

A MOBILE AUTONOMOUS HUMANOID ASSISTANT

M. A. Diftler, R. O. Ambrose, K. S. Tyree, S. M. Goza
Automation, Robotics, and Simulation Division
NASA/ Johnson Space Center
Houston, Texas, 77058
diftler@jsc.nasa.gov

E. L. Huber
Metrica
Johnson Space Center
Houston, Texas 77058
huber@jsc.nasa.gov

A mobile autonomous humanoid robot is assisting human co-workers at the Johnson Space Center with tool handling tasks. This robot combines the upper body of the National Aeronautics and Space Administration (NASA) / Defense Advanced Research Projects Agency (DARPA) Robonaut system with a Segway™ Robotic Mobility Platform yielding a dexterous, maneuverable humanoid perfect for aiding human co-workers in a range of environments. This system uses stereo vision to locate human teammates and tools and a navigation system that uses laser range and vision data to follow humans while avoiding obstacles. Tactile sensors provide information to grasping algorithms for efficient tool exchanges. The autonomous architecture utilizes these pre-programmed skills to form human assistant behaviors. The initial behavior demonstrates a robust capability to assist a human by acquiring a tool from a remotely located individual and then following the human in a cluttered environment with the tool for future use.

Keywords: dexterous robot, mobile, autonomy, hazardous environment.

1. Introduction

Humanoid robots offer great potential for assisting humans with a variety of tasks. By definition, they are designed to perform an ever increasing set of tasks that are currently limited to people. Of specific interest here are tasks that currently require human level dexterity while working in dangerous arenas. To this end, NASA and DARPA (Defense Advanced Research Projects Agency) are jointly pursuing the development of a mobile autonomous humanoid, Robonaut, for use in the hazardous environments of low earth orbit (LOE) and planetary exploration. Robonaut, which can also be teleoperated, is capable of interfacing with Extra-Vehicular Activity (EVA) systems that only have human interfaces and working with the same human rated tools designed for all NASA missions.

Humanoids are a relatively new class of robots. Two of the most well known are the self-contained Honda Humanoid Robot¹ and its most recent descendent, ASIMO² which are able to walk and even climb stairs. Another member of the Japanese Humanoid family also inspired by the Honda robot is the HRP-2P³ which is able to lie down and then stand back up and cooperatively move objects with humans. In the area of upper body capability several prototypes have been built that are designed to work with humans. One of the first, Greenman⁴, showed the benefits of a human teleoperating a humanoid robot. WENDY (Waseda Engineering Designed sYmbiont)⁵ has a full upper torso on a wheeled base and is a prototype for a possible domestic humanoid. Several

humanoids have been designed specifically to explore human-robot interaction. MIT's Cog⁶ and Vanderbilt's ISAC⁷ are both remarkable platforms for such work.

These are all impressive devices, but are still prototypes and of course evolving. Unlike natural evolution, researchers from around the world are experimenting with different techniques to improve their humanoids. Fukuda, et. al.⁸, provide an excellent survey of anthropomorphic robot evolution and suggest three characteristics that are most important towards making a better humanoid: human-like motion, human-like intelligence, and human-like communication.

Initially the NASA/DARPA Robonaut achieved these human-like characteristics solely through a human teleoperator directly controlling the system. Through an incremental process, more of the skills necessary to achieve the human-like capabilities necessary to perform EVA and planetary tasks are being reproduced within the Robonaut's control system. These skills combine new software, sensors, and, most recently, mobility to form the basis for a mobile autonomous assistant.

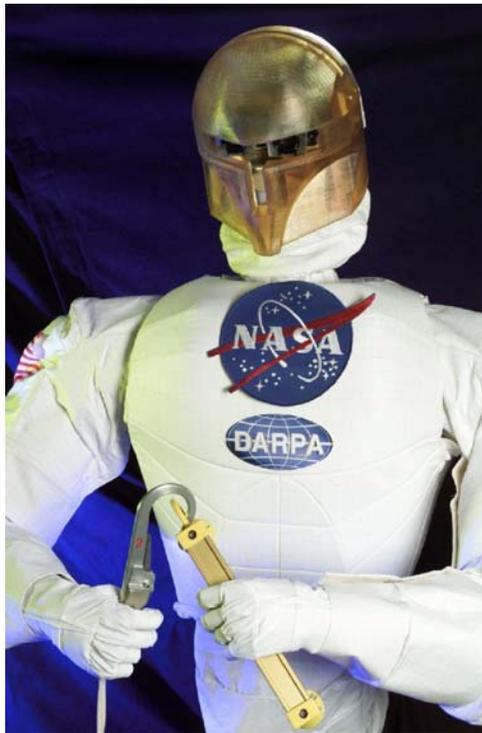


Figure 1: NASA/DARPA Robonaut

2. Robonaut

With 43 degrees of freedom, Robonaut (figure 1) is the first humanoid built for space⁹. It incorporates technology advances in dexterous hands, modular manipulators, and lightweight materials. Robonaut is human size (figure 2), has an articulated waist, and two seven-DOF arms, giving it an impressive work space for interacting with its environment. It has a pan/tilt stereo vision camera unit that provides ranging information

for both teleoperators and machine vision algorithms. In addition to having the correct anatomy to work with EVA equipment, the Robonaut system is designed with space operations in mind. Robonaut's single leg (figure 3) design includes a stinger to directly mate with the same ISS worksite interface (WIF) used by crew for stabilization.



