Working and Learning with Knowledge in the Lobes of a Humanoid's Mind

Robert Ambrose[1], Robert Savely[1], William Bluethmann[2], Eric Huber[3], David Kortenkamp[3],

[1] NASA Johnson Space Center, Houston Tx 77058 USA
{Robert.O.Ambrose and Robert.T.Savely}@nasa.gov
http://robonaut.jsc.nasa.gov/
[2] Hernandez 2400 NASA Rd 1, Mail Stop C35, Houston Tx 77058
William.j.bluethmann1@jsc.nasa.gov
[3] Metrica 1020 Hercules, Houston Tx, 77058, USA,
{huber and korten}@traclabs.com

Abstract. Humanoid class robots must have sufficient dexterity to assist people and work in an environment designed for human comfort and productivity. This dexterity, in particular the ability to use tools, requires a cognitive understanding of self and the world that exceeds contemporary robotics. Our hypothesis is that the sense-think-act paradigm that has proven so successful for autonomous robots is missing one or more key elements that will be needed for humanoids to meet their full potential as autonomous human assistants. This key ingredient is knowledge. The presented work includes experiments conducted on the Robonaut system, a NASA and the Defense Advanced research Projects Agency (DARPA) joint project, and includes collaborative efforts with a DARPA Mobile Autonomous Robot Software technical program team of researchers at NASA, MIT, USC, NRL, UMass and Vanderbilt. The paper reports on results in the areas of human-robot interaction (human tracking, gesture recognition, natural language, supervised control), perception (stereo vision, object identification, object pose estimation), autonomous grasping (tactile sensing, grasp reflex, grasp stability) and learning (human instruction, task level sequences, and sensorimotor association).

## 1 Introduction- Challenges for Autonomous Space Humanoids

NASA's motivation is to build humanoid class robots that can aid spacewalking astronauts, accelerating their work pace, and allowing them to be more productive while working in the extremes of space. Our approach has been to build multiple Robonaut units for developing the required mechanical and electrical technology, and then make these robots available as testbeds for university teams developing the autonomy required to complete the system. To meet our objectives, we will need to overcome many challenges, including facing new problems in autonomy, the environment, and safety around humans.

1.1 The challenge of working in 0g

Space walking is poorly named, as it has little in common with how animals walk on Earth. Space walking is more akin to mountain climbing in scuba gear, while parachuting in a freefall- an odd combination of effects and equipment to help people do a demanding job. Robots are now being studied for service in this same domain, working on large scale space structures like the Space Station, servicing science or military platforms in high orbit, or riding on the outside of a space craft in transit to Mars, the Moon or other destinations. Of particular interest are the problems of handling objects in the frictionless environment of 0g, where tools are tethered for crew use, and locomotion requires a hybrid of climbing, propulsion or other means.

NASA has a massive investment in Extravehicular Activity (EVA) tools, well documented in the EVA Tool Catalog [1] and countless mission logs. The depth and breadth of this tool set is immense, from space versions of conventional hand tools, to specialized tools and other EVA interfaces. NASA has accepted the challenge of building a robot that can work with these existing tools, as well as the other EVA infrastructure such as hand rails, EVA airlocks, tethers, and other installed hardware that can not be retroactively rebuilt to ease robot design. Figure 1 shows the first US EVA, a more recent EVA using a powered torque tool, and Robonaut emulating a climbing task on a 0g mockup at NASA's Johnson Space Center (JSC).
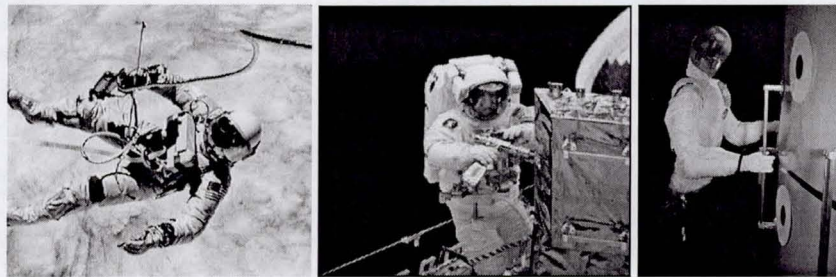
Fig. 1. Shown are images of the first US space walk by Ed White, a scene from Space Shuttle flight STS-103 of an astronaut using a power torque tool to tighten bolts, and Robonaut A climbing on the exterior of a 0g mockup that is free to move, emulating 0g conditions for a free climbing humanoid despite Robonaut A being mounted to a fixed pedestal in 1g

In order to move about in this environment, a robot must be able to climb autonomously, using gaits that smoothly manage its momentum and that minimize contact forces (walking lightly) while providing for safety in the event of an emergency requiring the system to stop [2]. All three of these objectives are now being explored at NASA's Johnson Space Center, using the Robonaut system and a set of mockups that emulate the 0g condition. NASA's goal for Robonaut is to develop the control technology that will allow it to climb on the outside of the Space Shuttle, the Space Station, and satellite mockups at JSC, enabling the robot to perform EVA task setups or serve as an astronaut's assistant.

1.2 The challenge of safely working with humans

Robots are dangerous. Decades of experience have shown that factories must be designed fundamentally to separate human and robot workspaces, with elaborate yet essential measures to track and protect humans from their own actions [3]. The safety challenge is even greater for working with astronauts, where the space suit and the thermal vacuum condition make crew even more fragile than factory workers. NASA's goal in developing Robonaut was to build a soft, smart, safe machine that could handle objects without leaving burrs or other damage that could later injure the gloves or suit fabrics of the astronauts. Robonaut's soft glove, and padded limbs and body are unusual in the world of robotics, having been designed for this added safety.

Mechanical padding is not enough. Robonaut is designed to move at a slow, deliberate rate that can easily be anticipated and avoided by a human, in sharp contrast to the blurring speed of factory manipulators. Robonaut's electromechanical power system is voltage limited, allowing the robot to be constrained in its top speed by the back EMF of the limb's motors, and limiting the stall strength that is available when the robot meets a stiff object and pushes with no velocity. But even this electromechanical safety is inadequate, since it assumes that slow motion or small forces are safe- an assumption that will be wrong for many circumstances. The third level of safety in NASA's humanoid class robots is a redundancy in power, sensing and control. This includes redundant communication pathways, redundant torque sensing, redundant position sensing, and redundant temperature sensing. Comparisons of proprioceptive sensors in the limbs with head mounted cameras that can see limb location can also be performed.

