HAND-EYE CALIBRATION FOR ROBONAUT

Kevin M. Nickels
Department of Engineering Science
Trinity University
San Antonio, TX 78212
August 8, 2003

Kenneth Baker
Dexterous Robotics Laboratory
Robotic Systems Technology Branch
Automation, Robotics, and Simulation Division
Engineering Directorate
NASA Johnson Space Center
Houston, TX 77058

_____ _____

Kevin M. Nickels Kenneth Baker

HAND-EYE CALIBRATION FOR ROBONAUT

Final Report

NASA Faculty Fellowship Program – 2003

Johnson Space Center

Prepared By:	Kevin M. Nickels, Ph.D.
Academic Rank:	Assistant Professor
University and Department:	Trinity University Department of Engineering Science San Antonio, Texas 78212

NASA/JSC

Directorate:	Engineering
Division:	Automation, Robotics, and Simulation
Branch:	Robotic Systems Technology
JSC Colleague:	Kenneth Baker
Date Submitted:	August 8, 2003
Contract Number:	NAG 9-1526

ABSTRACT

NASA's Human Space Flight program depends heavily on spacewalks performed by human astronauts. These Extra-Vehicular Activities (EVAs) are risky, expensive, and complex. In collaboration with the Defense Advanced Research Projects Agency (DARPA), NASA is developing a robotic astronaut's assistant called Robonaut that can help reduce human EVA time and workload.

This project designed and implemented a hand-eye calibration scheme for Robonaut, Unit A. The intent of this calibration scheme is to improve hand-eye coordination of the robot. The basic approach is to use kinematic and stereo vision measurements, namely the joint angles self-reported by the right arm and 3-D positions of a calibration fixture as measured by vision to estimate the transformation from Robonaut's base coordinate system to its hand coordinate system and to its vision coordinate system.

Two methods of gathering data sets have been developed, along with software to support each. In the first, the system observes the robotic arm and neck angles, and measures the 3-D position of a calibration fixture using Robonaut's stereo cameras, and logs these data. In the second, the system drives the arm and neck through a set of prerecorded configurations, and data are again logged.

Two variants of the calibration scheme have been developed. The full calibration scheme is a batch procedure that estimates all relevant kinematic parameters of the arm and neck of the robot. The "daily" calibration scheme estimates only joint offsets for each rotational joint on the arm and neck, which are assumed to change from day to day. The schemes have been designed to be automatic and easy to use so that the robot can be fully recalibrated when needed such as after repair, upgrade, etc, and can be partially recalibrated after each power cycle.

The system has been implemented on Robonaut Unit A and has been shown to reduce mismatch between kinematically derived positions and visually derived positions from a mean of 13.75cm using the previous calibration to means of 1.72cm using the full calibration procedure and 1.85cm using the "daily" calibration procedure. This improved calibration has already enabled the robot to more accurately reach for and grasp objects that it sees within its workspace. The system has been used to support an autonomous wrench-grasping experiment and significantly improved the workspace positioning of the hand based on visually derived wrench position estimates.

INTRODUCTION

The Dexterous Robotics Laboratory (DRL) at NASA Johnson Space Center (JSC) has developed a humanoid robot astronaut assistant called Robonaut [1]. Robonaut, shown in Figure 1, has been designed as an assistant to astronauts during EVA tasks, and has primarily been teleoperated by a remote human operator. The DRL is beginning to implement several semi-autonomous and fully autonomous controllers for Robonaut, necessitating improved hand-eye coordination for the system. This report documents two methods for gathering kinematic and visual data, and two automatic hand-eye calibration schemes developed for Robonaut in support of these modes.

Figure 1: Ground-based Robonaut system

Prior Work

Much previous work has been done on the self-calibration of redundant manipulators using internal or external kinematics constraints [2-5]. Of particular note is the treatment of Bennett and Hollerbach of a vision or metrology system as an additional kinematic link [5-6], allowing us to treat a one-arm plus vision setup as a closed kinematic chain. This approach allows us to leverage works on the automatic self-calibration of closed kinematic chains, such as [7].

One precondition of this approach is the accurate localization of a point or points of the arm's kinematic chain in the coordinate system of the eyes. Several other calibration schemes utilize special visual markers [4] or LEDs to localize points: we opted for a spherical calibration fixture for several reasons, described later in this paper. In order to accurately locate the spherical fixture in the image, a generalized Hough transform was used. The generalized Hough transform is described in [8,15].

The developed system is a closed-loop system. It reduces the errors between visually and kinematically derived predictions, but does not necessarily adjust these parameters to

match the workspace. If the visual system is not well calibrated, via for example, the procedures described in [9-11,19], the kinematic adjustments performed by this process will not correlate well with workspace positions.

Task Background

Robonaut has historically been operated by a human teleoperator. The DRL is increasing the autonomy level of the tasks performed by Robonaut. This includes, for example, the autonomous modification of previously trained behaviors such as wrench grasping. In this experiment, described in more detail later in this report, the teleoperator grasps wrenches in several different locations in the workspace. Robonaut then visually observes a wrench in a new location in the workspace and modifies and combines the trained behaviors to grasp this wrench. This task obviously requires good hand-eye coordination.

