workshop, participants differentiated between *enabling technologies* and *scientific challenges* (Section 4). The former facilitate the development, packaging, and deployment of AMM platforms, but fundamentally only improve on already existing capabilities. Scientific challenges, on the other hand, describe areas in which novel capabilities have to be created to enable aspects of AMM that cannot be addressed with existing technologies.

3.1 Hardware

Participants at the workshop suggest that several parts of an integrated AMM system are available, others need to be modified or improved, and still others required innovation. In the realm of mobility, the availability of adequate sensors, such as laser range finders and ladar, has resulted in significant scientific progress. Similar progress in sensory development for dexterous manipulation is a prerequisite for AMM. Mechanical mechanisms of sufficient sophistication to copy the capabilities of the human hand have been designed. However, comprehensive behavior using these mechanisms in response to environmental stimuli requires an adequate sensory ability. In the context of dexterous manipulation, this entails a suite of sensors that can differentiate haptic events that occur during the manipulation of objects in the environment. Current force sensors and tactile arrays do not provide sufficient sensory capabilities to achieve this goal.

Appropriate packaging of hardware represents another substantial challenge for Autonomous Mobile Manipulation. Safety, reliability, power requirements, and adequate form factor are necessary for the deployment of this technology in human environments. While research can be performed in the absence of these attributes, successful adaptation of the technology will depend on it.

3.2 Standardization of Sensors

The only sensor that is close to being cheap enough and standard enough for everybody to use is vision. Force sensing, tactile sensing, acceleration and thermal sensing are all comparatively non-standard and difficult to use, requiring effort to interface, amplify, filter, interpret, and make robust enough for robotic application. Because they are comparatively difficult to use, our robots have very few of them and there are no good standards to employ. The companies that manufacture these devices are either very small or see the robotics market as a very small part of their market. The miniaturization and standardization of these sensors would facilitate the use of multi-modal sensing in AMM.

3.3 Wiring

Wiring sounds prosaic but, as Steve Jacobsen said some years ago, it is perhaps the #1 problem in making an advanced hand for manipulation. Bus structures help, but raise their own problems regarding a lack of suitable standards for distributed processing and addressing many sensors with very different bandwidth requirements (from 10^1 to 10^3 Hertz). As an example, today's wireless "motes" are well suited for monitoring temperatures but not well suited for monitoring accelerations or transient forces.

Research on conductive polymers that can be integrated with 3D parts is promising and may provide practical alternatives to wires and flex circuits. Research on ways to fabricate sensors in-situ, directly deposited on parts, with local multiplexing, communications will ultimately help.

Beyond the robotics community, automotive companies like GM are very interested in eliminating the need for wiring harnesses. A particular challenge is to minimize the amount of power

wiring, as well as signal wiring. Self-powered sensors and ways of turning the entire car body into a power bus have been considered.

3.4 Embedded, Distributed Processing

Related to the wiring and sensing problems is the need for better, embedded processing environments. The solutions that exist today are not particularly easy to use. The software development environments need further development. Standards and libraries of solutions (e.g. for obtaining smooth force signals) should be developed.

3.5 Actuation

The leading technology in this sector still relies on DC motors and gear trains. It is difficult to make something that is back-drivable (for force control, for having low impedance and for being relatively robust with respect to unexpected external loads but that also has adequate torque and power for manipulation. Too much energy is absorbed in "isometric work" and in providing artificial damping via velocity feedback. Actuators need to be coupled with energy storage and dissipation elements so that one can better manage work. Such novel actuation mechanisms may also play an important role in the creation of inherently safe manipulators (see Section 4.7).

3.6 Power

Supplying sufficient power is still a problem for impressive autonomous performance (e.g. fast mobile robots), albeit progress over the last five years. Fuel cells are continuously improving and may one day represent an important power source. For current use, lithium polymer batteries are better than NiMH and lead acid batteries.

4 Scientific Challenges

To identify the principle dimensions of a research program aimed at deployable AMM systems, workshop participants were asked to form detailed summaries of the state of the art and scientific challenges facing the central dimensions of the overall task. These summaries assume strengths in the US regarding information technologies, mobility, navigation, dexterous robot hand, and manipulation planning and control technologies.

The break-out areas considered were:

1. Grasping and Manipulation
2. Control and Representation
3. Perception

4. Embodiment
5. Teaching and Learning

The presentations in these areas are available on-line at <http://www-robotics.cs.umass.edu/amm/>.

Discussions about these topics took place in the context of these break-out areas, plenary meetings, and break-out meetings focusing on a particular subjects. During these discussions, the participants identified the following scientific and technological challenges for AMM.

4.1 Mobility

AMM systems must be capable to navigate indoor, outdoor, in micro-gravity, or underwater as the case may be. Several approaches to mobility have been deployed and many common frameworks and tools are emerging. Mobility is a relatively mature area that can be harvested for AMM systems.

4.2 Representing Objects and Environments

Sharable and extendible representations that one can use for manipulating things in the world are required. Low-level perceptual apparatus are necessary for extracting attributes of the environment that influence the control of manipulation. Sometimes, these properties are observable only in the context of extended interaction, including geometrical information, object mass, friction, coefficient of restitution, bulk material properties, etc. Representations could be explicit (declarative, parametric) or implicit (procedural, behavioral). They should represent system dynamics in a natural way in order to facilitate coupled mobility and manipulation dynamics and provide a language of manual tasks with formal semantics.

4.3 Grasping and Dexterous Manipulation

There are many interesting and elegant basic algorithms for grasping and dexterous manipulation but they are hard to use and make too many assumptions regarding knowledge about the object. Acquisition of an object and incorporation of immediate sensory information in dynamically reformulating the grasp is a particularly central challenge. To support manual dexterity in tasks involving tool use and non-rigid bodies and to coordinate behavior for hands and eyes (preshape, grasp, and manipulate using visual, force and tactile feedback) new methods for incorporating environmental stimuli into hand control are necessary. A systems-level approach to representations for manipulation control, with shared libraries and languages to the extent possible, would facilitate the deployment of such systems and accelerate new development.