Figure 2: Robonaut – Astronaut size comparison

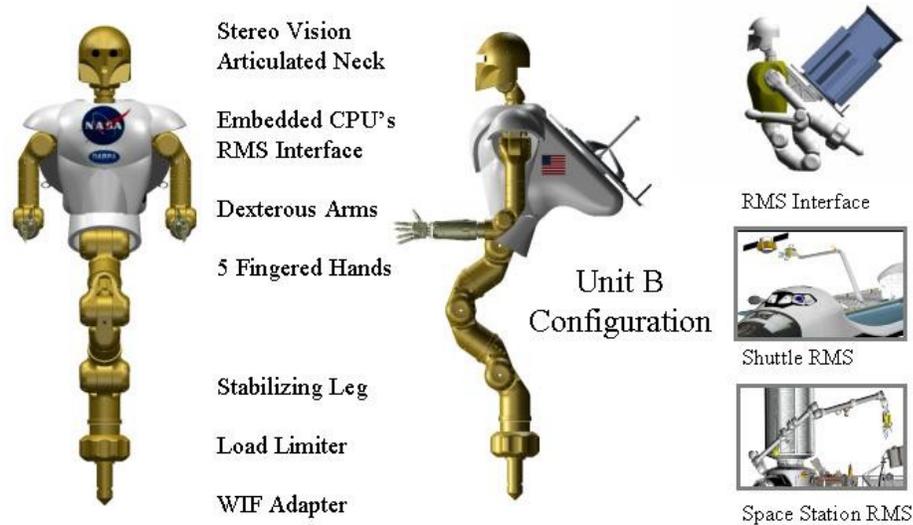


Figure 3. Robonaut in Zero-G Configuration

Robonaut's arms, neck and waist are human scale manipulators designed to fit within EVA corridors. Beyond its volume design, these appendages have human equivalent strength, human scale reach, thermal endurance to match an eight hour EVA, fine motion, high bandwidth dynamic response, redundancy, safety, and a range of motion that exceeds that of a human limb. Both the arms and waist have a dense packaging of joints and avionics developed with the mechatronics philosophy. The arm and waist house thermal-vacuum-rated motors, harmonic drives, fail-safe brakes and 16 sensors in each joint. The arm's small size, strength-to-weight ratio, density, and thermal vacuum capabilities make it the state-of-the-art in space manipulators today (figure 4).

Robonaut's hands set it apart from any previous space manipulator system¹⁰. These hands can fit into all the same places currently designed for an astronaut's gloved hand. A key feature of the hand is its palm degree of freedom that allows Robonaut to cup a tool and line up its long axis with the roll degree of freedom of the forearm, thereby, permitting tool use in tight spaces with minimum arm motion. Each hand assembly shown in figure 5 has a total of 14 DOFs, and consists of a forearm, a two DOF wrist, and a twelve DOF hand complete with position, velocity, and force sensors. The forearm, which measures four inches in diameter at its base and is approximately eight inches long, houses all fourteen motors, the motor control and power electronics, and all of the wiring for the hand. Joint travel for the wrist pitch and yaw is designed to meet that of a human hand in a pressurized glove.



Figure 4: Robonaut arm



Figure 5 : Robonaut hand

In keeping with the biological theme that is at the basis for developing humanoids, automated functions developed for Robonaut are distributed into various control system nodes that are analogous to the human brain's anatomy¹¹. The lowest level functions include: actuator control, motion control, safety, compliance control, tactile sensing, etc. All of these functions are implemented as part of Robonaut's brainstem. Higher level functions such as vision, memory, and grasping are located in other parts of Robonaut's brain. All communication between the distributed control system nodes passes through a well-defined Application Programmer's Interface (API) that is analogous to the thalamus in a human brain. This API is a living interface and accommodates communication requirements for additional capability as the robot evolves. One of the most recent additions includes definitions for mobility commands and data.

3. Mobility

Robonaut can be configured for a variety of mobility options. In its zero-g configuration, it can be moved by a larger manipulator that connects with the grapple fixture on Robonaut's back (figure 3). The arms can be used to climb along truss handrails and the leg can be used for repositioning of the body once it is stabilized to a space structure. On planetary missions, Robonaut can be integrated with a rover to make a modern day centaur (figure 6). For its first mobile autonomous assistant role, Robonaut is combined with a more readily available lower body.