1.3 New challenges in Autonomy

Twenty years ago, the intelligent robotics community was just forming, and there was little consensus on approaches or architectures for what we now call navigation of mobile robots. Today, this domain has greatly matured, with numerous architectures commercially available to upgrade a wheeled vehicle to a sophisticated thinking machine. But the class of interaction that such a machine can have with its environment is limited to perception, where physical contact is intentionally avoided. The early mobile robots could move through an arrangement of objects [4], they evolved into detailed mapping machines able to understand and report on their world [5], and now have attained seek or hide functions [6].

This has been a steady advance in the state of the art that is now impacting civilian and military applications of robotics. As complex as these cognitive functions are, the sophistication of their interaction with the environment pales in comparison to using a tool to modify the world [7]. Understanding the world well enough to know that a change is needed/possible, and then forming a plan to use a known tool to implement that change is an infinitely open challenge. Emerging theories on the role of tool use and weapon making in human cognition [8] [9] bode poorly for any robotics team that intends to quickly automate a humanoid as a competent tool user.

The existing world mapping techniques will be essential for the first phase of this effort, but must be combined with symbolic, relational, associative and generally qualitative representations of knowledge to complete the picture. A robot sees a box with rough texture on its top surface. A human looks at the same scene, and sees a workbench strewn with hand tools that bring back a lifetime of memories. Bringing the robot to the same perception level as the human tool user is the first most likely achievable step, making a humanoid equivalent to a human's apprentice. The next step, of planning the use of tools to solve particular problems, is most likely to be accomplished through relation (bolts are loosened with a wrench, screws are loosened with screwdriver). Improvisation (e.g. use a nearby wrench as a hammer, or go find a hammer) will remain the (often bad) domain of humans for many years to come.

But there is hope because the most sophisticated of today's humanoids, such as NASA's Robonaut, are already able to use tools while under the direct control of teleoperators [10]. The fact that the electromechanical system is physically able to use wrenches, screwdrivers, power drills, scissors, tweezers, forceps, syringes, scoops, sensor probes and complex EVA tools implies that the addition of a cognitive capability could bootstrap these existing machines into the productive human assistant that NASA desires. While they are currently used mostly through telepresence immersion of a human operator, this automation goal was in the initial design of the Robonaut systems. Automation for dexterous humanoids is today where mobile robot navigation was decades ago- open frontier for new architectures and approaches.

2 The Early Robonaut Prototypes

NASA has now completed two Robonaut prototypes. The first one, Robonaut A, was built as an evolving system over many years, as shown in Figure 2, with progression from arms and hands, to an integrated single limb system, to an upright, dual arm, upper torso. Robonaut B was completed in the Fall of 2002, and was built in a single production. These prototypes have been built following a philosophy that will enable future space flight of the design, by making careful design choices on materials.

Fig. 2. Robonaut development history, producing Unit A and Unit B

To save cost, many components have been selected based on there being flight qualified options available, though the prototypes use the less expensive, non-flight versions. Other subsystems, like arm joints, have been tested in thermal vacuum

chambers at JSC and found to work well across the extreme ranges of temperatures found in space. The goal has been to create a design that will need little change for flight, with only changes in procurement and assembly processes.

One of the fundamental goals in designing Robonaut was to allow the robot to accommodate existing interfaces built for humans. This reduces the need for a second set of tools for robot and human, and allows the robot to work in existing spacecraft. While more expensive than a robot designed for a robot-friendly environment, this is not an option since the spacecraft already exist, and are designed for humans. Figure 3 shows Robonaut A handling general objects designed for human use.

Fig. 3. Robonaut A being teleoperated in the handling of the zippers and Velcro on a backpack, and handling a squeeze bottle for dispensing liquids

2.1 Robonaut A

Figure 2 shows the evolution of Robonaut A. The primary work on A during 2001 was the early work on autonomy, with the addition of fabric coverings and the first generation of tactile skins. Since 2000, Unit A has been over 90% available, and has been used in numerous autonomy, human-robot interaction and telepresence studies.
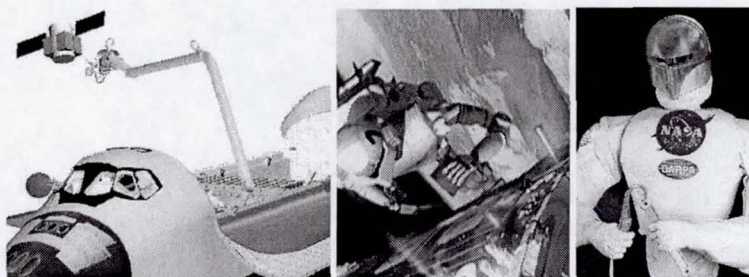
Fig. 4. Concepts of Robonaut being used on a Shuttle Remote Manipulator System, Climbing on the International Space Station, and a photo of Robonaut A handling a crew hand rail.

The original subsystems were a 5 fingered hand [11] and 7 degree of freedom arm [12] custom built at JSC in 1998. Integral to this early development was the design

and fabrication of embedded avionics that made the dense packaging of the Robonaut limb possible [13]. A major focus of the Robonaut project has been to design, construct, and control a bifurcating dexterous upper extremity. Robonaut's two upper extremities are seven-degree of freedom (DOF) arms with twelve DOF hands. Additionally, mounted on top of a two DOF neck (in a position analogous to the human head) is a sensor platform containing four cameras and an infrared temperature sensor. Each pair of cameras has two DOF. Despite being primarily focused on the upper body, Robonaut A has an articulating waist with three DOF that provides some amount of mobility. This results in a 47 DOF dexterous robot.

2.1.1 Development of Robonaut A's Hands

The requirements for interacting with space station EVA crew interfaces and tools provided the starting point for the Robonaut Hand design [11]. Both power (enveloping) and dexterous grasps (finger tip) are required for manipulating EVA crew tools. Certain tools require single or multiple finger actuation while being firmly grasped. A maximum force of 20 lbs and torque of 30 in-lbs is required for EVA assembly [12].

It is possible to either build interfaces that will be both robotically and EVA compatible or build a series of robot tools to interact with EVA crew interfaces and tools. However, both approaches are extremely costly and will add to a set of space station tools and interfaces that are already planned to be quite extensive. The Robonaut design will make all EVA crew interfaces and tools robotically compatible by making the robot's hand EVA compatible.