In particular, basic research is needed to devise methods capable of:

- ▶ using tools designed for humans,

- ▶ handling non-rigid bodies reliably,
- ▶ exploiting a visual sensor stream for the grasping approach phase and to pre-shape the hand,
- ▶ incorporating immediate sensory information to dynamically reformulate a grasp, and
- ▶ manipulating a grasped object robustly, while considering visual, force, and tactile feedback.

When performing grasping and manipulation tasks in an unstructured environment, a robotic agent is exposed to a large amount of variability and uncertainty. The successful execution of dexterous skills will depend on the agent’s ability to adapt its behavioral responses in accordance with environmental stimuli. To achieve the above requirements for grasping and manipulation, it will be necessary to develop methods to:

- ▶ represent generic skills that can be applied in a variety of circumstances,
- ▶ capture properties of objects in the environment that are relevant to performing these generic skills, and
- ▶ use those properties to model skills well enough to predict their effects.

4.4 Perception

The term perception should be viewed to include not only computer vision, but also tactile and force feedback, acoustic sensors, proprioceptive information, and any other sensor modality that can help to identify and differentiate environmental stimuli [2]. The respective sensor streams can be used to determine the relative location of a robotic agent, to identify objects in the vicinity to interact or avoid, and they have to be used to generate adequate behavioral responses in the presences of variability and uncertainty in the environment.

In spite of this necessary broader view of perception, the discussions at the workshop mainly focused on computer vision as a means of perception. However, it is very likely that—in addition to a working visual system—a multitude of sensor modalities is required to successfully achieve the elementary skills mentioned in the previous section. A single sensor stream is subject to interruptions (occlusion in vision), to limitations in accuracy, and to restrictions on which aspects of the state it can perceive. These shortcomings can only be compensated by integrating a wider range of sensor modalities. It is therefore important to embark on research initiatives that investigate the issues of perception in this broader sense.

In addition to taking a broader view of perception, it is also essential for perception research to be performed in an application-specific context. The research community in computer vision has developed a number of fundamental techniques that are able to address a variety of real-world problems. Workshop participants believed that these fundamental techniques cannot fully address the perceptual requirements of specific tasks within the domain of AMM. Instead, adequate sensory strategies have to consider sensory streams *in the context of the particular skill and the hardware required to perform that skill*. A consideration of these factors in conjunction seems to be necessary to achieve sufficient robustness in the presence of uncertainty or even hardware failure.

Note that the consideration of multiple sensor streams in the context of an application and

a specific hardware platform goes beyond mere sensor fusion. The specific requirements of the application and platform serve to identify relevant features in multi-modal sensor streams. In the absence of this context, it would be difficult to identify a set of generic features that can address a wide variety of applications and hardware platforms. Therefore, this integrated approach to research in perception has to play an important role in AMM. It will permit significant progress towards the autonomous execution of specific tasks, such as grasping and dexterous manipulation, without imposing the requirement of extracting a general-purpose model from the sensor streams.

4.5 Architectures

The term architecture has numerous connotations in the areas of computer science and in particular in robotics. Here, when referring to an architecture, we refer to the manner in which the components of an AMM system are organized, composed, and integrated. We consider architectures that compose specific skills to achieve robust and more complex behavior; we also consider architectures that compose these robust and complex behaviors in service of a higher-level objective or task. The first type of architecture aims to generate dexterous skills of AMM systems, whereas the second type uses these skills to achieve higher-level tasks in a robust fashion.

To implement robust, fault tolerant, and re-usable manual behavior, we have to develop architectures that exploit prior knowledge about interactions with the world. This knowledge can be learned from experience or captured in procedural representations. An acceptable architecture that exploits such knowledge will support complex, hierarchical organizations of control and perception, perhaps employing expensive (in terms of dollars and/or compute cycles) front end systems. Architectural support for learning models and automatically modifying behaviors must identify and exploit structure and the seamless relationship of planning and control is an important design consideration.

An architecture to combine dexterous skills to achieve higher-level objectives has to combine—in the broadest sense—three different aspects: an interface to permit programming or specification of tasks or high-level objectives, a set of basic motor and sensory skills, and a method to invoke these skills in accordance with the communicated objective. Since such a method has to consider the current state of the world and react to failures and unexpected situations, it must integrate learning, control, planning, and techniques for automated reasoning.

Techniques from classical AI may provide the necessary capabilities to reason about actions in the context of a task and the environment. Over the past decades, the fields of classical AI planning and robotics have progressed mostly independently. While initially the motivations were aligned, both fields have moved apart and today there is very little work at their intersection. Notable exceptions are a number of reasoning architectures in the context of mobile robotics. However, these architectures do not easily extend to AMM. New paradigms are needed that can leverage the achievements of classical AI in the real-world context of AMM. These paradigms need to be expressive and computationally tractable in an open and unstructured world. They have to represent a plan, monitor its execution, and correct it, in the face of errors.

4.6 Human-Robot Interaction

Autonomous mobile manipulators are intended to operate in human environments. Therefore, the ability to interact with humans represents a critical component of AMM systems. Humans communicate with gestures, language, or with implicit assumptions about how a particular task has to be accomplished. If a robot is to serve as an equal partner in this context, it has to be able to communicate at this level.

A variety of applications, such as planetary exploration, for example, may not require the interaction with humans. Here, instead of relying on social communication skills, it is critical to be able to specify tasks and objectives in an unambiguous and intuitive manner. Conventional paradigms of programming do not apply any more, since the variability in the environment cannot be anticipated exhaustively.

Assistive Robotics [4] aims to provide technologies for robotic systems that interact closely with human clients. Applications include: automated assistance for rehabilitation from stroke, injury, or disease; guidance and crowd control in disaster areas; therapy for the cognitively disabled or people with developmental disorders; and assistance for people with special needs. These applications do not necessarily require physical contact with the environment. Some of the required functionality can be performed by interacting with humans through voice and gestures. Assistive Interactive Robotics can thus be viewed as addressing aspects of human-robot interactions, while avoiding the complexities of dexterous manipulation.

