Proceedings, 2007 IEEE Aerospace Conference,
March 3-10, 2007, Big Sky MT, USA

Hand-Eye Calibration Using Active Vision

Kevin Nickels
Department of Engineering Science
Trinity University
One Trinity Place
San Antonio TX 78259-7200
210-999-7543
knickels@trinity.edu

Eric Huber
NASA/Johnson Space Center
2100 NASA Road One
eric@roboteyes.com
Houston TX 77058

Matthew DiCicco
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109, USA
Matthew.A.Dicicco@jpl.nasa.gov

Abstract— The project described in this paper designed and implemented a hand-eye calibration method for manipulators under observation by stereo cameras. This method has been utilized on Johnson Space Center's Robonaut, and on a planetary manipulator mock-up at the Jet Propulsion Laboratory. The intent of this calibration is to improve the manipulator's hand-eye coordination.

The approach uses kinematic and stereo vision measurements, namely the joint angles self-reported by the arm and 3-D positions of a calibration fixture as measured by vision, to estimate the transformation from the arm's base coordinate system to its hand coordinate system and to its vision coordinate system. In this formulation, the stereo measurements are assumed to be accurate, and any mismatches are absorbed in a modified model of the arm.

These methods have shown to reduce mismatch between kinematically derived positions and visually derived positions on Robonaut Unit A from a mean of 13.75cm to a mean of 1.85cm. Improved performance in semi-autonomous tasks is also described. On JPL's manipulator, with kinematics similar to that of the Mars Exploration Rover, the calibration reduced the mismatch from 15.26mm to between 3mm and 5.5mm.

1-4244-0525-4/07/$20.00/
IEEEAC paper # 1050

TABLE OF CONTENTS

1. INTRODUCTION

THE Dexterous Robotics Laboratory (DRL) at NASA Johnson Space Center (JSC) has developed a ground-based prototype humanoid robot called Robonaut[1], shown in Figure 1. Robonaut has been designed so that it could, for example, assist astronauts during EVA type tasks[1]. Its initial control has been by tele-operation, but the DRL is beginning to implement several semi-autonomous and fully autonomous controllers for Robonaut, necessitating improved hand-eye coordination for the system.

With the increasing demand for a higher level of science return in future surface missions, lander and rover-mounted robotic arms must exhibit a higher level of performance over current capabilities. As a recent example, the MER mission requirements set for the IDD includes a precision placement requirement of 10 mm in position and 10 degrees in orientation with respect to a science target when the IDD is deployed from a stationary rover base [2].

[1] There are two versions of Robonaut, referred to as Unit A and Unit B. In this paper, we will use the name Robonaut for Unit A.

Figure 1. Ground-based Robonaut system

This paper documents a method for automatic hand-eye calibration, originally developed for Robonaut, that improves the correspondance of the hand and eye coordinate systems.

Prior Work

Much previous work has been done on the self-calibration of redundant manipulators using internal or external kinematics constraints. [3], [4], [5], [6]. Of particular note is the treatment of Bennett and Hollerbach of a vision or metrology system as an additional kinematic link [6], [7], allowing one 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 [8].

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 [5] or LEDs to localize points: we opted for a spherical calibration fixture and visual measurements of this fixture. In order to accurately locate the spherical fixture in the image, a generalized Hough transform was used. The generalized Hough transform is described in [9].

While this system reduces vision-kinematic mismatch by adjusting the kinematic model, there are other options for reduction. A system called Hybrid Image Plane/Stereo (HIPS), developed at JPL, as well as its predecessor, Camera Space Manipulation, adjust the camera models to reduce this same mismatch [10], [11]. The End-effector Pose Error Compensation method [12] adds a workspace offset to the visually sensed position to bring the two coordinate systems into closer agreement.

All of these systems are closed-loop systems. They reduce the residuals between visually and kinematically derived predictions, but do not necessarily adjust these parameters to match the workspace. If the visual system is not well calibrated, via for example, the procedures described in [13], [14], [15], [16], the adjustments performed by this process will not cause the predictions to correlate well with workspace positions. Even with good calibrations, it is important to keep in mind that the adjustments made to the original models optimize the *agreement* of the kinematic and visual systems, each of which have been individually calibrated beforehand. The final visual and kinematic systems thus explicitly do not represent the closest approximation to "ground truth" that is possible, but rather various adjustments made for optimal agreement between models within a fixed workspace.