EVA compatibility is designed into the hand by reproducing, as closely as possible, the size, kinematics, and strength of the suited astronaut hand and wrist. The number of fingers and the joint travel reproduce the workspace for a pressurized suit glove. Staying within this size envelope guarantees that the Robonaut Hand will be able to fit into all the required places. The Robonaut Hand shown in Figure 5 also reproduces many of the necessary grasps needed for interacting with EVA interfaces.
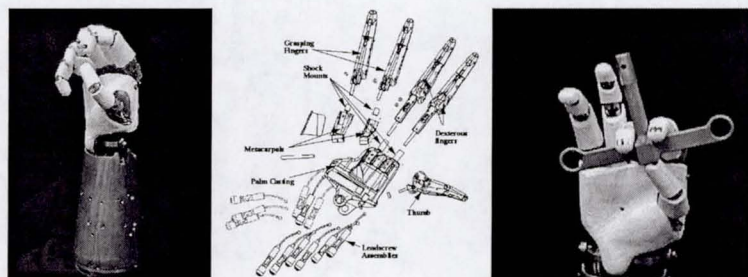

Fig. 5. The Robonaut hand, an exploded CAD model, and demonstrating an unusual grasp

Joint travel for the wrist pitch and yaw is designed to meet or exceed the human hand in a pressurized glove. Hand and wrist parts are sized to reproduce the necessary strength to meet maximum EVA crew requirements.

The Robonaut hand has a total of twelve degrees of freedom. The forearm houses the motors and drive electronics, a two DOF wrist, and a five fingered, twelve DOF hand. The forearm, which measures four inches in diameter at its base, is approximately eight inches long, houses all fourteen motors, twelve separate circuit boards, and all of the wiring for the hand.

In order to match the size of an astronaut's gloved hand and maintain the strength requirement, the motors are mounted outside the hand, and mechanical power is transmitted through a flexible drive train. Past hand designs have used tendon drives, which utilize complex pulley systems or sheaths, both of which pose serious wear and reliability problems when used in the EVA space environment. To avoid the problems associated with tendons, the hand uses flex shafts to transmit power from the motors in the forearm to the fingers.

Overall the hand is equipped with forty-three sensors. Each joint is equipped with embedded absolute position sensors and each motor is equipped with incremental encoders. Each of the leadscrew assemblies as well as the wrist ball joint links is instrumented as load cells to provide force feedback.

2.1.2 Robonaut Upper Arms, Leg and Neck

Robonaut has four serial chains emerging from the body: two upper arms for dexterous work, a neck for pointing the head, and a leg for stabilizing the body in micro gravity. These chains are all built with a common technology, best described as a family of modular joints, characterized by size and kinematic motion type. There are four torque ranges, from 10 ft-lbs to 200 ft lbs, and two motions types, roll and pitch. Other scales have been built for thermal vacuum testing, but currently are not integrated into the system. Figure 6 shows the upper five DOF of the Robonaut arm, several pitch joints, and hip joint compared to the human pelvis.
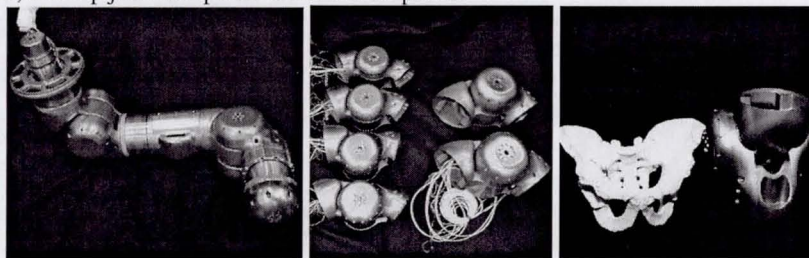
Fig. 6. The Robonaut arm joints are a modular set of roll and pitch joints, currently available in 4 scales including a 3", 4", 5" and 6" version, the 6" shown compared to a human pelvis

2.1.3 Robonaut A's Head

The head on Robonaut is configured to be a sensor platform, resting upon a two DOF neck. Having a head on Robonaut gives human co-workers a feel that Robonaut

is more human-like, allowing it to more readily become one of the team. The shell covering Robonaut's head is a prototype device, based on Roman Centurion headgear, grown using stereo lithography, Figure 7. This shell, however, is more than purely aesthetic; its serves as protection for the numerous sensors that reside inside.
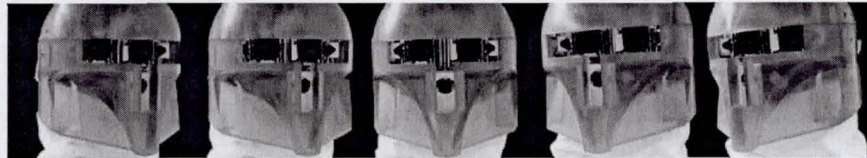

Fig. 7. The Robonaut head is a rich sensor platform that can be pointed by either a teleoperator or under automatic control by the robot's brainstem

The neck provides a range of motion that exceeds human head motion, allowing its 4 computer controlled cameras to be pointed and zoomed to provide viewing during manipulation. The head also contains an infrared, laser guided thermometer for estimating the temperature of objects in space, as well as audio communication interfaces for working with humans. The most noticeable sensors within the Robonaut head are its four cameras. The primary cameras are set apart a distance approximately the same as human eyes, allowing for stereo vision for both teleoperation and machine vision. Each of these cameras has two DOF, zoom and pan. The pan degrees of freedom for both cameras allow the combined system to set the verge point for stereo. This feature is especially useful for teleoperation where the teleoperator can optimize the stereo effect for a particular task. The secondary cameras are set farther apart and have a wider field of view.

EVA requirements state that hardware shall be designed to withstand temperature extremes in the -200 to +250° F range. As a result, the ability to sense the temperature of an object before coming into contact would be of great benefit. To this end, Robonaut has a low cost off-the-shelf infrared temperature sensor as part of its sensor suite. The temperature sensor has a low-power laser used to guide the sensor that can also be used to calibrate the verge mechanism in the eye cameras.