The subject of human-robot interaction was not one of the foci of this workshop, since it had been addressed by a previous DARPA/NSF-funded workshop held in September 2001 [3].

4.7 Safe Manipulators

Today's robots respond gracefully to anticipated forces routed through wrist force/torque sensors, but their behavior may be unpredictable when unexpected or undetectable contact situations arise. We need manipulators that respond appropriately to unexpected contacts and interactions—manipulators that can generate high forces without presenting high impedance and high inertia at all times. This problem needs to be addressed through novel mechanism and actuator designs. These designs should minimize distal inertia while still having high load ability and partly a sensor issue. For example, if it was possible to cover the robot with compliant skin that dissipates energy and detects contacts anywhere, then unanticipated contact with the world could be accommodated.

4.8 Commodity Hands

The availability of multi-fingered hands is an important prerequisite for research in dexterous manipulation. Several laboratory examples are available, but at large cost and setup time and with no existing standardized frameworks for control. Provisions must be available for interactions with the palm as well as the fingers. Subsystems for providing and maintaining the safe operation of these devices as they interact with non-stationary and uncertain environments become an important consideration.

5 Recommendations

5.1 Areas of Focus

There is consensus among the workshop participants that the research community should focus on a common set of goals. To determine these goals, we should identify current strengths and build on them, rather than compete with research groups that have a clear lead in a particular area.

With respect to anthropomorphic AMM hardware, many Asian research labs, in particular in Japan and Korea, have a strong lead. Rob Ambrose, after completing a tour of approximately 25 leading research laboratories in Asia, suggests that an appropriate area of focus for US-lead research would be the upper body, including bi-manual dexterity, but ignoring issues associated with legged locomotion.

In other research sectors, US strengths were seen in the areas of information technologies for mobility (as distinct from integrated systems), navigation, hands, manipulation planning, control technologies, and computer vision. These strengths are consistent with the aforementioned focus on upper bodies in the hardware sector: General research problems in mobility are addressed without dependency on a particular mode of mobility and core strengths in manipulation can be co-developed with the corresponding hardware platforms.

The research area of Autonomous Mobile Manipulation differs from most ongoing robotics research efforts in that it requires a significant level of system integration. The autonomous operation of a dexterous, mobile robotic agent requires the coordination of perception, control, reasoning, and interfaces as well as a well-integrated hardware platform. It is therefore of critical importance to initiate multi-disciplinary, integrative research issues to address these requirements.

5.2 Cross-Fertilization

A significant part of the discussion surrounded evidence that the contributing research communities must find opportunities to cross current boundaries between research areas—particularly those between robotics, computer vision, AI, and learning—to create constructive forums for integrated AMM systems. Much of this goal lies in the hands of the research community. This workshop was a useful first step, but the larger research community must respond as well, by continuing such a dialog in the form of additional workshops, special tracks at conferences, and special issues of journals.

Participants perceived the need for large, center-scale grants for integrated and cross-disciplinary efforts. This was a strong opinion of the workshop participants to facilitate cross fertilization among research groups and disciplines. These kinds of programs force the issue of integration and software infrastructure and provides the opportunity for more researchers to get access to more functional experimental platforms.

Some of the scientific challenges described above can be supported via existing programs and relatively small, 1-2 PI efforts. A broad and integrated AMM initiative, however, cannot be supported exclusively by the existing Robotics and Computer Vision programs. To ensure success of

such an initiative, CISE will have to provide broader programmatic support for research in AMM and coordinate various NSF funding programs, potentially beyond the boundaries of CISE.

Overall, the research community has to initiate a concerted effort of raising awareness for AMM among funding agents (NSF, DARPA, NIH, and others) as well as policy makers to elevate AMM to the level of visibility and importance it has on other continents and in other countries. For example, Korea has declared humanoid robotics to be one of its top ten national priorities. Further discussion on how to achieve this goal is necessary.

As a positive outcome of this workshop, John Hollerbach and Jean Ponce began the planning of a joint issue between the IJRR and the IJCV. Initiatives like this have the ability to bring attention to particular research problems relevant to AMM.

5.3 Recommended Research Initiatives

Section 4 describes the most important challenges that were identified by participants in the context of AMM. Many of these challenges lie outside of the scope of the current NSF Robotics and Computer Vision programs. Based on the participants' discussions, we recommend three main programmatic thrusts to be included in the current NSF programs:

- ▶ **Skills:** One of the defining characteristics of AMM is the variability in environmental situations. Robust dexterous manipulation skills are needed that represent general categories of behavior, rather than specialized skills. A research program in this area should emphasize robustness in real-world scenarios, generalizability of behavior, tight integration of multi-modal sensory feedback, and the ability to compose elementary skills into higher-level behavioral units.
- ▶ **Perception:** A research program in the area of perception should address the integration of a variety of sensor streams in a task-specific context. Sensory requirements differ from task to task and adequate and robust closed-loop behavior can only be achieved in the presence of appropriate sensory feedback.
- ▶ **Architectures:** A robotic agent has to be capable of performing a series of skills to accomplish a specified task. This requires architectures for representing closed-loop behavior in a fault-tolerant and robust fashion. Such an architecture should enable the organization of these closed-loop controllers into higher-level skills. At a higher level of abstraction, architectures are needed that generate a plan (choose appropriate sequences of skills) to accomplish tasks, given environmental circumstances. Such an architecture has to realize the goals of planning in classical AI in the context of autonomous mobile manipulation in open environments.

These programmatic recommendations are concerned with research directions that relate to AMM and do not attempt to suggest an allocation of resources to other research areas.

5.4 Infrastructure and Programmatic Recommendations

5.5 Hardware Infrastructure

The availability of an adequate hardware platform for AMM was widely considered as a necessary prerequisite for scientific progress in this area. Just as the proliferation of personal computers aided to jump start much progress in the area of computer science, such a hardware platform would enable researchers to undertake directed research efforts. The availability of such platforms to private people and pre-graduate educational institutions, it was argued, would create a new generation of highly-motivated researchers and a more general awareness for the field.