Figure 6: Conceptual Robonaut Centaur Configuration

3.1 Robotic mobility platform

Mobility for Robonaut is provided by the Segway™ Robotic Mobility Platform (RMP), as shown in Figure 7. It is a derivative of the Segway™ Human Transporter (HT). The HT is a two wheeled motorized vehicle for transportation. It is capable of traversing a multitude of terrains. DARPA commissioned Segway™ to develop a computer-controlled version capable of balancing large payloads. This became the Segway™ RMP. The RMP is controlled via computer. Velocity and turning rate are the primary controls. When these values are set to zero, the RMP will hold position even when external forces are applied. The RMP is capable of a zero turning radius and +/- 8 mph velocities. One of these very special devices was delivered to NASA – JSC.

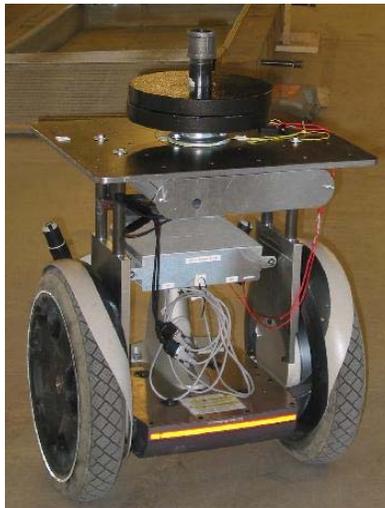


Figure 7 : Segway™ RMP

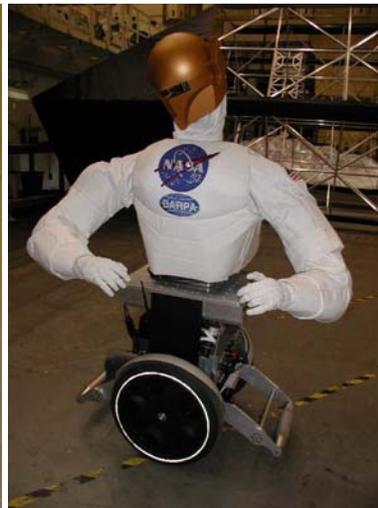


Figure 8: Robonaut waist motion

The RMP was tested extensively to determine its capabilities. Control software was tested and the stock hardware was modified to suit the needs of Robonaut. A battery power distribution system was added and training wheels were added for initial development and testing. One shortfall of a two-wheeled platform is its inability to stay

upright if drive power is severed for any reason. The training wheels prevent this failure mode from causing damage to its robotic payload.

The leg of Robonaut Unit B, except for a single roll joint, was removed in preparation for mounting on the RMP. This joint provides a single waist DOF and allows Robonaut to pivot on top of its mobile platform providing more flexibility to the teleoperator. Figure 8 depicts this DOF. Robonaut on the RMP combines human like dexterity with low profile mobility making for an impressive and capable humanoid¹². Sensing in the form of vision, a laser rangefinder, and tactile feedback provide the data needed to yield the first Robonaut version of a mobile autonomous humanoid.

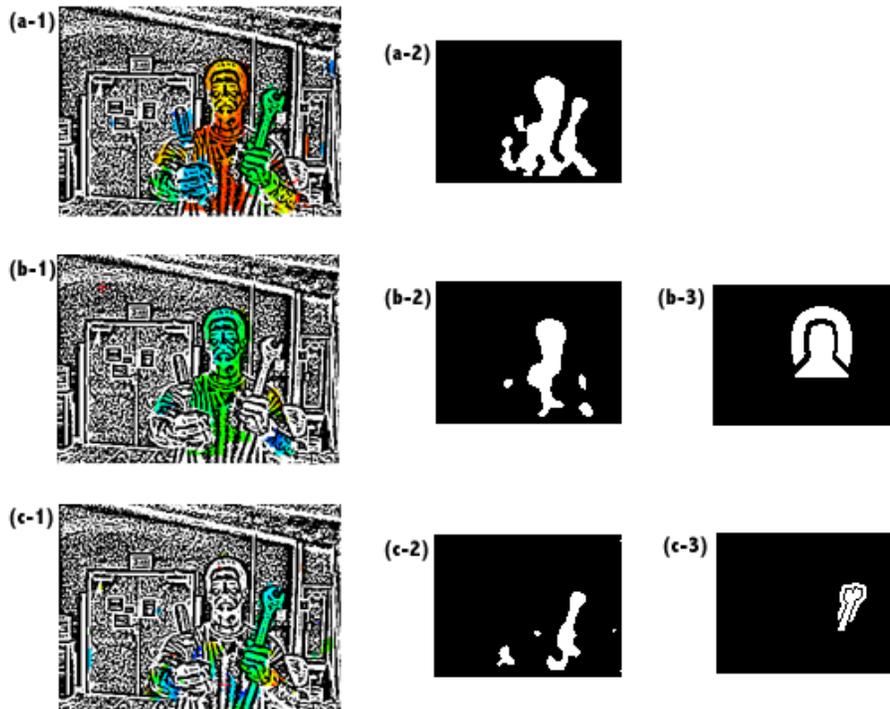


Figure 9: Band Filtered Depth Map

a-1. Color-coded depth map of human, wrench, and screwdriver a-2. Binary depth map of (a-1)
 b-1,2. Human depth band-filtered maps b-3. Matching human head silhouette template
 c-1,2. Wrench depth band-filtered maps c-3. Matching wrench silhouette template