The last components in the Robonaut head are a pair of stereo microphones and a speaker. While this arrangement would not be available in the vacuum on space, its utility has shown the necessity of communications between the robot, its human partner and a teleoperator.
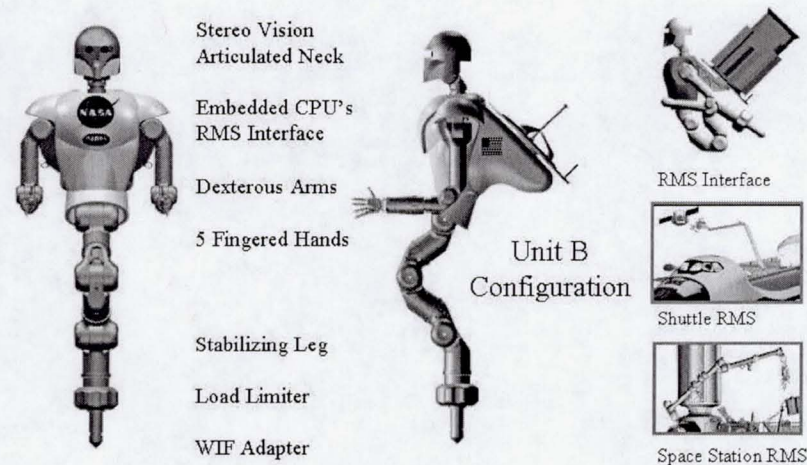

2.2 Robonaut B

During EVA, crewmembers often place both legs into a portable foot restraint connected to either a crane or a ground point on the spacecraft. This creates a path to react the loads generated during work. In its microgravity configuration, Robonaut B works on the same principle. Instead of a foot restraint to react loads, Robonaut B has a single seven DOF leg with the same interface to the spacecraft foot restraints as an end effector, Figure 8. Once anchored to a spacecraft, the multi-jointed leg provides a

greater amount of body mobility than even a human crewmember. Additionally, Robonaut B is designed to have a grapple fixture on its back, allowing it to be maneuvered by an RMS. When grounded in this manner, the leg is a third arm.

Robonaut B has this added 0g stabilizing leg for use on the exterior of spacecraft, a fully integrated avionics suite in its torso, an interface on its back for the larger Space Shuttle and Space Station cranes (RMS and SSRMS) and added dexterity in its neck.

Stereo Vision
Articulated Neck

Embedded CPU's
RMS Interface

Dexterous Arms

5 Fingered Hands

Unit B
Configuration

Stabilizing Leg

Load Limiter

WIF Adapter

RMS Interface

Shuttle RMS

Space Station RMS

Fig. 8. Robonaut B is shown with labeled anatomy, mounted on an RMS interface, and positioned by both the RMS and SSRMS
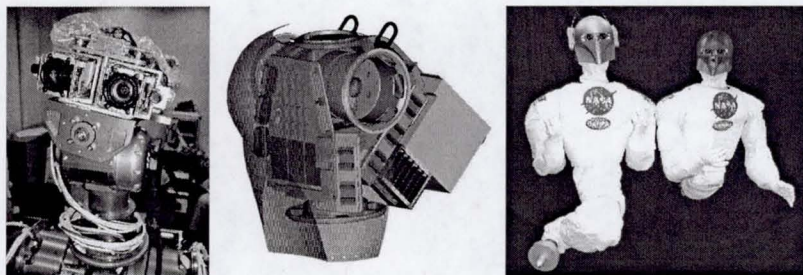
Fig. 9. Robonaut B neck is shown without the helmet, in a pose that slightly rolls the new third axis to its left. Unit B's body is shown in a 3D CAD drawing with transparency to show embedded avionics. Both Robonaut B (left) and Robonaut A (right) are shown side by side

Robonaut B is not merely a copy of Unit A. The team had many new ideas on component technologies, and the challenge of making it portable drove it to new designs. In particular, B has fully integrated avionics, with the large rack behind Robonaut A now miniaturized, and completely packaged in the robot's torso, as shown in Figure 9. A cart makes Robonaut B portable, allowing the team to move the robot out of the lab and transport it off-site. Due to the embedded avionics, the interface to the

robot is now simplified to raw power and Ethernet, a dramatic change over Robonaut A's many hundreds of cables. The embedded avionics also includes a complete, 8 slot cPCI chassis that is unallocated, ensuring that the robot can grow with new autonomy requirements now being researched using Robonaut A.

3 Autonomy Architectures for Dexterous Humanoids

Three guiding principles have grounded the development of Robonaut's autonomy. The first is that any control system must allow and adapt for a broad mix of human inputs, from direct teleoperation, to supervision, to gestures and natural commands expressed by adjacent human team mates. The second philosophy is that the design of the robot and software development strategies must not constrain researchers with limits on their choices for CPU's, operating systems, programming languages, or the use of high level software packages and development tools. The third principle is that developing a tool using system will push cognitive technology to enable other functions such as human-robot interaction, planning and health monitoring.

3.1 An adaptable Architecture

The fact that Robonauts are wired for full immersion teleoperation is a double edged sword. The positives are manifest, in that teleoperation has allowed the robot to be used prior to an autonomy system being available, and has shown that a dexterous robot, if endowed with the cognitive abilities of a human, is able to work with tools and interfaces that are designed for humans. Furthermore, this teleoperation has opened a new opportunity, allowing a human to instruct the robot in a very fundamental manner, as will be discussed later in this section. Other positives include a very thorough checkout of the design, where each Robonaut has been run efficiently for many hours, allowing the system to be tuned and serviced for optimal performance.

But negatives of having a telepresence option exist. The primary one is that choices about changes to the system must not be done in a manner that precludes the teleoperator control modes. Optimal for autonomy may not be optimal for teleoperation, and so sub optimal choices will be inevitable. This has a significant impact on the visual system for the robot. Likewise, an over focus on teleoperation can quickly rationalize away the need for the key sensing modalities that enable automation. The Robonaut team has had the discipline to avoid this pitfall, building a sensate humanoid limb with over 170 sensors per upper extremity, and vision that meets both needs.

Figure 10 shows an adaptable architecture that can be operated in 4 degenerate cases: Operation with an adjacent or remote human, while either working or learning. For all cases, the sensori-motor level I/O is identical for the actual robot or a simulation called RoboSIM [12] that emulates the brainstem interface and API. Section 4 will presents results of each of these cases in early experimental results.
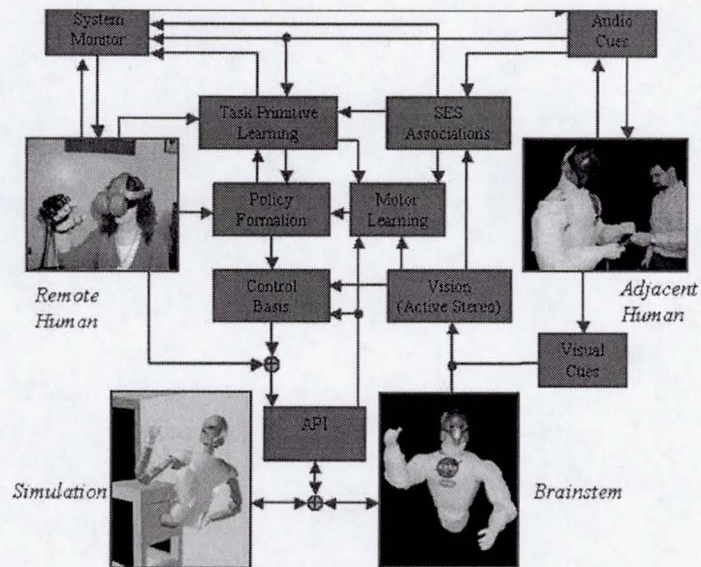
Fig. 10. The general model allows for the robot to operate in a working or learning mode, with either a remote or adjacent human. In any of the four cases, many of the functional blocks in the above architecture are inactive