An ideal hardware platform would be cheap to permit broad distribution among industrial, research, and educational institutions, as well as private homes (comparable to the Sony AIBO). Participants viewed it as a worthwhile endeavor to investigate a minimal set of capabilities such a hardware platform has to exhibit to still permit useful research.

Centralized infrastructure and shared facilities were also considered as a model of making AMM hardware platforms available to the research community. This model looks more like the “super computer” model than the distributed “personal computer.” Instead of or in addition to creating a cheap platform and making it widely available, it may be possible to create experimental centers that are equipped with a costly but capable experimental platform. This platform could be shared remotely by several institutions or researchers, much like super-computers were time-shared in the early days of computing. An important caveat, however, is the fact that robots—in contrast to super computers—may break when interacting with the environment. Furthermore, teleoperation of robots and live video feeds are subject to bandwidth constraints that make true experimentation difficult.

Short of providing an entire AMM platform, the availability of new commodity hands would boost research in manipulation and control. These hands should have provisions for interactions with the palm as well as the fingers, and intrinsic safety features. Autonomous manipulation represents an important domain of AMM in which much additional progress is needed. The availability of commodity hands could therefore be viewed as a first step towards an integrated hardware platform for AMM.

5.6 Software Infrastructure

Successful research in AMM has to combine techniques and components from different laboratories and across different scientific disciplines. Such an integration is beyond the scope of most ongoing research efforts. It will therefore be necessary to develop a broad range of methods and principles to support such integration, in particular in the context of research. This will involve standards, protocols, shared architectures, common interfaces, representations and file formats, libraries, data sets, and many additional means of facilitating integration and sharing of experiences and resources.

The importance of software infrastructure and software architecture is generally underestimated in the academic world, in spite of many “best practices” that have proven extremely suc-

cessful in industry over the course of several decades. In this domain, researchers should turn to these practices and adopt them to accommodate a level of integration among homogeneous systems that is commonly performed among practitioners in industry.

5.7 Intermediate Goals

Participants discussed which intermediate goals would be most appropriate to guide ongoing research activities. Doug Gage emphasized that DARPA is focused on specific capabilities that can be delivered within a time-frame of 18 months. Such capabilities have to relate to applications of interest to DARPA.

Potential intermediate (middle-term) goals—including some of interest to DARPA—include:

- ▶ logistics, supply chain applications, loading pallets/trucks, WalMart-scale distribution systems, re-fueling operations
- ▶ battlefield support: deployment of sensors, handling explosive ordnance, rescue operations
- ▶ agriculture and construction
- ▶ exploration - assisting astronauts, taking samples, assembling, repairing, inspecting, rescuing
- ▶ elder care - assistive/service robotics, cleaning, companion, gofer, medical interfaces, cognitive and physical prosthesis

Certain technological milestones are necessary to meet these goals:

- ▶ navigate indoor, outdoor, in microgravity, or underwater
- ▶ robust motion capabilities for dexterous manipulation
- ▶ state estimation robust with respect to hardware failure, operating conditions, with obscuration and uncertainty
- ▶ eye-hand coordination - preshape, grasp, and manipulate with visual, force and tactile feedback
- ▶ new commodity hands, with provisions for interactions with the palm as well as the fingers, and intrinsic safety features.
- ▶ tool use and non-rigid bodies
- ▶ mechanisms for actively modeling controlled environmental interactions
- ▶ human-machine communication

6 Acknowledgments

The organizers would like to express their gratitude to Yunku Juh (National Science Foundation, NSF) and Bob Savely and Robert Ambrose (National Aeronautics and Space Administration, NASA). Without their financial support and organizational assistance this workshop would not have been possible.

The members of the steering committee were instrumental in the preparation of the technical content for this workshop. The organizers would like acknowledge their contributions to the success of this event. The steering committee consisted of Rob Ambrose (Platforms and International Initiatives), Rod Grupen (Grasping and Manipulation), Vijay Kumar (Control and Representation), Greg Hager (Perception), Kenneth Salisbury (Embodiment), and Stefan Schaal (Teaching and Learning).

The quality and relevance of a workshop critically depend on the participants. We would like to thank all participants for their lively participation and intellectual contributions (see Appendix A). Their involvement, interest, and dedication lead to the insights and recommendations described in this report.

Finally, the organizers are indebted to Priscilla Coe from UMass Amherst and to Melita Stubblefield from NASA Johnson Space Center for their help in organizing and executing this event. In addition, Priscilla Coe administered the financial support provided by NSF and NASA, which permitted us to focus on the technical content. We are grateful to work with such dedicated and competent people.

References

- [1] R. Ambrose. Platforms and international initiatives. Presentation at the NSF/NASA Workshop on Autonomous Mobile Manipulation, March 2005.
<http://www-robotics.cs.umass.edu/amm/presentations/ambrose.pdf>.
- [2] G. Hager. Preception. Presentation at the NSF/NASA Workshop on Autonomous Mobile Manipulation, March 2005.
<http://www-robotics.cs.umass.edu/amm/presentations/hager.pdf>.
- [3] E. Rogers and R. Murphy. DARPA/NSF Study on Human-Robot Interaction.
<http://www.csc.calpoly.edu/~erogers/HRI/>, September 2001.
- [4] S. Schaal. Teaching, learning, and developmental programming. Presentation at the NSF/NASA Workshop on Autonomous Mobile Manipulation, March 2005.
<http://www-robotics.cs.umass.edu/amm/presentations/schaal.pdf>.
- [5] S. Singh, A. G. Barto, and N. Chentanez. Intrinsically motivated reinforcement learning. In *Proceedings of the International Conference on Neural Information Processing Systems (NIPS)*, Vancouver, Canada, December 2004.