4. Vision

4.1 Philosophy and Scope

The objective for our vision system is to develop object recognition and pose estimation algorithms that aid robotic systems in performing practical tasks in real-world domains. Real-world in this context includes 3D environments which include:

- Multiple objects of different known types
- Objects which are geometrically complex

- Objects possessing free-form contours
- Clutter, both foreground and background
- Specular, dark, bright, or translucent surfaces

Solutions to the challenges in machine vision are strongly tied to the form(s) of data collected from the environment. If the environment is substantially enhanced (adorned with targets, fiduciary marks, artifacts, etc.), simplified, or otherwise modified, then the resulting algorithms will correspondingly stray from the goal stated above. Therefore, it is our belief that to make significant strides in machine vision, it is necessary to limit the target domain to real-world objects and cluttered environments.

We are targeting real world objects with potentially free-form contours. This constraint implies that objects may include: straight lines; corners or intersections conforming to a specific type; lines of constant curvature; long continuous contours; easily detected colors; motion; specific or special shapes such as circles, squares, cylinders, cubes, etc. By the same token, objects of interest in the real-world may possess none of these special features. For this reason we have adopted a doctrine which prohibits reliance on special features for determining pose.

The vision system used by Robonaut employs Laplacian of Gaussian (LoG) pre-filtering making the correlation process sensitive to natural surface (visual) texture. Most objects used by robots provide significant texture over a large portion of their surface. By our approach, we capitalize on this wealth of data by matching to the entire visible surface of the object – i.e. its shape.

4.2 Vision Data Flow

Our stereo-based vision system capitalizes on object shape to track the pose (position and orientation) of well-defined objects, such as wrenches and tables, as well as variable-form structures, such as a human torso and limbs. To achieve tracking, vision operates over a series of stages, one stage cascading into the next. Stages of processing are hierarchical, with each stage typically composed of a series of sub-processes. At the highest level, there are currently three stages of processing:

- 1) Image Server: Takes the stereo video streams from a pair of firewire cameras mounted in Robonaut's head and generates sign-of-LoG (Laplacian of Gaussian) convolved (filtered) image pairs. LoG convolution accentuates any spatial variations (texture) in the grayscale image to promote matching between the stereo image streams.
- 2) Stereo Tower: Takes the binary images from stage (1) and performs patchwise (area) correlation to generate silhouette and range maps in real-time. A range map is a two-dimensional array of distance measurements corresponding to object surface points within a scene. Silhouette maps are binary derivatives of range maps. Each bit of the silhouette map, corresponding to a point in the scene, indicates whether surface material is measured within a specifically targeted distance range. The silhouette map provides a simple means of segmenting out objects of interest from the rest of the scene (See Figure 9).
- 3) Multi Tracker: Takes silhouette and range maps from stage (2) and matches them against large sets of 2D templates for acquisition, and 3D templates for pose estimation.

Pose estimation is the job of Multi Tracker, which is composed of a set of Object Trackers, each seeking a specific type of object. Each Object Tracker is composed of a series of cascading filters, designed to match sets of object-specific templates against incoming silhouette maps.

Match correlation values are simply computed by summing the XOR results between individual binary pixels. By keeping data compact and the operations simple, this approach to matching templates and depth maps is fast. Using the Multi-Media registers available on conventional desktop processors, entire rows of a binary-packed template can be accessed with a single instruction, and bitwise matching can be performed in parallel. The following sections discuss how template matching filters are used to determine object pose in real-time.

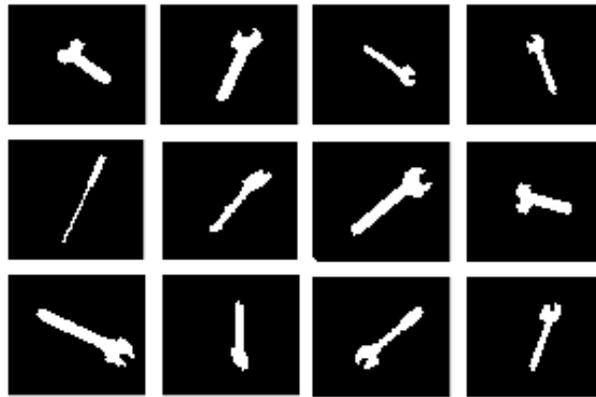


Figure 10: Examples of Wrench Silhouette Templates:
2D binary templates of an adjustable wrench representing its silhouette as viewed from different distances and orientations

4.3 Pose Estimation Method

It is difficult to match the orientation of an object if its position and scale within the scene are not known. Yet it is difficult to apply templates to finding an object's position without knowing what it looks like – and its appearance is dependent on orientation (Figure 10). In short, the problem of template-based, pose estimation is one of bootstrapping.