In the simplest case, the robot is teleoperated by a remote human, in a working mode with no task level or sensori-motor learning. In this mode, shown in the lower right quadrant of Figure 11, most of the functional blocks of the architecture are inactive. This is an example of how the functional blocks of the control system can be on-line, running, but not involved in the particular activity. While not required for the simplest teleoperated case, these other functions can be running, for example observing and tracking humans in the robot's proximity (vision), or keeping track of tools and their locations (SES, Sensori Ego Sphere). In section 4, Figure 15 shows an example of a human working with Robonaut in this teleoperated mode.

The second working case has the robot performing a previously learned task, working with an adjacent human. In this mode, shown in the upper right quadrant of Figure 11, the functional blocks associated with the remote human are not involved, and the task level learning functions are also inactive. Even in this working case, the sensori-motor level learning is still active, allowing the robot to make small refinements in the previously learned skills. Examples of this level of learning might be accommodating for tools that are slightly different from past tools, or adapting to changes in the robot's performance due to age, temperature, or unmodeled factors. This sensori-motor learning approach [15][16] has been demonstrated by Robonaut collaborators at UMass [17] using a humanoid testbed as a precursor to integration with NASA's humanoid. In this control mode, the sensor-motor learning allows a further advancement, in enabling a common experience base and task description be sent to similar, but distinct robots, who can then each refine the policies by adapting to their unique conditions. This resulting robustness across a population of Robonauts allows for

learning and instruction to be shared between robots on earth and in space. In section 4, Figure 15 shows an example of a human working with Robonaut in this mode.
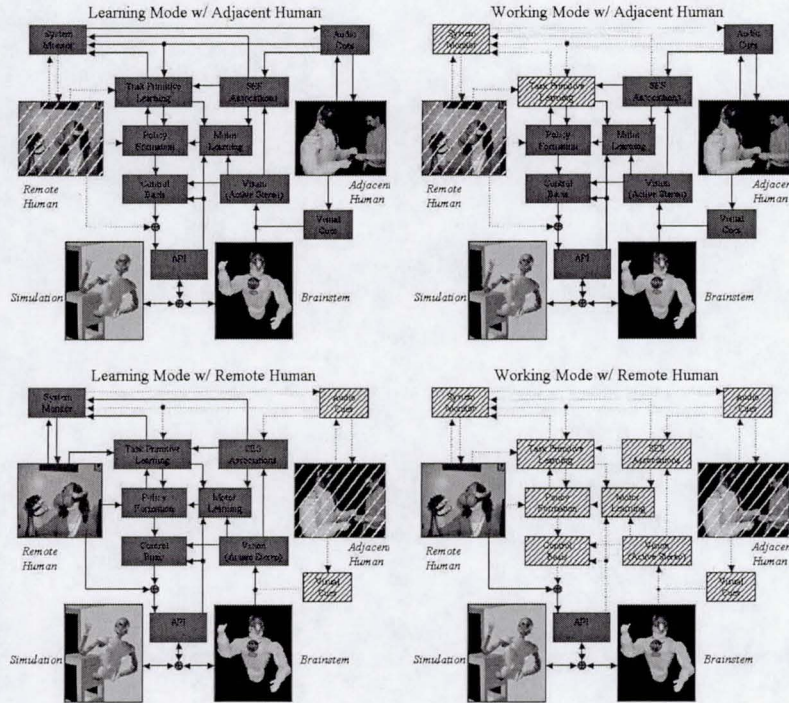

Fig. 11. The four cases of the adaptive architecture

The next two cases involve task level learning, where either a remote or adjacent human can instruct the robot as it acquires new autonomous behaviors. The case shown in the upper left quadrant of Figure 11 has the robot working with an adjacent human. Here, the robot is able to reactively refine its sensori-motor skills, while now also being instructed at the task level. We define task level instruction as a logical, or symbolic sequencing of actions that must be followed in a order subject to some arrangement of conditional branching. An example of this level of instruction is an astronaut telling the robot, "Robonaut, hand me the bolt, and then pick up the wrench". The following steps and functional blocks make this happen:

Task Sequence	Function Block (Fig. 10)	Failure modes
Locate Human	Vision, neck control basis	Keeps looking for human
Hear voice command	Audio	Keeps listening
Look at old object location	SES, neck control basis	Scan for object
Find object pose	Vision	Scan for object
Grasp object #1	Arm / hand control basis	Feel for object
Locate Human	Vision, neck control basis	Keeps looking for human

Hand off object #1	Arm / hand control basis	Feel for human pull
Look at old object location	SES, neck control basis	Scan for object
Find object pose	Vision	Scan for object
Grasp object #2	Arm / hand control basis	Feel for object

In section 4, Figure 16 shows an example of a human working with Robonaut in this autonomous, learning mode.

The last case of Figure 11 has a remote human instructing Robonaut through full immersion using telepresence [18] gear. This form of instruction is not natural for human-to-human coaching, since we cannot "step into" another human and show them how to do a task. While we try to show others that they might learn through imitation, we are limited in that we can't show force, or tactile experience. Humans overcome this with a complex set of mirror neurons, but these leverage, and require, prior experiences, where a robot may be fully naïve [19]. A humanoid fully wired for immersive telepresence control can go beyond these limits, with the human "stepping into" the robot and instructing from within. A human can instruct a robot in a task by directly using the robot, as the machine tracks user commands, and the vast array of sensori-motor experiences that will be used as the robot subsequently controls itself through the very same task.

This sensori-motor learning is based on associations that can be made in the SES developed by Peters [20][21], and related to the specific task. The task space can be segmented into sensori-motor episodes [22] by using the large and multi channel data stream to identify motion terminators shown by Mataric. The key step is to then associate both sensori-motor data, instructor command stream, and non-proprioceptive data such as the visually located object that is causing the sensor event. Peters has shown that this approach [23] can be used to instruct new manipulation skills in a humanoid, and images of early Robonaut experiments are shown in Figure 16.