Task Background

Robonaut has historically been operated by a human teleoperator. The DRL is increasing the autonomy level of the tasks performed by Robonaut [17], [18]. 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.

Robonaut Unit A is constructed with relative joint encoders. As the arm is powered down in the evening and restarted the next morning, the position of all joints on the 7DOF right arm and the 2DOF neck[2] 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.

Kinematic Model

Robonaut's arm is a redundant manipulator with 7 degrees of freedom. This manipulator can be described by 7 homogeneous 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_j axis followed by a screw about the x'_j axis (the rotated x_j axis) [19], [20] and one involving a screw about the x_j axis followed by a screw about the z'_j axis [21]. In addition to the screw order, these systems differ in their conventions for placing coordinate axes relative to links.

[2]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.

DHPs for equivalent manipulator systems will therefore differ in the two structures. In this work, we utilize the former, where each transformation is defined as

$$A_j = Trans(x'_j, a_j)Rot(x'_j, \alpha_j)Trans(z_j, d_j)Rot(z_j, \theta_j),$$

where Rot implies a rotation about an axis and Trans implies a translation along an axis[21]. The position and orientation of the last link can be computed by a sequence of DH transformations defining the kinematic model

$$T_c = A_1 A_2 A_3 \dots A_{nf}$$

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 [3] from the chest coordinate system to the eye coordinate system are parameterized in this way.

2. Calibration Procedures

Operationally, two methods are used to gather 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.

A set of DHP values is derived from these data in one of two ways. In a daily calibration, estimates are generated only for the joint angle offsets. In a full calibration, estimates are generated for all relevant DHPs.

Daily Calibration

The daily calibration method is intended to be a lightweight process that is performed frequently. First, a prerecorded set of configurations is loaded into the system. The robot is sent to each configuration and a new set of kinematic and visual measurements (see Figure 2is taken.

As the individual measurements are updated, summary of the entire data set (figure 3) is also updated, including the average residual between the visual measurements and the kinematically derived predictions. This process takes approximately 10 minutes for a 65-element calibration set. Most of this time is spent moving the robot, not performing calculations.

An optimization algorithm is then used to estimate a set of joint angle offsets θ_i, $i = 0 \dots 10$ 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. A 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.

[3] 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.

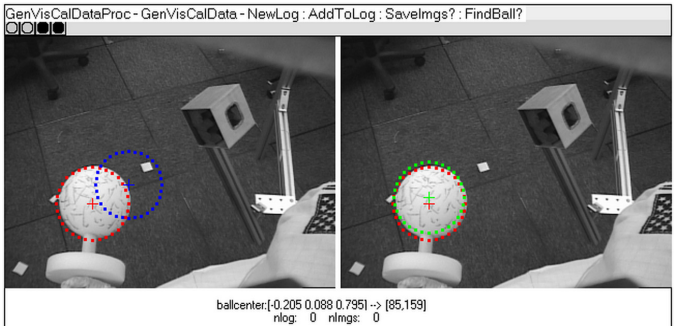

Figure 2. Comparing Kinematic and Visual Predictions of the location of a sphere. Blue markings in the left frame indicate the projection into the image of the sphere based on the original kinematic model. Green markings in the right frame indicate the projection into the image of the sphere based on the updated kinematic model. Red markings in each frame indicate the visually detected location of the image of the sphere. These colors are utilized in all displays in this paper.

Full Calibration

A full calibration method has also been developed. 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 an optimization algorithm is used to find a set of DHPs that best explain 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 to accurately map between the manipulator workspace and the visual workspace.