A List of Participants

A.1 Steering Committee

1. Robert Ambrose, Johnson Space Center, NASA
Platforms and International Initiatives
2. Rod Grupen, University of Massachusetts Amherst
Grasping and Manipulation
3. Greg Hager, Johns Hopkins University
Perception
4. Vijay Kumar, University of Pennsylvania
Control and Representation
5. Kenneth Salisbury, Stanford University
Embodiment
6. Stefan Schaal, University of Southern California
Teaching, Learning, and Programming

A.2 Focus Groups

1. *Platforms and International Initiatives*
Robert Ambrose, Johnson Space Center, NASA
2. *Grasping and Manipulation*
Rodney Brooks, Massachusetts Institute of Technology
Mark Cutkosky, Stanford University
Rod Grupen, University of Massachusetts Amherst
Andrew Miller, Columbia University
Nancy Pollard, Carnegie Mellon University
Jeff Trinkle, Rensselaer Polytechnic Institute
3. *Perception*
Narendra Ahuja, University of Illinois at Urbana-Champaign
Yiannis Aloimonos, University of Maryland
Dana Ballard, University of Rochester
Chris Brown, University of Rochester
Greg Hager, Johns Hopkins University
Martial Hebert, Carnegie Mellon University
Martin Jägersand, University of Alberta
Jean Ponce, University of Illinois at Urbana-Champaign

4. *Control and Representation*
John Hollerbach, University of Utah
Oussama Khatib, Stanford University
Vijay Kumar, University of Pennsylvania
Al Rizzi, Carnegie Mellon University
Daniela Rus, Massachusetts Institute of Technology
5. *Embodiment*
John Evans, President, John M Evans LLC
Akhil Madhani, Disney Imagineering R&D
John Morrell, Segway LLC
Kenneth Salisbury, Stanford University
6. *Teaching, Learning, and Programming*
Christopher Atkeson, Carnegie Mellon University
Drew Bagnell, Carnegie Mellon University
Geoff Gordon, Carnegie Mellon University
Jessica Hodgins, Carnegie Mellon University
Maja Matarić, University of Southern California
Andrew Ng, Stanford University
Brian Scasselatti, Yale University
Stefan Schaal, University of Southern California

A.3 All Participants

1. Narendra Ahuja, University of Illinois at Urbana-Champaign
2. Yiannis Aloimonos, University of Maryland
3. Robert Ambrose, Johnson Space Center, NASA
4. Dana Ballard, University of Rochester
5. Drew Bagnell, Carnegie Mellon University
6. Oliver Brock, University of Massachusetts Amherst
7. Rodney Brooks, Massachusetts Institute of Technology
8. Chris Brown, University of Rochester
9. Mark Cutkosky, Stanford University
10. John Evans, President, John M Evans LLC
11. Andrew Fagg, University of Oklahoma
12. Doug Gage, XPM Technologies

13. Rod Grupen, University of Massachusetts Amherst
14. Greg Hager, Johns Hopkins University
15. Martial Hebert, Carnegie Mellon University
16. John Hollerbach, University of Utah
17. Martin Jägersand, University of Alberta
18. Oussama Khatib, Stanford University
19. Vijay Kumar, University of Pennsylvania
20. Akhil Madhani, Disney Imagineering R&D
21. Giacomo Marani, University of Hawaii
22. Maja Matarić, University of Southern California
23. Andrew Miller, Columbia University
24. Mark Moll, Rice University
25. John Morrell, Segway LLC
26. Andrew Ng, Stanford University
27. Una-May O'Reilly, Massachusetts Institute of Technology
28. Alan Peters, Vanderbilt University
29. Nancy Pollard, Carnegie Mellon University
30. Jean Ponce, University of Illinois at Urbana-Champaign
31. Al Rizzi, Carnegie Mellon University
32. Stephen Rock, Stanford University
33. Daniela Rus, Massachusetts Institute of Technology
34. Kenneth Salisbury, Stanford University
35. Bob Savely, Johnson Space Center, NASA
36. Brian Scassellati, Yale University
37. Stefan Schaal, University of Southern California
38. William Townsend, Barrett Technology Inc.
39. Jeff Trinkle, Rensselaer Polytechnic Institute

40. Ben Wegbreit, Strider Labs, Inc.
41. Thomas Willeke, NASA Aimes
42. Junku Yuh, National Science Foundation

B Whitepaper: Integrating Manual Dexterity with Mobility for Human-Scale Service Robotics

B.1 Summary

This position paper argues that a concerted national effort to develop technologies for robotic service applications is critical and timely—targeting research on integrated systems for mobility and manual dexterity. This technology provides critical support for several important emerging markets, including: health care; service and repair of orbiting spacecraft and satellites; planetary exploration; military applications; logistics; and supply chain support. Moreover, we argue that this research will contribute to basic science that changes the relationship between humans and computational systems in general.

This is the right time to act. New science and key technologies for creating manual skills in robots using machine learning and haptic feedback, coupled with exciting new dexterous machines and actuator designs, and new solutions for mobility and humanoid robots are now available. A concentrated program of research and development engaging federal research agencies, industry, and universities is necessary to capitalize on these technologies and to capture these markets.

Investment in the US lags other industrialized countries in this area partly because initial markets will probably serve health care and will likely appear first outside the borders of the United States. It is the considered opinion of the signatories of this document that this situation must be reversed. To capitalize on domestic research investment over the past two decades, and to realize this commercial potential inside of the United States, we must transform critical intellectual capital into integrated technology now. Our goal is to ensure that the US economy and scientific communities benefit as this nascent market blossoms and we will outline the economic risks of allowing other nations to continue to go it alone.