One unfortunate aspect of the design of the arms for Robonaut Unit A is that all joint encoders are relative. As the arm is powered down in the evening and restarted the next morning, the positions of all joints on the 7DOF right arm and the 2DOF neck can change, leading to errors in self-reported joint angles. Errors in these angles, as well as uncertainty in the as-built kinematic parameters of the arm, have lead to workspace errors of up to 10-15cm in various situations. While human teleoperators are very good at correcting for this type of systematic error, it is unacceptable for the degree of autonomy now being required of Robonaut.

The goals for this project were to develop a closed-loop automatic kinematic-visual calibration scheme for Robonaut Unit A. Two variants on this scheme were developed. In the first, a set of measurements is used to estimate all relevant kinematic parameters of the robot. In the second, a much smaller set of measurements is used to estimate only the offsets in reported joint angles for the arm and neck. The intent is to perform the full calibration when the unit is upgraded, changed, repaired, etc, and the partial calibration after each power cycle.

These procedures are integrated with Robonaut's visual cortex, but can also be used to update the procedures used to perform motion planning and execution for the arm when it is being commanded by visually derived data.

Kinematic Model

Robonaut's arm is a redundant manipulator with 7 degrees of freedom. This manipulator can be described by 7 homogenous transformations A_j from link j to link j-1 as defined by the Denavit-Hartenberg (DH) convention. There are two common structures for the definition of Denavit-Hartenberg Parameters (DHPs): one involving a screw about the z_i axis followed by a screw about the x_i' axis (the rotated x_i axis) [12,13] and one involving a screw about the x_i axis followed by a screw about the z_i' axis [14]. In addition to the

screw order, these systems differ in their conventions for placing coordinate axes relative to links. DHPs for equivalent manipulator systems will therefore differ in the two structures. In this work, we utilize the former, where $A_j = Trans(x,a_j)Rot(x,?_j)Rot(z,?_j)Trans(z,d_j)$, *Rot* implies a rotation about an axis, and *Trans* implies a translation along an axis. The position and orientation of the last link can be computed by a sequence of DH transformations defining the kinematic model $T_c = A_1A_2A_3\ldots A_{nf}$ where *nf* is the number of degrees of freedom. Both the 7DOF transformation from the chest coordinate system to the hand coordinate system and the 2DOF transformation[1] from the chest coordinate system and the eye coordinate system are parameterized in this way. Figure 2 shows the coordinate systems used in Robonaut.

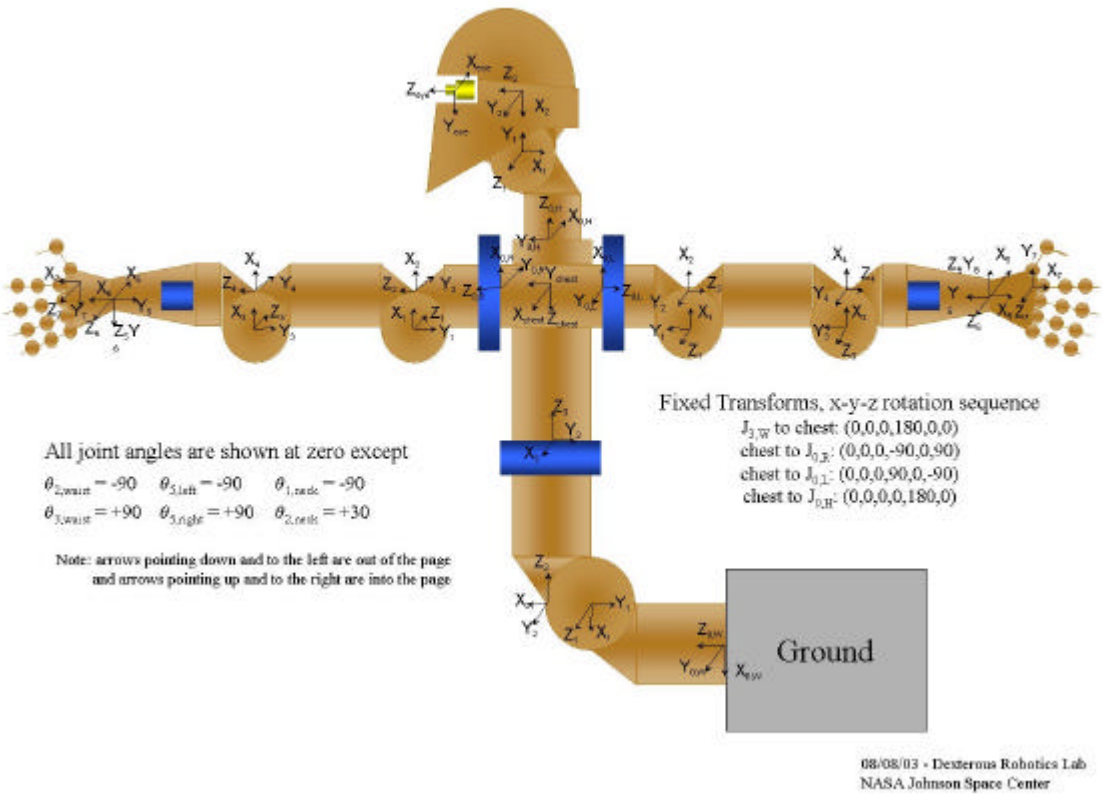

Figure 2: Robonaut's Coordinate Systems

[1] For simplicity, and to allow for automatic calibration of the helmet-camera transform, we use 3 degrees of freedom in the chest-head transform. On Unit A, the joint angle will always be zero for the third DOF, but Unit B has active head roll as well as pitch and yaw.