The Robonaut approach to this problem employs several successive stages of filtering which employ templates. It starts by finding a small set of templates that will likely capture the target object in any pose within the given domain. This set of templates is generic (liberal) in form, and as a side effect, non-targeted objects may also match. Successive stages use templates that are increasingly specific to the target object. As templates become more specific, they increase in fidelity; shapes are sharper, making matching requirements more precise. Upon each stage, foreign (non-targeted) objects are “weeded-out” and only target objects remain (Figure 11).

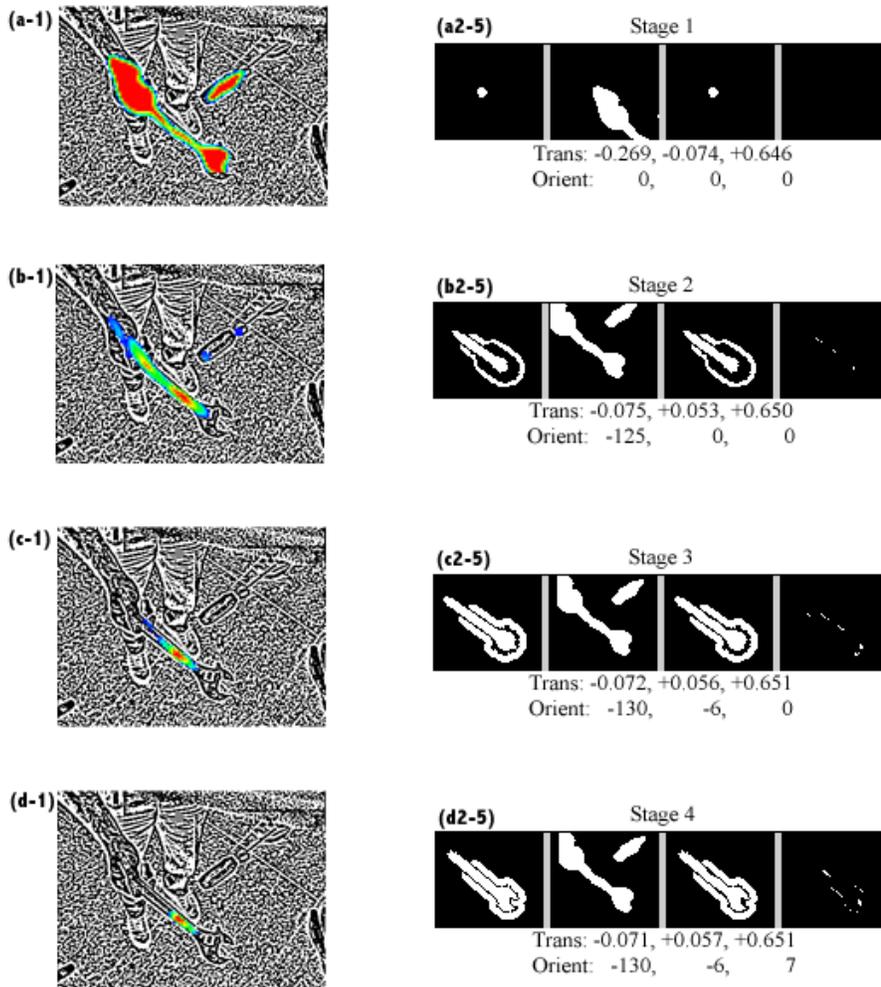


Figure 11: Multi-Stage Pose Matching of Adjustable Wrench
 a-1. Color-coded confidence map for scale (distance (S)) match
 a-2. Template for matching scale at any orientation about Z-Y-X
 a-3. Best match patch from binary depth map
 a-4. Correlation between template (a-2) and patch (a-3)
 a-5. Anti-correlation between (a-2) and (a-3)
 b-1. Confidence map for Z-rotation (in-plane)
 b-2 Template matching S, Z-rotation and any orientation about Y-X
 c-1. Confidence map for Z-Y rotation (in-plane)
 c-2 Template matching S, Z-Y rotation and any orientation about X
 d-1. Confidence map for Z-Y-X rotation (in-plane)
 d-2 Template fully matching the pose of the imaged wrench

High fidelity matching occurs after significant pruning is performed by earlier stages. Many more templates are required to interrogate a candidate location, but only a small fraction of image pixels remain as candidates.

By this method, templates are applied through successive stages to filter out target

object candidates within the scene until only the “true” candidate(s) remain. Template fidelity is increased at each stage to gain an increasingly accurate estimate of object pose. Each stage re-assesses match candidate locations within the scene and passes only the best remaining candidates to the next stage. Each stage narrows down the pose search by at least one degree of freedom. Figure 11 shows a pose estimation sequence for an adjustable wrench.

Each stage of pose estimation employs templates designed to “capture” a specific degree of freedom (component of orientation). A stage must be capable of capturing its target DOF while remaining tolerant to any remaining undetermined DOFs. To achieve this flexibility, early stages must employ liberal silhouettes. Later stages, which have fewer undetermined DOFs, can afford to apply higher fidelity templates, which more accurately reflect the appearance of the target object. In the final stage, the templates are true 2D silhouettes of the target object, providing the greatest pose estimation precision in all degrees of freedom.