Robert Ambrose

Christopher Atkeson

Oliver Brock

Rodney Brooks

Chris Brown

Joel Burdick

Mark Cutkosky

Andrew Fagg

Roderic Grupen

Jeffrey Hoffman

Robert Howe

Manfred Huber

Oussama Khatib

Pradeep Khosla

Vijay Kumar

Lawrence Leifer

Maja Matarić

Randal Nelson

Alan Peters

Kenneth Salisbury

Shankar Sastry

Robert (Bob) Savely

Stefan Schaal

B.2 Commercial Potential

The service robotics industry is projected to be a huge commercial opportunity with products going to market at an accelerating rate over the next 20 years. Service robotics currently shares some important characteristics with the automobile industry in the early twentieth century [2] and the home computer market in the 1980's [6]. We argue that many of the new opportunities that exist rely on technologies supporting manual dexterity. Specifically, the marriage of new research on manual dexterity involving grasping and manipulation with more mature technologies for signal processing and mobility can yield new integrated behavior that supports applications heretofore unattainable.

We are advancing this argument now because new developments regarding actuation and sensing promise to make robots more responsive to unexpected events in their immediate surroundings. This is a boon to mobility technology and is the “missing link” to producing integrated manipulation systems. New, bio-inspired robots are demonstrating impressive performance and better robustness than their traditional robotic forbears. We have discovered that legged locomotion need not be as difficult and complex as we had thought. Therefore, we can afford to add new capabilities, and complexity, on top of legged platforms. Grasping and dexterous manipulation still await comparable insights and the technological foundations are now in place.

Health Care: Health care is the largest segment of the US economy and is becoming too expensive to deliver. We follow closely on the heels of Asia and Europe where demographic pressure is forcing technology to meet the demand for a more efficient means of “in-place” elder care now. This pressure is due largely to a precipitous decline in the ratio of wage earners to retirees and the prospects are very nearly the same in much of the industrialized world. So far, the US is forestalling the problem by holding the birth rate slightly above 2.0 (2.07 in 2004 [1]) but inevitably, the same challenge faces the US health care system in the future [5]. The prospects for large scale institutionalization of the elderly population are daunting, both in terms of the investment in infrastructure required and in the quality of life issues as older people are moved out of their homes and into centralized facilities.

The answer is technology for “aging in place.” The centralized “mainframe” approach to health care for baby boomers around the globe must be augmented with information technology and assistive devices that promise to be the health care equivalent of low cost personal computers [6]. In the long term, we must “consumerize” health and wellness technologies and make it practical and affordable to push them into existing homes. The goal must be to provide cognitive and physical assistance to the elderly and infirm. Dexterous machines are an important facet of this armamentarium. In the shorter term, mobile manipulators can make significant contributions to health care in existing hospitals for services to convalescence and recovery operations before they make it into people's homes. There are millions of stroke and heart attack patients that are not currently getting adequate post-surgical followup and quality-of-life support. These systems must share critical geometry with humans to co-exist in human environments and to serve as assistants to clients across a wide spectrum of service specifications.

In-Orbit Servicing of Satellites: The dream of a re-usable space shuttle that can service the International Space Station (ISS) and important and unique orbiting platforms like the Hubble space telescope is waning as administrators at NASA struggle with the exposure of humans to a 1:50 risk of catastrophic failure of the spacecraft. The design of the shuttle is driven by the configuration required for a space plane with the significant overhead of human life support. The support chain of maintenance and supply at the threshold of space should not expose humans to unnecessary risk. Robotic maintenance missions are the answer. Collaboration between robots and humans in such missions is facilitated when both humans and robots can operate the same tools and have overlapping sensory viewpoints, accessible workspace, and force and velocity capacity. This argues for anthropomorphic robot design to achieve these specifications and once again, robot hands and dexterous manipulation are an important key to success.

Planetary Exploration: On January 14, 2004, the President outlined his goals to return to the Moon and then push onto Mars. These goals will require the construction of habitat, and the maintenance and operation of science labs, geological exploration crews, chemical processing plants, etc. The human pioneers that first undertake this mission will be exposed to tremendous risk while outside of protected habitat and yet such activity cannot be avoided. There is a clear role for robots that can both navigate and manipulate with some degree of autonomy. Non-dexterous mobile manipulators capable of excavation and resource extraction will partner with dexterous mobile manipulators to mine raw materials and to dig trenches, install habitat modules, and then cover them with regolith to protect them from radiation. The same machines will transition over time to assist humans that occupy these habitats, and will also serve as caretakers in between human crews.

Computer systems that act as cognitive and physical prosthetics for astronauts in these hostile environments are feasible and necessary to reach these ambitious goals. The round trip communication latency can vary between 2 seconds (low Earth orbit) to upwards of 30 minutes (Mars depending on where the planets are in their orbits) making it impossible for controllers on Earth to react to problems on the space vehicle or in Martian habitats in a timely manner. Intelligent systems with the capacity for collaborative and independent problem solving become critical to mission success. Rather than simply follow preprogrammed commands, robots must be able to assess a situation and recommend a course of action without human intervention every step of the way and then effect a solution involving a spectrum of human-robot collaborations.

Manual dexterity and autonomous mobility are key elements of this vision. The coupling between systems that are designed to avoid some forms of contact with the environment while seeking others—often simultaneously and in service to multiple objectives—is critical to mission success.

Military Applications: The Pentagon spent \$3 billion on unmanned aerial vehicles between 1991 and 1999 and is reportedly prepared to spend \$10 billion by 2010 under a Congressional mandate that one third of its fleet of ground vehicles should be unmanned by 2015 (National Defense Authorization Act for Fiscal Year 2001, S. 2549, Sec. 217). The same impact is expected for pilotless air and water vehicles, where drone aircraft for reconnaissance and air to ground

missile deployment is already becoming accepted military doctrine. Boeing, Northrop Grummond, and Intel (among many others) are currently assembling infrastructure to support these significant markets.

A similar revolution in military technology, one that exploits new technology for manual dexterity, finally promises to replace human *hands* in dangerous environments as well. With the ability to manipulate, autonomous machines may one day serve to reduce the exposure of human soldiers in combat, in the supply chain (re-fueling, ordnance), in BSL4 facilities for handling dangerous substances. Moreover, this new technology can provide mobile information gathering agents with the ability to probe environments, dig, and sample soil.