CALIBRATION PROCEDURES

We have developed two procedures for gathering calibration data from the robot. In the first, the robot is observed under external control and data are logged. In the second, the robot is actuated to each of a set of prerecorded target configurations, and data are again logged. We have developed two variants of the procedure to establish a set of DHP values. In the daily calibration variant, estimates are generated only for the joint angle offsets. In the full calibration variant, estimates are generated for all relevant DHPs.

Supporting Programs

We have developed several computer programs to support these calibration procedures. This section reviews the programs and foreshadows their use in the calibration process. Throughout the user interfaces, the results of image measurements are shown in red (or gray, if this report is printed in grayscale). Blue/dark gray markings indicate kinematically derived predictions using the previous calibration, and are shown in the left image of the split display. Green/light gray markings indicate kinematically derived predictions using the updated calibration, and are shown in the right image of the split display.

The GenVisCalData (Generation of Visual Calibration Data) program, whose display is shown in Figure 3, is used to observe the robot under external control, and to log relevant kinematic and visual data. In the background image, Robonaut's hand can be seen holding the calibration fixture. The red/gray circle represents the continually updated visually located position of the ball in the image (this measurement is actually a 3D measurement based on depth-from-disparity over the image.) The blue/dark gray circle in the left frame represents the projection of the calibration fixture into the image as predicted by the arm and neck kinematics (using the as-built kinematic parameters). The green/light gray circle in the right frame represents the projection of the calibration fixture into the image as predicted by the arm and neck kinematics (using the updated kinematic parameters). The program continually queries the robot for the current joint angles, allowing a live qualitative review of the quality of the current calibration.

The SumVisCal (Summary of Visual Calibration) program, whose display is shown in Figure 4, is used to automatically cycle through each configuration of the robot in a data set and update all measurements in the set. It will also summarize a calibration set, and is used to estimate the optimal joint angle offsets for a given data set (the "daily calibration").

For each robot configuration in the set, this display shows two red/gray dots, one in each frame, representing the visually-located positions of the calibration fixture, and a blue/dark gray dot in the left frame and a green/light gray dot in the right frame, representing the kinematically-derived position predictions of the calibration fixture. A

white line connects each pair of dots. If the updated calibration exactly predicted the visual measurements, the pairs of dots in the right hand frame would coincide in all cases.