Under teleoperator control, even in instructional sessions, the robot still has an active refinement function, comparing grasp and other control bases with informed policies from prior experiences. Platt has also shown that grasping can be informed [24] by prior knowledge of the object, and be coordinated with a larger scheme for handling object of many scales by employing more of a humanoid's bifurcating chain system [25]. Figure 16 shows results of early experiments using Robonaut's tactile glove, and simple geometric models limited to primitives alone. The Robonaut team has collaborated with Peters, Grupen, and Mataric and their students on an initial implementation of this architecture, working in all four cases as shown in Figure 11.

3.2 The Humanoid Mind as a Collection of Lobes

Our second development philosophy has been to avoid constraining our researchers with enforced choices on CPU's, operating systems (OS), languages, and other engineering details that curb creativity. In many cases, these choices were made for no good reason, familiarity alone, or even historical reasons that no longer make sense. In other cases, the choices are quite rational, and are proven to be optimal for the specific function at hand. In either case, making a researcher rewrite code for the sake

of commonality, or consistency, puts creativity in opposition to integration. Is there an analogy in nature that provides a solution to this dilemma?

One need only look at brain physiology to see an alternate approach. The mammalian brain is segmented into lobes that are topographically mapped to cognitive, perceptive, interactive and motor control functions. Making the analogy to Robonaut's mind, do the stereo algorithms running on the visual cortex need the same CPU or OS as the database handling routines running on the SES? Clearly the answer is no, as these two functional lobes have very different bandwidth, loop rate and storage needs. Neither lobe's needs is similar to the I/O and real time requirements of the brainstem, where vast amounts of sensor and motor signals are running at hundreds of Hz.

Taking this one step further, does the communication between any two lobes have to be the same as between any two other lobes? Again, the clear answer is no, since the specific data being transferred is so dissimilar, and the data rates vary from low, or asynchronous functions, to high synchronous rates for closed loop control.

From these two observations, Robonaut's mind has been designed as a set of functional lobes that can be connected into "patterns" of activity for specific tasks. For example, when instructing the robot, Robonaut's brainstem (position, force, tactile data and control signals), Robonaut's ear (voice recognition), Robonaut's visual cortex (stereo vision, human and object tracking) and Robonaut's telepresence workstation (human command stream) are all tracked. On reflection, the learning algorithm can compare the broadband data stream and identify which functional lobe was involved, not by prior knowledge, but through an investigation of the data. A "pattern" is therefore the combination of lobes that had active and event based data that were identified during task segmentation [26] of sensori-motor episodes.

Our initial goal of freeing researchers to select and use the computers of their choice has both accelerated their work as well as expanded the Robonaut mind to an evolving set of functional lobes. Where mammalian evolution is measured in eons, the Robonaut team, with its university collaborators, is able to quickly try new functions while leaving old lobes intact. The following functional lobes are now online:

Functional Lobe	Connections	Computer Specifics
Brainstem	Visual Cortex, Ear, Mouth, Nose, Eye, Teleoperator, SES, SIM	cPCI, 4xPPC, VxWorks, RTI Control Shell, C++.
Visual Cortex (stereo, object / human tracking)	Brainstem, SES, Simulation, Policy Learning	PCI, Pentium, Win2000, Visual C++
Ear (voice recognition)	Brainstem, Teleoperator, SES, Mouth	PCI, Pentium, Win2000, Visual C++, ViaVoice
Mouth (voice synthesis)	Brainstem, Teleoperator, Ear	PCI, Pentium, Win2000, Visual C++, ViaVoice
Eye (Zoom, focus, iris)	Brainstem, Teleoperator, Ear	PCI, Pentium, Win2000, Visual C++
Nose (IR thermometer)	Brainstem, Teleoperator, Ear	PCI, Pentium, Win2000, Visual C++
SES	Visual Cortex, Brainstem, Simulation, Policy Learn-	PCI, Pentium, Win2000, Visual C++

	ing	
Grasp Policy Learning	Brainstem, SES	PCI, Pentium, Red Hat, C++
SIM	Visual Cortex, Brainstem, SES	PCI, Pentium, Win2000, Visual C++
Telepresence Workstation	Brainstem, Ear, Mouth	2x Laptop, Pentium, Win2000, Visual C++

### 3.3 Driving Humanoid Cognition Towards, and Through, Tool Use

Few animals meet the mark as tool users [7][8][9]. Animals that have great prowess in navigation, such as mosquitoes that can fly between raindrops, have little or no self-awareness. Their understanding of the world is limited to seek and avoid missions- tasks they do quite well and with grace with amazing (and painful) skill.

A quick look at the functional lobes in Figure 11 shows a structure that allows "patterns" of activity much like human brain imaging can relate cognitive function with brain metabolism or EEG. Sensori-motor learning "lights up" certain lobes, whereas task level execution uses other combinations of the functional blocks. Unlike the linear sense-think-act sequence, the Robonaut architecture is centered on knowledge, the key and missing ingredient in the flight of the mosquito. This knowledge is found in two forms- self-knowledge and world knowledge – but is brought together in tandem in the SES and in the informed policies that develop through the robot's life experiences.

In both cases, the nexus of these associations is in the contact and interaction with the environment. Where the mosquito avoids contact with the environment (a rain drop will kill such a small insect) the mammalian brain, and evolving mind of Robonaut, thrives on these events. The SES associates both the object being touched with the sensor and motor states of the event [21]. The policies that form and are refined during grasping and handling of an object [24] associate a manifold of sensori-motor state spaces, allowing for refinement and even better- choice.