5. Navigation

For Robonaut, vision provides the goal and navigation is needed to reach the goal. To this end, obstacle avoidance, mapping and tracking functions were ported to Robonaut from another JSC robot designed for planetary surface exploration, the EVA Robotic Assistant.^{13,14} The modular software from this project was easily transferred to the Robonaut platform, with only a few modifications required. Most of the ported navigation software is self-contained, but a few pieces have to communicate directly with Robonaut: sending driving and steering commands, receiving roll and pitch values, and receiving the human target position from the vision system. To enable the required communication, the navigation software was augmented to communicate with Robonaut’s API.



Figure 12: Robonaut with SICK laser

In order to add obstacle avoidance to Robonaut, a sensor had to be added to detect obstacles. A SICK laser rangefinder was thus mounted on the front of the RMP (figure 12), slanting slightly downward such that the laser's rays hit the ground a few feet in front of the robot. By using the current roll and pitch values of the RMP, the intersections of the laser's rays and the ground (or objects) can be converted into the robot's coordinate system. An obstacle map is generated from the resulting sensor data, filling in "goodness" and "certainty" values for each cell of a discretized 2D map. These values are calculated based on height differences between adjacent laser rays and height values of individual laser rays relative to the robot. Adjustable parameters determine the threshold values which indicate obstacles.

For the EVA Robotic Assistant, the global position of the robot is also taken into account each time the laser data is obtained, based on GPS readings. Each generated laser map is merged into a global map for future use. For Robonaut, working primarily indoors, localization information is not available, so only the instantaneous laser maps are used. Obstacle avoidance in this case is more reactive to the immediate environment, with older data not available.

To navigate to a chosen target – a human in the case of our scenario – two steps are needed. First, the tracking software requires knowledge of the target location. For Robonaut's local-only system (no global information) the target's position is given relative to the robot itself. This information comes from the vision tracking software described above. The mobile base tracking software then takes this relative position and determines the steering and velocity commands needed to travel towards that target.

The second step in the navigation is to check for obstacles. The ground track arc that would result from the desired driving command is overlaid on the laser obstacle map, to determine if the arc intersects any obstacles. If not, the driving command is allowed through and the robot proceeds on its course. If the arc does intersect an obstacle, then a modified driving command must be generated. A set of alternative arcs is defined for the robot ahead of time, and all of these arcs are overlaid on the laser obstacle map to check for obstacle intersections. Those arcs that are safe are also compared to the desired arc, and the closest safe match is selected. Point turns are also considered, if needed, as the RMP is quite adept at turning with a zero turning radius. The modified arc (or point turn) is then sent to the RMP, via the tracking software and the API, and the robot proceeds toward the target, avoiding any obstacles. The entire arc evaluation process occurs with a 5 Hz frequency, so once an obstacle is avoided, the robot can proceed back on course straight toward the target.

As mentioned earlier, parameters can be set to define what constitutes an obstacle for a given robot. The correlation between driving commands and ground track arcs are also parameterized and are set specifically for a given robot. These correlations had to be defined for Robonaut and the RMP at the start of the project. Additional parameters are also available to help adjust for the size of the robot, the desired buffer zone around the robot, the amount of space in which the robot can maneuver, and the existence of global localization. All these parameters are kept in a configuration file, to be read into the software at run time. Most of the parameters can be adjusted during run time as well.

6. Grasping

The mobile Robonaut employs a tactile glove (see figure 13), for autonomous grasping.¹⁵ This glove is instrumented with 19 moderate resolution force sensors. These sensors measure a resistance change that is converted to a force through a voltage divider and a calibrated A/D interface. Each finger joint is equipped with one sensor, and the thumb

has multiple sensors to distinguish between the different thumb contacts that can be achieved. Three sensors are strategically located across the palm, and are very useful for determining contact when making tool and power grasps. In addition to providing good tactile data, the glove is rugged and designed to protect the sensors, provide excellent gripping surfaces and take the abuse associated with doing a wide range of EVA and planetary tasks.



Figure 13: Robonaut tactile glove



Figure 14: Reflexive grasp

A “grab reflex” commands the fingers to autonomously close when tactile stimulation is detected on the glove’s palm sensors (figure 14). Upon contact with the object, the fingers continue to apply force. This is similar to the distal curl reflex observed in human infants. This simple primitive is one of many grasp primitives¹⁶ based on tactile and force feedback available to build autonomous sequences.

The “grab reflex” along with the vision based primitives and the navigation routines form the basis for Robonaut’s initial mobile autonomous behavior. A sequencer combines these capabilities and coordinates Robonaut’s actions.