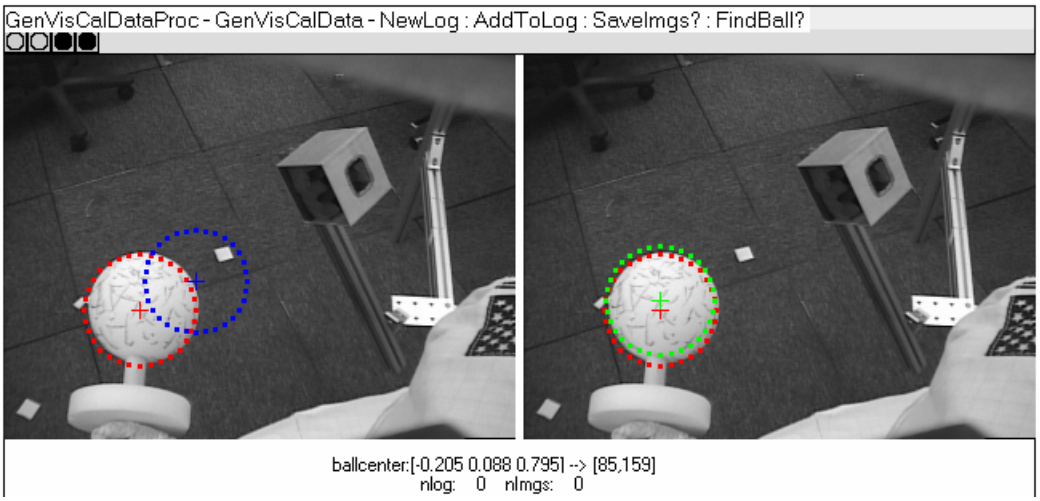

Figure 3: The GenVisCalData Display, ready to log data

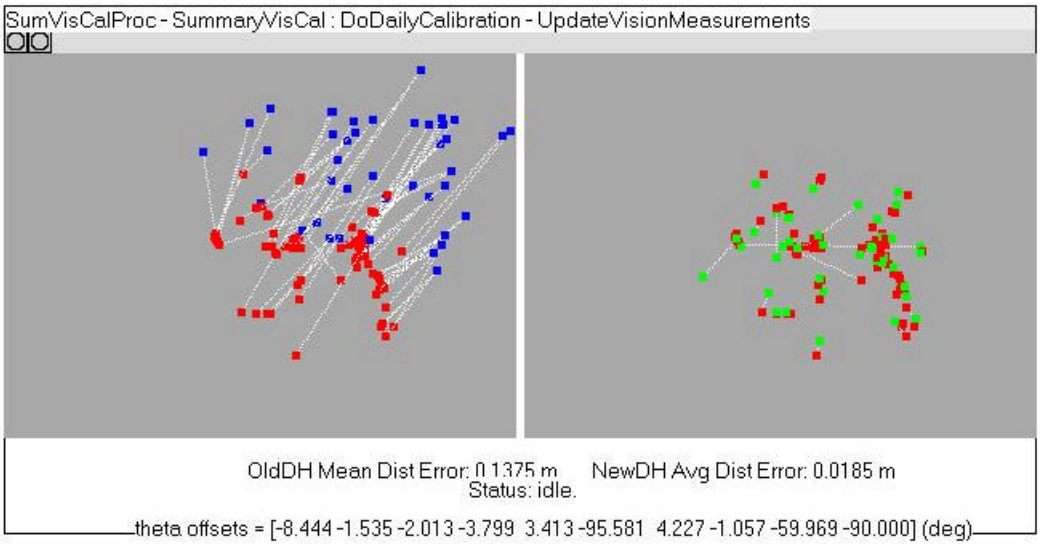

Figure 4: SumVisCal Display showing a typical calibration set summary

The RoboVisCal (Robotic Visual Calibration) program is used to review a stored calibration set frame-by-frame. It is also used either alone or as a slave to SumVisCal to actuate the robot to a particular recorded configuration and update the kinematic and visual measurements for this point in the data set. Figure 5 presents RoboVisCal's display showing a typical configuration that is a member of one of these sets. The graphical conventions are the same as for GenVisCalData.

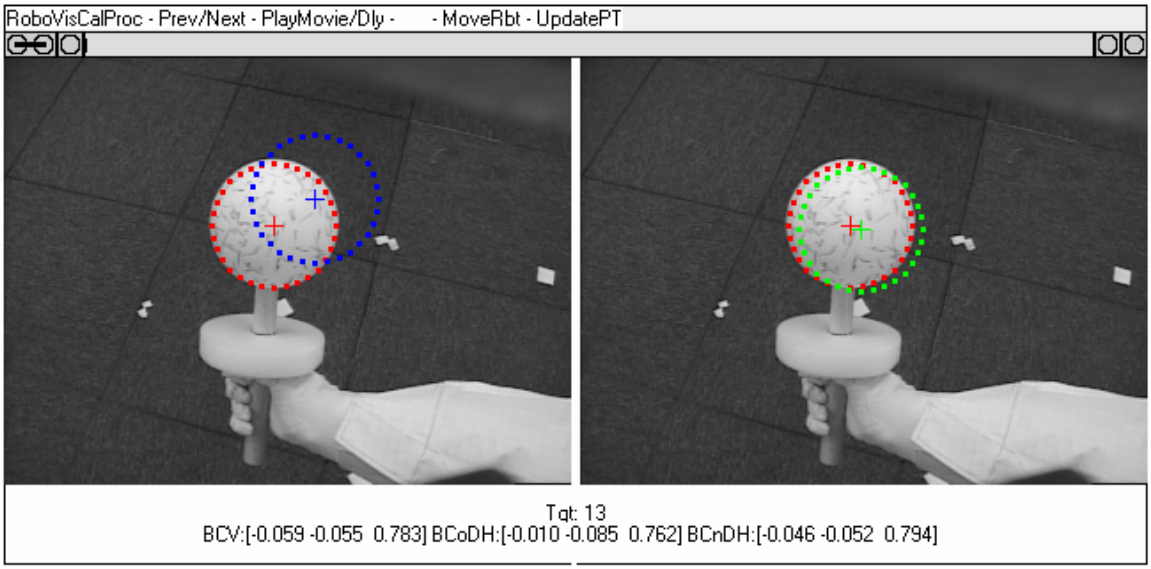

Figure 5: The RoboVisCal Display showing a typical target in a calibration set

Daily Calibration Procedure

A daily calibration procedure has been developed using these software modules. First, a prerecorded set of configurations is loaded into the system. The SumVisCal program is used to drive the robot to each configuration and to update the kinematic and visual measurements. As this happens, the red/gray dots shown in Figure 4 will dynamically update, and the average distance error between the visual measurements and the kinematically derived predictions will be dynamically updated. This process takes approximately 10 minutes for a 65-element calibration set, primarily time to move the robot. A dataset could also be captured using GenVisCalData as described above, but this is more time-consuming.

A Nelder-Mead optimization algorithm is then used to estimate a set of joint angle offsets $?_i$ $i=0...9$ for the arm and neck based on the current set of visual measurements. This process takes approximately 5 seconds per iteration, and can be done repeatedly to improve the estimate. As this is an iterative search with a random initial value, repeated optimizations on the same data may improve the results. The daily calibration thus consists of updating a set of visual measurements, followed by estimating the joint angle offsets. Currently, these estimates are manually input to Robonaut's control system.