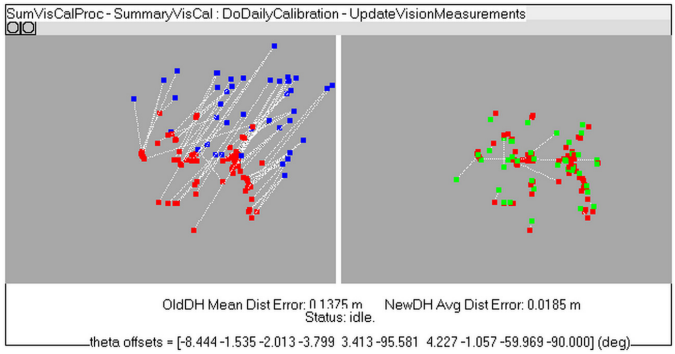

Figure 3. Summary of a calibration data set.

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 2 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. 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.

3. THEORY

This section describes the theoretical underpinnings of the above methods 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, shown in Figure 2, in the visual coordinate system. We utilize the existing depth-from-disparity stereo algorithms developed by the DRL 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 vote 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 [9] for more detail on the standard Hough Transform. This algorithm can be extended to describe many types of parameterized shapes, such as circles [22] or 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 $P_{sc,i} \quad i = 1 \ldots M$, 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 4 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.

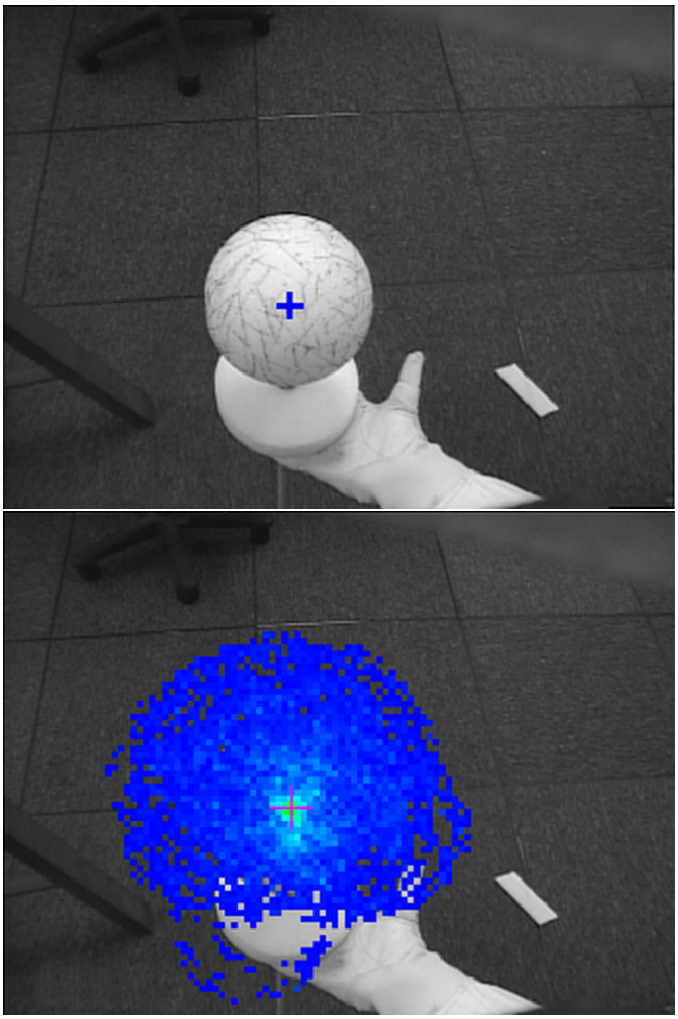

Figure 4. Sphere-Hough Transform Results. The winning location is projected into the top image, the 3D slice containing the winning location projected into the bottom image.

Optimization

The daily and full calibrations described above differ only in which parameters are optimized. In this section, we will describe how this calibration is posed as an optimization problem. As described above, a set of joint measurements and visual measurements of the calibration fixture is generated. 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 Denavit-Hartenberg Parameters (DHPs) as well as the joint angles of the arm:

$$P_{c,i} = A_1(DH, \mathbf{q}_i)A_2(DH, \mathbf{q}_i)\dots A_7(DH, \mathbf{q}_i)P_e,$$

where 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[4]. 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:

$$T_{ce,i} = A_{1N}(DH, \mathbf{q}_i)A_{2N}(DH, \mathbf{q}_i)A_{3N}(DH, \mathbf{q}_i).$$

These transformations are used to create a kinematic estimate of the position of the calibration fixture in the eye coordinate system: $P_{e,i} = (T_{ce,i})^{-1}P_{c,i}$. 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 $P_{v,i}$.