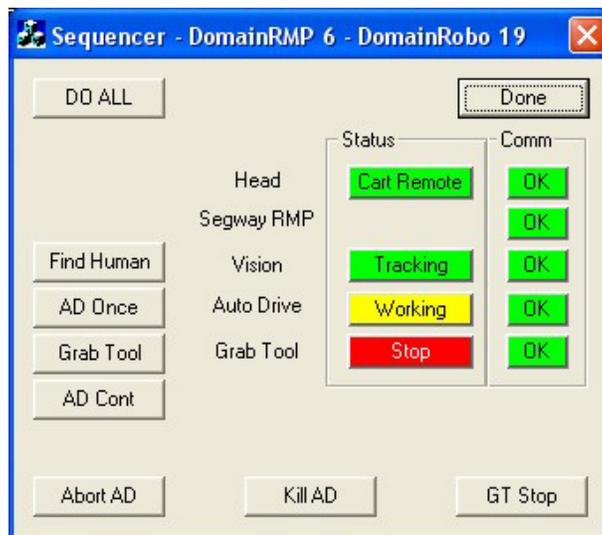


Figure 15: Sequencer Display

7. Sequencer

The Sequencer monitors the status of the various processes and controls the activation and deactivation of various autonomy routines through Robonaut API calls. If

communication is lost during part of the run or before a process is started, then the Sequencer flags the error and does not allow the routine to progress. In order for this to occur, all processes produce a heartbeat. This heartbeat is a simple data packet that emits a status at a constant deterministic rate.

The Sequencer also acts as part of the safety system. Processes can be aborted at any time with buttons located on the dialog. In addition, if the routine aborts mid stream, the Sequencer allows the operator to restart processes and continue from the abort point.

In normal operation, the operator presses one start button and the Sequencer coordinates the entire autonomous behavior. It starts and stops processes as required without operator intervention until the routine is complete. While it is running, the buttons on the dialog display red, yellow, and green colors and text status describing the state of the various components. Red indicates an abort or not ready state. Yellow means that work is in progress. Green indicates the process is ready to take action.

The Sequencer provides the operator with a simple means of monitoring system status and current state, figure 15. The color coding gives a quick visual indication of state and where problems exist. If a problem occurs, the Sequencer allows quick recovery.



a: Avoid obstacles

b: Approach human

c: Acquire tool



d: Stow tool

e: Follow human to worksite

Figure 16: Sub-behaviors

8. Initial Behavior

Using the sequencer described above, Robonaut is able to combine basic behaviors to form its more complex mobile autonomous behavior. In this case the task is to act as a tool handling assistant with the following goals:

1. Scan the room searching for human heads and acquire one as a goal.
2. Proceed towards the human.
3. Avoid obstacles along the way.
4. Stop in front of the human.
5. Scan for a tool in the human's hand.
6. Follow the human hand as the tool is positioned for hand off.
7. Reach out for the tool.
8. Close hand upon making palm contact with the tool and maintain a secure grasp.
9. Move the acquired tool to a stowed position.
10. Reacquire the human head.
11. Follow the human with tool to work site.

Figure 16 shows various sub-behaviors that make up this task. Robonaut has successfully performed this task more than 50 times in multiple environments. While mobility provides exceptional benefits: a large work area, an unlimited number of potential tasks, etc..., it also comes with a number of associated challenges: uncontrollable lighting, unpredictable obstacles, and radio frequency interference.

8.1 Lessons learned

These challenges were apparent when Robonaut performed this task at the 2004 DARPA Tech conference in Anaheim, California. Lighting conditions varied significantly throughout a three day period during demonstrations. Parameter and camera tuning were required and in certain lower lighting situations, the human leader was required to stay closer to the robot so that it could maintain visual contact. This data is being fed back into future camera selection criteria and vision development priorities.

Visitors to the Robonaut demonstration area would occasionally be picked up by the vision system as the human teammate. The actual teammate could easily regain Robonaut's attention by stepping between the robot and the visitor and becoming the most prominent human in the scene. A console operator monitoring the vision system could provide the same correction, by momentarily overriding the visually controlled head control system and redirecting Robonaut's view. A similar vision challenge involved mirrors where Robonaut mistakenly would pick up the reflection of its teammate. Reflective surfaces can be found in a variety of space and planetary applications, and a mobile autonomous humanoid assistant must be able to deal with this challenge along with other confusing scenery. Additional sensing that will augment vision to help identify image features is currently being investigated.

Another issue not normally associated with well structured space and planetary activities is radio interference. Using a robotic assistant in environments where both digital and analog wireless communications can not be controlled obviously requires more care than is normally provided in a laboratory environment. When Robonaut's wireless Ethernet system is overwhelmed by interference on its current operating channel, the robot will safe itself and wait for communications to clear up. This is a reasonable interim solution while more fault tolerant communications are investigated.

While the single SICK laser rangefinder works well for this task, this initial implementation has limitations. With no global localization, a persistent map is not available in case the robot needs to back up. The laser rangefinder also only has a single scan line. Objects below the beam of the laser cannot be seen. Even if the Robonaut originally sees a short object, if it gets too close, or turns suddenly to face the short object

at close range, the laser cannot detect the obstacle. A related challenge dealt with point turns. When Robonaut found itself too close to an obstacle, the software determined that a point turn was necessary. However, without a persistent global map, the robot soon turned so far that the obstacle no longer appeared in the laser scan. The software would then decide to turn back the way it had come, generating an oscillating behavior. Modifications had to be made to correct this problem, insisting that a point turn be carried out long enough to clear the obstacle completely before allowing a different point turn or driving direction to be selected.