Full Calibration Procedure

A full calibration procedure has also been developed using these software modules, in conjunction with some Matlab code. An updated set of kinematic and visual measurements is taken as described above. This data set is saved to a text file and taken to matlab, where a Nelder-Mead optimization algorithm is used to find a set of DHPs that

explain best this set of measurements. This process takes from 25-120 minutes, depending mostly on the computational hardware. The results from this search are a full set of DHPs that can be used in Robonaut's control software, as well as in the visual cortex, to accurately map between the manipulator workspace and the visual workspace.

Calibration Fixture

While hand-eye calibration could be performed using visual measurements of any point on the kinematic chain (or many points on the chain), we designed the calibration fixture shown in Figure 5 for several reasons. A visual measurement point distant from the wrist axes gives good observability for motions in the wrist roll and yaw axes. This particular fixture does not give good observability of wrist pitch. A sphere is a relatively simple target, visually. The center of a sphere is observable and well defined regardless of the relative pose between the cameras and fixture. The fixture also has a hand-guard to ensure a relatively repeatable grasp. This prototype fixture should be replaced with a more robust fixture that exhibits a very repeatable grasp and significant distance from the wrist axes in each of the wrist DOFs.

APPROACH/THEORY

This section describes the theoretical underpinnings of the above procedures and presents the algorithms used in the calibration. First, the Sphere Hough Transform and its use in locating the calibration fixture in the eye coordinate system are described. Then, the setup for the nonlinear optimization at the heart of the hand-eye calibration system is described.

Finding a Sphere in a depth dataset

Central to this task is the accurate localization of the calibration fixture in the visual coordinate system. We utilize the existing depth-from-disparity stereo algorithms developed by Eric Huber at the Johnson Space Center, and perform a search for a sphere-shaped object in the depth map (a 2D array of depths measured from the visual coordinate system origin).

The BallFinder algorithm begins with a seed location. This location is currently set to the kinematically derived prediction of the calibration fixture location, expressed in the visual coordinate system.

Points outside a large spherical region centered at this location are rejected from consideration. This pruning step rejects distant points, such as the floor, from further consideration as possible members of the sphere surface. Next, a minimal surface area test is performed on all surviving points. Based on the distance of the seed location from the camera, the expected number of points on the sphere's surface is computed.

Locations under consideration that are not members of a contiguous set of some fraction of this size are rejected from consideration. This pruning step eliminates small isolated regions. All remaining points participate in a voting scheme based on the generalized Hough Transform described below.

The Hough Transform is a classic computer vision algorithm in which lines are located in an image by allowing each point that is a member of a line to vote for some set of M lines that could have created this point. Lines which truly exist in an image will accrue more votes, and the top vote getters are very good candidates for lines in an image. See [15] for more detail on the standard Hough Transform. This algorithm can be extended to describe many types of parameterized shapes, such as circles [8] or spheres. In the Sphere-Hough Transform, each point P_{ss} that survives the pruning algorithms described above votes for a set of M spheres (centered at $P_{sc,i}$ $i=1..M$, points randomly sampled from the surface of a sphere centered at P_{ss}) of which this point could be a surface point. Each of these points P_{sc} represents one vote, in the Hough Transform paradigm, for a sphere centered at point P_{sc}. In our case, the voting is in Cartesian space, since the radius of our calibration fixture is known. The location with the most votes is deemed the most likely to contain the actual sphere center. Figure 6 depicts a slice of the voting results from an example image (the depth slice that contains the winning vote) on the right, and on the left the input image with the winning 3-D location projected into it using the current camera calibration.

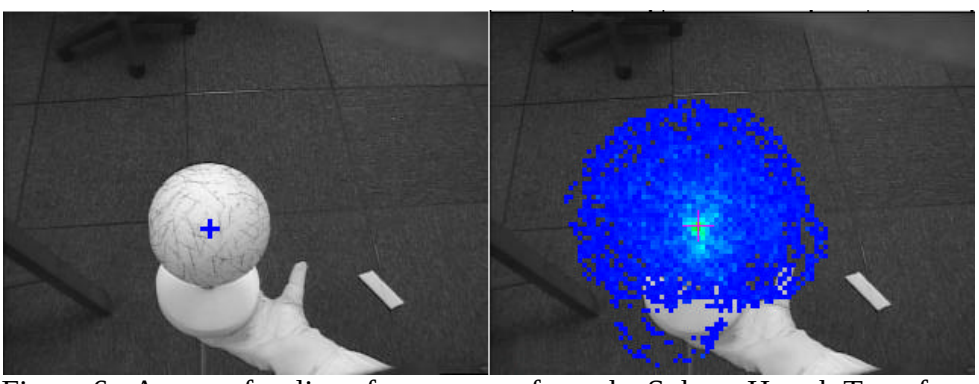

Figure 6: A map of a slice of vote-space from the Sphere-Hough Transform

Hand-Eye Calibration

The constrained and full variants of the calibration procedure described above differ only in which parameters are optimized. In this section, we will describe the setup for this optimization procedure. As described above, SumVisCal or GenVisCalData generates a set of joint measurements and visual measurements of the calibration fixture. For each configuration i in the calibration set, the kinematic model is used to predict the location of the calibration fixture in the chest coordinate system. This is a function of the Denevit-Hartenberg Parameters (DHPs) as well as the joint angles of the arm:

$$\mathbf{P}_{c,i} = A_1(DH,\mathbf{q_i})*A_2(DH,\mathbf{q_i})*\ldots*A_7(DH,\mathbf{q_i})*\mathbf{P}_e,$$

where $\mathbf{P}_e$ is the (fixed) position of the calibration fixture in the hand coordinate system, DH contains the DHPs for the arm and neck, and $\mathbf{q_i}$ contains the joint angles for the arm and neck[2]. The kinematic model for the neck is used to predict the transformation from the chest coordinate system to the eye coordinate system in the same way:

$$T_{ce,i} = A_{1N}(DH, \mathbf{q_i}) * A_{2N}(DH, \mathbf{q_i})$$

These transformations are used to create a kinematic estimate of the position of the calibration fixture in the eye coordinate system:

$$\mathbf{P}_{e,i} = (T_{ce,i})^{-1} * \mathbf{P}_{c,i}$$

We also have for each configuration *i* in the calibration set the visual measurement of the 3-D position of the calibration fixture, also in the eye coordinate system, that we call $\mathbf{P}_{v,i}$

The optimization procedure attempts to minimize the difference between $\mathbf{P}_{v,i}$ (fixed) and $\mathbf{P}_{e,i}$ (function of DHPs) over all *i* in the calibration set by search in DHP space. Our objective function (the function to minimize) for this search is the sum of the distances between point pairs in our calibration set. A Nelder-Mean simplex method [17] is used to minimize this function by search in the DHP space. For the constrained, or "daily" calibration, the joint angle offsets (*?$_i$ =0..9*) are optimized. For the full calibration all nonzero (and non-?/2) DHPs are optimized. Several pairs of offsets, designed to be symmetric, are constrained to be equal and only contribute one dimension to the DHP search space.

EXPERIMENTAL RESULTS

Four experiments were performed to validate this procedure. In Experiment 1, the mean error between a set of visual observations and kinematically derived predictions is reduced using the updated DHPs. In Experiment 2, the updated DHPs derived from the data in Experiment 1 are used in conjunction with a "daily calibration procedure" to reduce the mean error between a set of visual data and kinematically derived predictions over the as-designed DHPs. In Experiment 3, the "daily calibration" procedure is performed on one-half of a dataset, and the errors in both this set and the half of the dataset not used for training are evaluated. Finally, in Experiment 4, an improvement in Robonaut's performance on an autonomous wrench-grasping experiment is observed.

Experiment 1 – Effect of Full Calibration

The as-designed DHPs for the arm and neck are presented in Table 1. As noted in the table, there is some disagreement in the DRL about the as-designed values for these parameters. A set of 67 robot configurations were chosen, and the reported joint angles and visual measurements logged. *The mean distance between the kinematically derived prediction for these measurements and the actual visual measurements was 13.75cm.*

[2] For convenience, we take $\mathbf{q_i} = [q_{1,arm}\ q_{2,arm} \ldots q_{7,arm}\ q_{1,neck}\ q_{2,neck}]^T$, and similarly concatenate the DHPs.

Table 1: As-Designed D-H Parameters for Robonaut, Unit A.
Angles are in degrees and lengths in cm.

	Shoulder Roll	**Shoulder Pitch**	**Elbow Roll**	**Elbow Pitch**	**Wrist Roll**	**Wrist Pitch**	**Wrist Yaw**	**Neck Yaw**	**Neck Pitch**	**Neck Roll**
? $_j$	0	0	0	0	?	-??	0	0	???[1]	0
$\mathbf{d_j}$	30.48[2]	0	36.83	0	36.83	0	-1.27	28.575[3]	0	2.92
? $_j$	-??	??	-90	??	-90	??	0	??	??	0
$\mathbf{a_j}$	-6.35	6.35	-5.08	5.08	0	0	3.81	-5.08	-11.96	0

1 – a slight head-tilt is more comfortable for teleoperation
2 – in some designs, this is 32.94 cm
3 – in some designs, this is 27.31 cm

The full calibration procedure described above was performed on this data set. The DHPs shown in Table 2 were found. For the same set of 67 configurations, the *mean distance between the kinematically derived prediction (using the updated DHPs) for these measurements and the visual measurements was 1.85cm.* These data are summarized in Figure 4.

Table 2: D-H Parameters for Robonaut, Unit A, after Visual-Kinematic Calibration
Angles are in degrees and lengths in cm.

	Shoulder Roll	**Shoulder Pitch**	**Elbow Roll**	**Elbow Pitch**	**Wrist Roll**	**Wrist Pitch**	**Wrist Yaw**	**Neck Yaw**	**Neck Pitch**	**Neck Roll**
? $_j$	-8.444	-1.535	-2.013	-3.780	?????	-95.58	4.227	-1.057	?????	-90.0
$\mathbf{d_j}$	31.856	0	35.498	0	35.498	0	-0.053	28.292	0	2.537
? $_j$	-??	??	-90	??	-90	??	0	??	??	0
$\mathbf{a_j}$	-5.056	5.056	-0.99	.99	0	0	11.358	-6.773	-12.667	0

Experiment 2 – Effect of Daily Calibration

The DHPs shown in Table 2 were used to predict the location of the calibration fixture in a set of 150 unique configurations, *with an average error of 7.94cm.* This was several days (and several power cycles) after Experiment 1, so we expect that the reported joint angles differed from the actual joint angles by different amounts than estimated in Experiment 1. The "Daily Calibration Procedure" described above was used to compute the updated joint angle offsets shown in Table 3. The remainder of the DHPs were as shown in Table 2. *The average error was reduced to 2.02cm over this dataset.*

Table 3: D-H Parameters for Robonaut, Unit A, after "Daily Calibration"
Angles are in degrees.