Logistics and Supply Chain Support Almost every aspect of product distribution is automated with two notable exceptions: transportation and load/unload at distribution centers. Loading and unloading shipping containers, and inventory control in warehouse and distribution operations can be automated in the near term by mobile manipulation systems. For example, if there is sufficient demand volume, there is a significant cost advantage to using shipping containers to transport materials by land and sea. As a result, large distribution systems, like Walmart, can reduce costs significantly by making inventory management and distribution more autonomous. Mobile manipulation technologies can support automating the rest of distribution, logistics, and material supply chain reducing costs, enhancing inventory tracking and supply chain security.

Contributions to Basic Science: Many of the traits we consider uniquely human stem not from great capacity for strength, speed, or precision, but instead from our adaptability and ingenuity—our dexterity. When we move from laboratories and simulation into the real world, the merits of flexibility and adaptation and the cognitive representations that support these processes are clearly justified [7].

The human hand and the neuroanatomy that co-evolved to support it are critical to the success of human beings on earth and our distinctive cognitive ability. In addition to creating integrated mobile manipulators and an array of autonomous manual skills, research on mechanisms, control, and representations for robot hands has the potential to advance our understanding of the computational processes underlying cognition. Specifically, the process of grounding knowledge has important implications in language, human development, and man-machine interfaces.

The result will be practical implementations of machines and computational decision making that responds to changing situations and complicated environments. Mobile manipulators exploit structure in the form of Newtonian mechanics. We may exploit rules governing other domains as well: in bioinformatics, molecular forces and reaction dynamics govern behavior; in enterprise systems, business rules form categories of transactions and documents. A focused and integrated research initiative in this area will prepare for emerging commercial markets, lead to new kinds of adaptable machines, and influence the future relationship between networks of machines and human societies.

B.3 Technological and Economic Risk

Technological Leadership Commercial versions of mobile manipulation systems will support service robotics, health care, military and space applications—markets that can transform economies. Moreover, virtually any computational system that interacts with complex and open environments or datasets will benefit.

Europe and Japan are investing tremendous resources in the development of this technology with \$30 billion dollars of investment planned over the next 5 years in Japan alone to prepare for the nascent service robotics market aimed at elder care. By way of comparison, the total budget for the National Science Foundation, including operations and all areas of supported research is approximately \$5.6B/year in 2004 (NSF PR 04-12 - February 2, 2004). Honda, Sony, and Toyota are making significant investments into humanoid robotic technology. Toyota is launching a service robotics division to respond to the R&D challenges posed in this new domain [3].

So far, most efforts in humanoid robotics have focused largely on walking. It stands to reason that new technologies for manipulation and manual dexterity are next. With stakes of this magnitude, the US must take measures to mobilize its resources by actively building consortia of industry and academia to meet this challenge. If this is not afforded the priority it deserves, then we will have squandered our technological advantage.

Educational US academic institutions have been the torch bearer for high technology training for the entire world community for several decades. Sadly, the United States is losing that distinction. Applications for graduate school in the US from Europe and Asia are down starkly in the past couple of years. This is due in part to restricted access of foreign-born students to our educational market since 9/11/01, and also partly to the massive investment by these nations into research, development, and education. US educational institutions are an effective pipeline for creative young researchers that can be emulated in other parts of the world to serve their economies. We are being outspent and it will take much less time to lose our advantage in education and training than it took to create it. We argue that a concerted investment in technologies for mobility and machine dexterity involving all branches of engineering, materials science, computer science, and cognitive science will serve to shore up this slowly eroding infrastructure and attract the world's best young minds into areas of critical future economic value to our nation.

B.4 Research and Technical Challenges

- Embodiment - Power, actuation, packaging, mobility, mechanisms, sensors
 - Reliable integrated packages for sensing (tactile, visual, auditory, and proprioceptive) and actuation (power source, power-to-weight ratio, volume, controllability) systems must be developed to meet these goals.
 - Simple, robust, cost effective mechanical systems - combining safety, load carrying capacity and speed, dexterity and power. Hands are an essential sensorimotor component for achieving the applications cited.

- A new approach encompassing embodiment, control, and cognitive organization is necessary to fuel critical future applications.
- Grasping and Manipulation
 - New control techniques are required for robots to interact purposefully with the environment at scales representing the human niche (ranging approximately from $10^{-2} m$ to $10^1 m$, from $0.01 N$ to $10^2 N$, and over durations ranging from milliseconds to hours).
 - New techniques are required to model and reason about complex systems and “system of systems” ranging from coordinating multiple limbs, large scale mobility, multiple robots, and human-robot teams.
- Control/Perception/Representation/Cognition
 - New approaches to representing sensorimotor interaction are needed at several levels (feature, object, context) and at several spatial and temporal scales.
 - Incomplete world state must be addressed with intelligent, active information gathering technologies that recover critical context on a task-by-task basis.
 - New approaches are needed for modeling “activity” in sensor data and discrete event feedback.
 - Representations employed by robots must be *grounded* in natural phenomena accessible directly to humans and robots alike.
- Teaching, Learning, and Developmental Programming -
 Interactions between body parts, sensors, archival information, other robots, human collaborators, and an unstructured environment form hierarchies of complex systems that challenge traditional approaches to programming.
 - New approaches to instruction, imitation, and exploration must be incorporated into machine learning techniques to acquire the building blocks of cognitive systems.
 - Formal models of generalization, and processes of assimilation and accommodation.
 - New programming techniques are needed that incorporate lifelong training and instruction.
 - Methods for transferring experience earned by one agent (human or robot) into meaningful and actionable knowledge by another agent or agents.

B.5 Action Items

The immediate agenda involves using this document to address the community, including funding agencies, industry, and academia, in order to direct attention to this critical technical, economic, and scientific challenge. The signatories of this document would be happy serve help this role.

This introduction will be followed by presentations at workshops, symposia, and panels to elaborate on the critical technical challenges and opportunities, to create a more detailed research agenda, and to create an organized advocacy group and fund-raising strategy. A joint NSF/NASA Workshop on Autonomous Mobile Manipulation is in the planning stages. and will lead to a proposal soon to kick this process off. We invite outside participants to help in the organization of such events and we look forward to serving the information needs of industry and federal agencies in realizing this important milestone in service robotics applications.