The Robonaut glove and visual cortex are the two most obvious players in this humanoid behavior. It is not surprising that primate classification [27] has identified enhanced binocular vision, and hands with opposable thumbs and flat nails as the key distinguishing features of primates. The Robonaut glove shown in Figure 12 has gone through three generations of development. The latest glove includes over thirty tactile sensors, integrated with textured surfaces and anchored onto the distal finger tip in the rough location of the human finger nail, keeping the fabric (skin) from moving freely relative to the endoskeleton. Tactile arrays across the palmer and sides of the fingers, as well as a more extensive array of tactile sensors on the palm complete the glove. The skin on Robonaut A has now been upgraded with tactile elements up the limbs, and on the torso to allow arm and body grasps, in extended grasps [25] with multiple points of contact. Figure 12 shows some of the tactile arrays just under the subcutaneous layer on Robonaut A's torso.
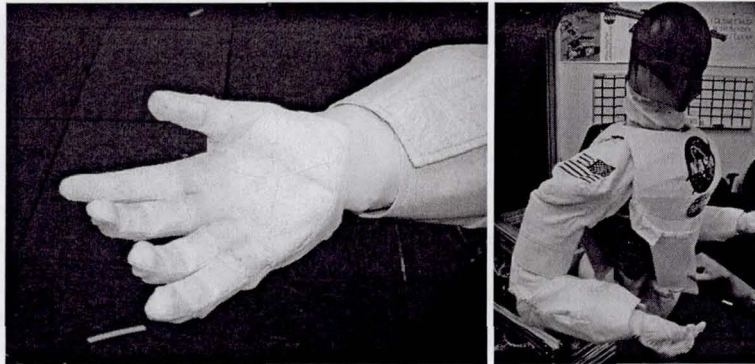
Fig. 12. The Robonaut Glove, and tactile array on arms and torso, under layer of skin

Stereo vision completes the eye hand coordination set that primates have at their disposal. The Robonaut eyes, described in section 2.1.3, meet the primate criteria as binocular, forward looking, and with enhanced depth perception. This visual cortex can build depth maps at arms length that have better resolution than finger scale, allowing for a natural combination of sensori-motor coordination in visual servoing [23].

Figure 13 shows the results of this anatomy, with Robonaut A handling a broad mix of objects and tools. These include a mix of texture maps, depth, and colorized segmented depth.
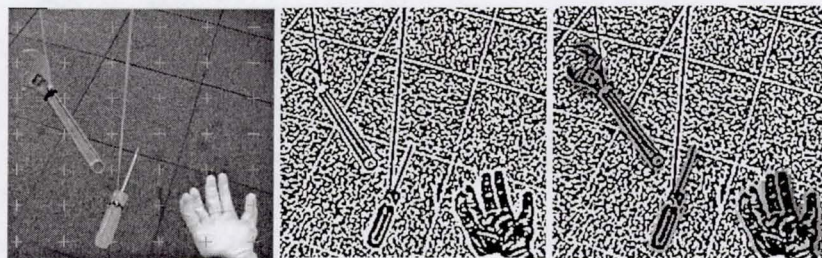
Fig. 13. Stereo depth maps of Robonaut hand and two tools hanging on tethers to emulate 0g

As Robonaut handles objects, the forces imparted are registered in the multi axis load cells embedded in Robonaut's limbs. The light contact along surfaces of the gloves are also detected as a unique signature that is distinct for various tools, and even specific grasps of tools. Figure 14 shows the Robonaut simulation, providing visualization of the tactile sensor states of the hand and glove as a tool is grasped. This grasp signature, along with limb force, position and velocity information, has over 170 channels of data for association with the object being handled. The circular dots on the palmer surfaces of the animated fingers show the sensor state, with habituation and other algorithms online to enhance tactile acuity.
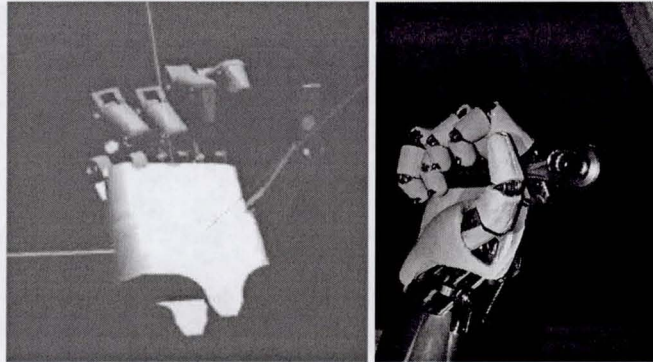
Fig. 14. Robonaut hand on tool, and sensor signature from tactile array

4 Results from Early Autonomy Experiments

Experiments have been completed that exercise all four of the adapted architecture cases presented in section 3.1. The simplest case of teleoperation is shown in Figure 15, where the robot is working under direct human control. This case has information flowing between the remote human and robot, but no learning by the machine.
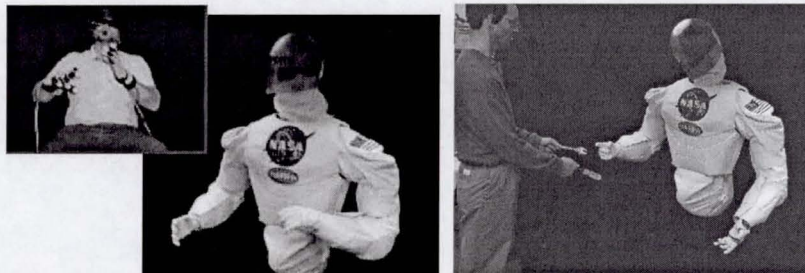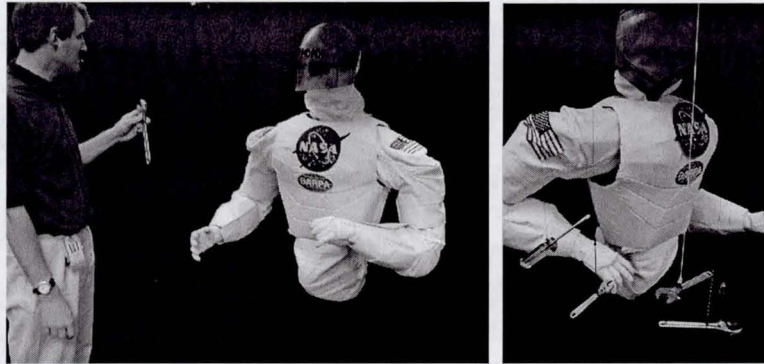

Fig. 15. Robonaut working with a remote and adjacent human, with no task level learning. The right photo has the robot following a previously taught task sequence, initiated by voice reflex, to visually find and point at the requested tool held by the human team mate.

The second case of an adjacent human working with the robot operating in an autonomous mode is also shown in Figure 15. In this case, there is also no learning at the task level, though sensori-motor learning is available for refinement. In this task, the human asked the robot to identify a tool, and track it by pointing at it.

The third case has a human teaching a robot a new task, by composing a sequence of previously mastered policies. In the case shown in Figure 16, the robot is asked to find and then take a wrench from the human's hand.