	Shoulder Roll	**Shoulder Pitch**	**Elbow Roll**	**Elbow Pitch**	**Wrist Roll**	**Wrist Pitch**	**Wrist Yaw**	**Neck Yaw**	**Neck Pitch**	**Neck Roll**
? $_j$	-1.70	0.243	-0.149	2.323	?????	-90.634	9.340	0.601	?????	-93.078

Experiment 3 – Effect of Daily Calibration on Non-Training Data

In this experiment, half a data set was used to tune the DHPs, and the power of these parameters to predict the position of the calibration fixture in the other half of the data set was tested. The DHPs shown in Table 2 were used to predict the location of the calibration fixture in a set of 75 unique configurations, the first half of the configurations from Experiment 2, *with an average error of 7.87cm.* This was several days (and several power cycles) after Experiment 1, so we expect that the reported joint angles differed from the actual joint angles by different amounts than estimated in Experiment 1. The "Daily Calibration Procedure" described above was used to compute the updated joint angle offsets shown in Table 4. The remainder of the DHPs were as shown in Table 2. After this calibration, *the average error was reduced to 2.04cm over this dataset.* This set of DHPs was then used with no further optimization to predict the position of the calibration fixture in 75 configurations that had not been used in training, the second half of the configurations from Experiment 2. Over this dataset, the DHPs from Tables 2 and 4 produced an *average prediction error of 2.35cm.*

Table 4: D-H Parameters for Robonaut, Unit A, after "Daily Calibration"
Angles are in degrees.

	Shoulder Roll	**Shoulder Pitch**	**Elbow Roll**	**Elbow Pitch**	**Wrist Roll**	**Wrist Pitch**	**Wrist Yaw**	**Neck Yaw**	**Neck Pitch**	**Neck Roll**
$?_j$	0.186	0.153	0.404	3.198	?????	-88.780	8.946	-0.121	?????	-93.620

Experiment 4 – Performance in Wrench-Grasping Experiment

As an example of the types of tasks that the DRL is demanding of Robonaut, this section presents the contribution of this visual calibration to an experiment run by a team from Vanderbilt University on autonomous wrench-grasping. In this experiment, a teleoperator is observed grasping wrenches in nine different workspace locations. Figure 7 shows the physical setup for this experiment. Robonaut's vision system is used to observe the wrench in a unique location, and a learning algorithm [16] is used to grasp the wrench in this location.

In this experiment, described in more detail in [18], a 6DOF Cartesian-space vision-workspace correction was implemented. This correction was a linear interpolation between vision/workspace mismatches recorded at several locations using teleoperator data. This workspace correction reduced vision/kinematic mismatches, but not enough to enable Robonaut to grasp the wrench. The updated DHPs shown in Table 2 were experimentally placed into the inverse kinematics procedures for Robonaut, and the workspace correction was removed. The system was immediately able to grasp wrenches at several different positions in the workspace.

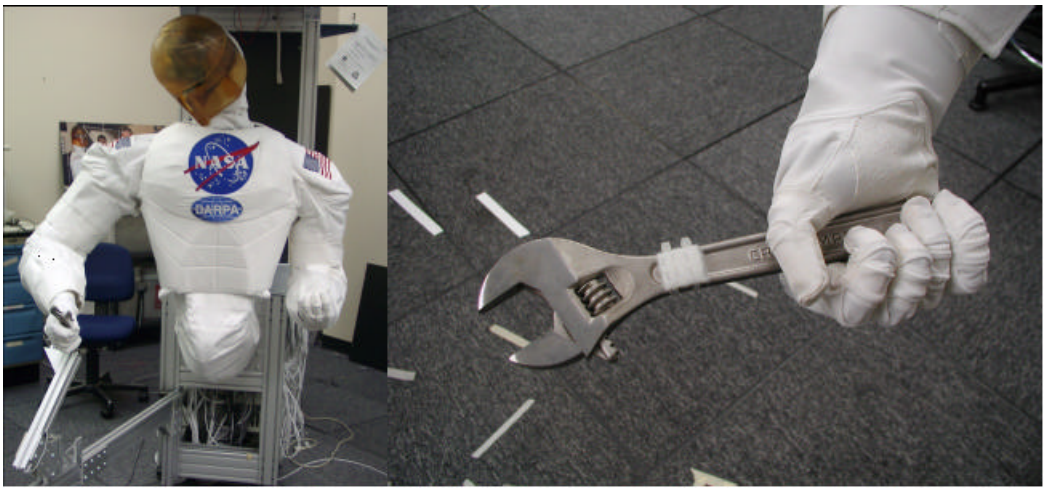

Figure 7: An overview of the Vanderbilt wrench-grasping experiment

DISCUSSION AND CONCLUSIONS

Closed-loop self-calibration of the combined kinematic/visual systems for Robonaut Unit A has been performed. This calibration procedure does not explicitly register the visual or the kinematic system with ground-truth, but modifies the perceptions associated with the kinematic movements to match the perceptions of the vision system. Procedures and algorithms have been developed that will enable the robot to be recalibrated when necessary. These procedures have reduced vision-kinematic mismatch from 13-15cm to 2-3cm in various situations, and have enabled the DRL team to continue increasing the autonomous capability of Robonaut.