References

- [1] Central Intelligence Agency. The world factbook. <http://www.cia.gov/cia/publications/factbook/geos/us.html>, 2004.
- [2] Uutinen Julkaistu. Service robotics defines future of man-machine interaction. Tekes Technical Programmes, <http://akseli.tekes.fi/Resource.phx/tuma/kone2015/en/robotics-uutinen.htx>, 2004.
- [3] D. Kara. Sizing and seizing the robotics opportunity. In *COMDEX*, 2003.
- [4] R.W. Pew and editors Van Hemel, S.B. *Technology for Adaptive Aging*. The National Academies Press, Washington, DC, 2003.
- [5] Nicholas Roy, Gregory Baltus, Dieter Fox, Francine Gemperle, Jennifer Goetz, Tad Hirsch, Dimitris Margaritis, Mike Montelermo, Joelle Pineau, Jamie Schulte, and Sebastian Thrun. Towards personal service robots for the elderly. In *Workshop on Interactive Robots and Entertainment (WIRE)*, 2000. <http://web.mit.edu/nickroy/www/papers/wire2000.pdf>.
- [6] B. Schlender. Intel's Andy Grove: The next battles in tech. *Fortune*, pages 80–81, May 2003.
- [7] M. Swinson and D. Bruemmer. Expanding frontiers of humanoid robotics. *Intelligent Systems*, 12(17), 2000.

B.6 Affiliations of the Signatories

This document was prepared by Professors Rod Grupen and Oliver Brock at the Laboratory for Perceptual Robotics, University of Massachusetts Amherst.

Ambrose, Robert - Robonaut Team Leader, Automation, Robotics and Simulation Division, NASA Johnson Space Center

Atkeson, Christopher - Professor, Robotics Institute, Carnegie Mellon University

Brock, Oliver - Assistant Professor of Computer Science and Co-Director of the Laboratory of Perceptual Robotics, University of Massachusetts Amherst

Brooks, Rodney - Professor of Computer Science and Engineering and Director of the Computer Science and Artificial Intelligence Laboratory (CSAIL) at the Massachusetts Institute of Technology, and co-founder and Chief Technology Officer of iRobot Corporation.

Brown, Chris - Professor, Computer Science, University of Rochester

Burdick, Joel - Professor of Mechanical Engineering and Bioengineering; Deputy Director, Center for Neuromorphic Systems Engineering, California Institute of Technology

Cutkosky, Mark - Professor and Associate Chair Design Division, Department of Mechanical Engineering, Stanford University

Fagg, Andrew - Associate Professor of Computer Science, University of Oklahoma

Gruppen, Roderic - Professor of Computer Science and Director of the Laboratory of Perceptual Robotics, University of Massachusetts Amherst

Hoffman, Jeffrey - Professor, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology

Howe, Robert - Professor of Engineering, Director of the BioRobotics Laboratory, Harvard University

Huber, Manfred - Assistant Professor, Department of Computer Science and Engineering, University of Texas at Arlington

Khatib, Oussama - Professor of Computer Science, Stanford University

Khosla, Pradeep - Professor and Head, Department of Electrical and Computer Engineering, Philip and Marsha Dowd Professor of Electrical and Computer Engineering and Robotics, Carnegie Mellon University

Kumar, Vijay - Professor, Mechanical Engineering and Applied Mechanics, and Computer and Information Science, University of Pennsylvania

Leifer, Lawrence - Professor, Department of Mechanical Engineering, Director, Stanford Center for Design Research Director, Stanford Learning Laboratory, Stanford University

Matarić, Maja - Associate Professor, Computer Science Department, Director, USC Center for Robotics and Embedded Systems (CRES), University of Southern California

Nelson, Randal - Associate Professor, Department of Computer Science, University of Rochester

Peters, Alan - Associate Professor of Electrical Engineering, Department of Electrical and Computer Science, Vanderbilt University

Salisbury, Kenneth - Departments of Computer Science and Surgery, Stanford University.

Sastry, Shankar - NEC Distinguished Professor of Electrical Engineering and Computer Sciences and Bioengineering, Berkeley University

Savely, Robert (Bob) - Senior Scientist for Advanced Software Technology, Chief Scientist Automation, Robotics and Simulation Division, NASA Johnson Space Center

B.7 The US Moon-Mars Initiative

The Moon-Mars initiative includes a new space vehicle to return astronauts to the Moon as early as 2015. Highlights of President Bush's space exploration goals include:

- completing work on the International Space Station by 2015;
- developing and testing a new manned space vehicle(the crew exploration vehicle) by 2008 and conducting the first manned mission by 2014;
- returning astronauts to the moon as early a 2015 and no later than 2020;
- using the Moon as a stepping stone for human missions to Mars and worlds beyond; and
- allocating \$11 billion in funding for exploration over the next five years, which includes requesting an additional \$1 billion in fiscal 2005 (Congress responded in July by recommending a \$220 million reduction)

B.8 American Demographic Trends

The United States has seen a rapid growth in its elderly population during the 20th century. The number of Americans aged 65 and older climbed to 35 million in 2000, compared with 3.1 million in 1900. For the same years, the ratio of elderly Americans to the total population jumped from one in 25 to one in eight. The trend is guaranteed to continue in the coming century as the baby-boom generation grows older. Between 1990 and 2020, the population aged 65 to 74 is projected to grow 74 percent.

The elderly population explosion is a result of impressive increases in life expectancy. When the nation was founded, the average American could expect to live to the age of 35. Life expectancy at birth had increased to 47.3 by 1900 and the average American born in 2000 can expect to live to the age of 77.

Because these older age groups are growing so quickly, the median age reached 35.3 years in 2000, the highest it has ever been. West Virginia's population is the oldest, with a median age of 38.6 years. Utah is the youngest, with a median age of 26.7 years.