Fig. 16. Robonaut working with a remote and adjacent human, being taught new tasks. The right photo has the robot being instructed by a remote human to reach out and grasp different tools, each hanging on a tether to emulate the weightless conditions in 0g

5 Conclusions

An adaptable architecture was developed that arranges the brain of a humanoid in a set of functional lobes that can be taught and used to execute prior lessons. The approach can adapt to working with remote and adjacent humans, taking instructions at the task level, or through a novel form of instruction using telepresence control of the machine. By stepping into the machine using VR gear, a remote human can show NASA's Robonaut system how to handle new tools, grasp new objects, or perform other functions that will allow it to assist astronauts in space. Adjacent humans, such as crew working shoulder to shoulder with the humanoid, can direct it through a new sequence of tasks, built upon primitives of previously mastered eye-hand coordination skills.

The robot's interaction with the environment are central to the control and learning approach. Working with DARPA/MARS researchers at UMass, Vanderbilt, USC and MIT, NASA's Robonaut team has implemented the beginnings of an autonomous control system that can work with tools and other objects that are intended for human use. The deep and broad mix of sensors and high dexterity of the Robonaut system has been found to enable a motion segmentation approach, which can take natural human cues to learn new sensori-motor control policies. The architecture is now undergoing further extensions, with functional lobes of the humanoid brains being developed for long term memory, collective cognition between multiple humanoids, and system state monitoring and display to human team mates.

References

1. EVA Catalog: Tools and Equipment, NASA JSC Internal Document (JSC-20466)
2. Whittaker, R., Staritz, P. Ambrose, Robert O., et.al., "Robotics for Assembly, Inspection and Maintenance of Space Macrofacilities", AIAA Space 2000 Conference, September 2000.
3. ITOU Kiyoshi, TAKAHASHI Kohji, "Methodologies for Robot Safety", Vol. 29-02-004. IPSJ Magazine, Tokyo Japan
4. Hans Moravec, "Towards Automatic Visual Obstacle Avoidance", Proceedings of the 5th International Joint Conference on Artificial Intelligence, MIT, Cambridge, Mass., August 1977, p. 584
5. Sebastian Thrun, "A Probabilistic Online Mapping Algorithm for Teams of Mobile Robots", International Journal of Robotics Research, 2001.
6. Cassimatis, Trafton, Schultz, Bugajska and Adams." A Task Domain for Combining and Evaluating Robotics and Cognitive Modeling Techniques". PerMIS 2002, August 13, 2002.
7. Richard Karel, "Tool Using Crow Gives New Meaning to Term 'Bird Brained'", Psychiatric News, March 1, 1996.
8. Renato M.E. Sabbatini, "The Evolution of Intelligence", Brain and Mind Magazine, February 2001.
9. S.T Boysen and G.T. Himes, "Current Issues and Emergent Theories in Animal Cognition", Annual Reviews in Psychology, Vol 50, Pg 683-705, 1999.
10 Robert Ambrose, Christopher Culbert, and Frederik Rehnmark," An Experimental Investigation of Dexterous Robots working with EVA Hardware", AIAA Space 2001, Albuquerque Nm, August 2001.
11 Christopher Lovchik and Myron Diftler, "The Robonaut Hand: A Dexterous Robot Hand For Space," Proceedings of the IEEE International Conference on Automation and Robotics, Vol. 2, pp 907-912, Detroit, Michigan, May, 1999.
12 Robert Ambrose, R. Scott Askew, William Bluethmann, Myron Diftler, S. Michael Goza, Darby Magruder, Frederik Rehnmark, "The development of the Robonaut System for Space Operations", ICAR 2001, Invited Session on Space Robotics, Budapest, August 20, 2001.
13 William Bluethmann, Robert Ambrose, et.al., "Robonaut: A Robot Designed to Work with Humans in Space", Journal of Autonomous Robots, Vol 14, pg 179-197, 2003.
14 Extravehicular Activity (EVA) Hardware Generic Design Requirements Document, 1994. NASA-Johnson Space Center, Houston, Texas. JSC 26626.
15 Elizeth G. Araujo and Roderic A. Grupen. *Learning control composition in a complex environment*. In Pattie Maes, Maja Mataric, Jean-Arcady Meyer, Jordan Pollack, and Stewart W. Wilson, editors, From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior, pages 333-342, Cambridge, Massachusetts, September 1996. MIT Press.
16 Huber, M. and Grupen, R., ``Feedback Control Structure for On-Line Learning," Journal of Robotics and Autonomous Systems, Special Issue on Robotic Learning: The New Wave, Noel Sharkey, Editor, 1997.
17 Wheeler, D., Fagg, A. H., and Grupen, R., *Learning Prospective Pick and Place Behavior*, Proceedings of the IEEE/RSJ International Conference on Development and Learning, 2002.
18 Jenks, Diftler and Williams, "Robonaut: A Telepresence Base Astronaut's Assistant, SPIE 2001, Boston, November 2001.
19 Brooks, R.R. (1996) *"Behavior-Based Humanoid Robotics"* Paper presented at the IEEE/RSJ IROS96.
20 Johnson, C., A.B. Koku, K. Kawamura, R.A. Peters II, "Enhancing a Human-Robot Interfacing using Sensory EgoSphere", Proc. of IEEE International Conference on Robotics and Automation (ICRA), May 11-15, 2002, Washington, DC, 2002.

21 Peters II, R.A., K. E. Hambuchen, K. Kawamura, and Wilkes, D. M., "The Sensory EgoSphere as a Short-Term Memory Humanoids," Proceedings of the IEEE-RAS International Conference on Humanoid Robots, pp 451-459, Waseda University, Tokyo, Japan, 22-24 November. 2001.

22 Ashley Tews, Maja Mataric and Gaurav Sukhatme, "A Scalable Approach to Human-Robot Interaction", ICRA 2003, Taipei, Taiwan, May 2003

23 Richard Alan Peters II, Christina L. Campbell, William J. Bluethmann, and Eric Huber, "Robonaut Task Learning through Teleoperation", ICRA 2003, Taipei, Taiwan, May 2003.

24 Diftler, M., Culbert, C., Ambrose, R., Platt, R., Bluethmann, W., "Evolution of the NASA/DARPA Robonaut Control System", ICRA 2003, Taipei, Taiwan, May 2003.

25 Robert Platt Jr., Andrew H. Fagg, Roderic A. Grupen, "Extending Fingertip Grasping to Whole Body Grasping", , ICRA 2003, Taipei, Taiwan, May 2003.

26 Jenkins, O. C., Matari'c, M. J. & Weber, S. (2000), *Primitive-Based Movement Classification for Humanoid Imitation*, in `Proceedings, First IEEE-RAS International Conference on Humanoid Robotics', Cambridge, MA, MIT.

27 Brian Schwimmer Primate Taxonomy, 1996