9. Conclusions and Future Work

Robonaut has successfully demonstrated its initial abilities as a mobile autonomous humanoid assistant. Robonaut's laboratory tactile and grasping capabilities easily transferred to a mobile base. Navigation software using a laser rangefinder developed for NASA's EVA Robotic assistant also integrated well with Robonaut's RMP base. The most challenging transfer from the laboratory to the field involved stereo vision, a demanding field in any arena. Much was learned from this integration experience and is being used to plan future work.

Work is underway to provide additional sensing for Robonaut's RMP base. A rear facing laser rangefinder is being considered along with a pan-tilt unit that could produce a more complete environmental map. Other proximity sensors that will provide coverage between the lasers are being investigated and will be integrated soon for evaluation.

New firewire cameras that provide a global shutter which exposes all image lines at the same time are being tested. By capturing the image in such a fashion, camera motion does not bias data within each image. This is expected to significantly improve stereo performance.

A next generation tactile glove with significantly more coverage will replace the existing one and provide data for more complex grasps. Additionally, force/torque sensors located within Robonaut's forearms will provide data to better communicate intent between the human and the robot.

All of these additional capabilities will lead to a more robust version of Robonaut's initial autonomous behavior and augment the basis for new behaviors.

Acknowledgements

NASA and the Mobile Autonomous Robot Software (MARS) program in the DARPA Information Processing Technology Office (IPTO) sponsored this work.

References

1. Hirai, K. et al., The development of Honda Humanoid Robot. *Proceedings of the IEEE International Conference on Robotics and Automation*, Leuven, Belgium, 1321-1326, 1998.
2. Sakagami, Y. et al., The Intelligent ASIMO: System Overview and Integration, *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, 2478-2483, October, 2002.

3. Yokoyama, K. et. al., Cooperative Works by a Human and Humanoid Robot, *Proceedings of the IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, 2985-2991, September, 2003.
4. Shimamoto, M.S., TeleOperator/telePresence System (TOPS) Concept Verification Model (CVM) Development, in Recent Advances in Marine Science and Technology, '92, Saxena, N.K., ed., Pacon International, Honolulu, HI, pp. 97-104.
5. Morita, T., Iwata, H., Sugano, S., Development of Human Symbiotic Robot: WENDY. *Proceedings of the IEEE International Conference on Robotics and Automation*, Detroit, MI, 3183-3188, 1999.
6. Brooks, R.A., Breazeal, C., et. al., The Cog Project: Building a Humanoid Robot, Computation for Metaphors, Analogy, and Agents. C. Nehaniv (ed), Lecture Notes in Artificial Intelligence 1562. New York, Springer, 52-87, 1999.
7. Peters, R. A., et. al., A Software Agent Based Control System for Human-Robot Interaction. *Proceedings of the Second International Symposium on Humanoid Robot*, Tokyo, Japan, 1999.
8. Fukuda, T., et. al., How Far Away is "Artificial Man"?, *IEEE Robotics and Automation Magazine*, 7(3), 66-73, 2001.
9. Diftler, M., Ambrose, R., "Robonaut: A Robotic Astronaut Assistant", *Proceedings of the 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Montreal, Canada, June, 2001.
10. Lovchik, C.S., and Diftler, M. A., "The Robonaut Hand: a Dexterous Robot Hand for Space", *Proceedings of the IEEE International Conference on Robotics and Automation*, Detroit, Michigan, 907-912, 1999.
11. Bluethmann, W., et al., "Robonaut: A Robot Designed to Work with Humans in Space", *Autonomous Robots*, Kluwer Academic Publishers, Hingham, MA, v. 14, no. 2/3, 2003, pp. 179-197.
12. Goza, S., Ambrose, R.O., Diftler, M., Spain, I., "Telepresence Control of the NASA/DARPA Robonaut on a Mobility Platform", *Proceeding of the Conference on Human-Computer Interaction*, Vienna, Austria, April 2004.
13. Burrige, R., Graham, J., Shillcutt, K., Hirsh, R., and Kortenkamp, D., "Experiments with an EVA Assistant Robot," *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Nara, Japan, May 2003.
14. Graham, J. and Shillcutt, K., "Robot Tracking of Human Subjects in Field Environments," *Proceedings of the International Symposium on Artificial Intelligence, Robotics and Automation in Space*, Nara, Japan, May 2003.
15. Diftler, M.A., Culbert, C.J., Ambrose, R.O., Platt Jr., R., Bluethmann, W.J., "Evolution of the NASA/DARPA Robonaut Control System", *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, Taipei, Taiwan, September, 2003.
16. Martin, T., Ambrose, R.O., Platt, Jr., R., Diftler, M., Butzer, M., "Tactile Gloves for Autonomous Grasping with the NASA/DARPA Robonaut", *Proceedings of the 2004 IEEE International Conference on Robotics & Automation*, New Orleans, Louisiana, New Orleans, Louisiana, April 2004