Future work

There are several directions in which this work could be improved. The most obvious is to do a careful extrinsic calibration of Robonaut's vision system so that this closed-loop procedure will more accurately reflect distances and rotations in the workspace. Also useful would be to systematically study the number of measurements required to calibrate the system, both in the reduced and full cases. The system should also be extended to calibrate the left arm of Unit A and each arm of Unit B. The method described in this report should extend to these situations in a very straightforward manner. With some extension, this method could be extended to the simultaneous calibration of the vision system and both arms of Robonaut.

REFERENCES

[1] R.O. Ambrose, H. Aldridge, R.S. Askew, R.R. Burridge, W. Bluethmann, M. Diftler, C. Lovchik, D. Magruder, and F. Rehnmark, "Robonaut: NASA's space humanoid," Humanoid Robotics, Vol. 15(4), July/August 2000, pp. 57-62.

[2] H. Zou and L. Notash, "Discussions on the Camera-Aided Calibration of Parallel Manipulators," in Proceedings, 2001 CCTOMM Symposium on Mechanism, Machines, and Mechatronics,

[3] M. Meggiolaro, GT. Scriffignano, and S. Dubowsky, "Manipulator Calibration using a single endpoint contact constraint," in Proceedings 2000 ASME Design Engineering Technical Conferences, Baltimore MD, Sept 2000

[4] T. Huntsberger, Y. Cheng, E. Baumgartner, M. Robinson, P. Schenker, "Sensory Fusion for Planetary Surface Robotic Navigation, Rendezvous, and Manipulation Operations," in Proceedings 11th International Conference on Advanced Robotics, Coimbra, Portugal, June 30-July 3, 2003, pp 1417-1424.

[5] J. M. Hollerbach and C. W. Wampler, " The Calibration Index and the Role of Input Noise in Robot Calibration", in Robotics Research; 7th Int. Symp., G. Giralt and G. Hirzinger, eds., Springer, 1996, pp. 558-568.

[6] D. J. Bennett, M. M. Hollerbach and D. Geiger, "Autonomous Robot Calibration for Hand-eye Coordination", Int. J. of Robotics Research, vol. 10, no. 5, pp. 550-559, October 1991;

[7] D. J. Bennett and J. M. Hollerbach, "Self-calibration of Single-loop, closed kinematic chains formed by dual or redundant manipulators," Proceedings of the 27th Conference on Decision and Control, Austin TX December 1988, pp 627-629.

[8] M. Wohlfart, "Hough transform applications in Computer Graphics (with focus on medical visualization)," unpublished, Spring 2003

[9] R. I. Hartley, "Self-calibration of Stationary Cameras", Int. J. of Computer Vision, vol. 22, no. 1, 1997, pp, 5-23.

[10] E. Hayman, K. Knight and D. W. Murray, "Self-Alignment of an Active Head from Observations of Rotation Matrices", Proc. 15 Int. Conf. on Pattern Recog., Barcelona, Sept. 2000

[11] Z. Zhang, "A Flexible New Technique for Camera Calibration," IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(11), 1330-1334, 2000.

[12] A. J. Koivo, "Fundamentals for control of robotic manipulators, " New York, N.Y. : Wiley, 1989.

[13] M. W. Spong and M. Vidyasagar, "Robot dynamics and control," New York : Wiley, 1989.

[14] J. Craig, "Introduction to robotics : mechanics and control," 2nd ed, Reading, Mass. : Addison-Wesley, 1989.

[15] M. Sonka, V. Hlavac and R. Boyle, "Image Processing , Analysis, and Machine Vision," 2nd ed, Pacific Grove, CA: Brooks/Cole, 1999

[16] Cambron, M.E., and Peters II, R.A., "Learning sensory motor coordination for grasping by a humanoid robot," Proceedings of the 2000 IEEE International Conference on Systems, Man and Cybernetics, Nashville, TN, v.5, pp. 870-875, October 2000.

[17] W. Press, S. Teukolsky, W. Vertterling, and P. Flannery, "Numerical Recipies in C: The art of scientific computing," 2nd edition, New York, NY : Cambridge University Press, 1992

[18] M. Diftler, R. Platt, C. Culbert, R. Ambrose, W. Bluethmann, "Evolution of the NASA/DARPA Robonaut Control System," in Proceedings 2003 IEEE International Conference on Robotics and Automation, Sept 15-19, 2003, Tiapei, Taiwan.

[19] Z. Zhang, O. Faugeras and R. Deriche, "An Effective Technique for Calibrating a Binocular Stereo Through Projective Reconstruction Using both a Calibration Object and the Environment", Videre, vol. 1, No. 1, pp. 57-68, 1997.