The optimization attempts to minimize the difference between $P_{v,i}$ (fixed) and $P_{e,i}$ (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. We currently use a Nelder-Mead simplex method [23] to minimize this function by search in the DHP space. For daily calibration, the joint angle offsets ($\theta_i = 0 \dots 10$) are optimized. For a full calibration all nonzero (and non-$\pi/2$) 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.

4. RESULTS - ROBONAUT

Four experiments were performed on Robonaut to validate our calibration procedures. These are summarized in Table 1.

In the first experiment, the mean residual between a set of visual observations and kinematically derived predictions is compared with the existing and revised DHPs. In the second experiment, the updated DHPs derived from the data above are used in conjunction with a daily calibration on a different day, and the residuals are again compared. In the third experiment, a daily calibration is performed on one-half of a dataset, and the residuals in both this set and the half of the dataset not used for training are evaluated. Finally, the effect of using an updated calibration in an autonomous wrench-grasping experiment is described.

Experiment 1 — Effect of Full Calibration

The as-designed DHPs for the arm and neck are presented in Table 2. A set of 67 robot configurations were chosen, and the reported joint angles and visual measurements logged. This data set will be referred to as DS1. The mean 3D distance between the kinematically derived prediction for these measurements and the actual visual measurements was 13.75cm.

A full calibration was performed on DS1. The DHPs shown in Table 3 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 3.

Experiment 2 — Effect of Daily Calibration

The DHPs shown in Table 3 were used to predict the location of the calibration fixture in a set of 150 unique configurations (referred to as DS2), with an average residual of 7.94cm. This was several days (and several power cycles) after the experiment described in Experiment 1, so it is expected that the reported joint angles deviated from the actual joint angles by different amounts than estimated in Table 3. A daily calibration was used to compute the updated joint angle offsets shown in Table 4. The remainder of the DHPs were as shown in Table 3. With the new offsets, the average residual was reduced to 2.02cm over this dataset.

Experiment 3 — Effect on Novel Data

In this experiment, the 150 element DS2 data set was randomly split into two 75 element subsets, DS2A and DS2B. The data in DS2A was used to tune the DHPs, and the power of these parameters to predict the position of the calibration fixture in the DS2B was tested.

The DHPs shown in Table 3 were used to predict the location of the calibration fixture in each of the 75 unique configurations in DS2A, with an average residual of 7.87cm. This was several days (and several power cycles) after Experiment 1, so the reported joint angles likely deviated from the actual joint angles by different amounts than estimated in Table 3. A daily calibration on DS2A was used to compute the updated joint angle offsets shown in Table 4. The remainder of the DHPs were as shown in Table 3. After this calibration, the average residual was reduced to 2.04cm over the DS2A dataset. This set of DHPs was then used with no further optimization to predict the position of the calibration fixture in the 75 configurations that had not been used in training (DS2B). Over the DS2B dataset, the DHPs from Tables 3 and 5 produced an average prediction error of 2.35cm.

[4]For convenience, we take $\mathbf{q}_i = [q_{1,arm} \dots q_{7,arm} q_{1,neck} q_{2,neck} q_{3,neck}]^T$, and similarly concatenate the DHPs.

Table 1. Quantitative Summary of Hand-Eye Experiments on Robonaut.

Experiment	DH Set	Training Sample	Test Sample	Mean Error (mm)
Exp. 1	As-Designed (Tbl. 2)	–	DS1	13.75 cm
Exp. 1	Full (Tbl. 3)	DS1	DS1	1.85 cm
Exp. 2	Full (Tbl. 3)	DS1	DS2	7.94 cm
Exp. 2	Daily (Tbl. 4)	DS2	DS2	2.02 cm
Exp. 3	Full (Tbl. 3)	DS1	DS2A	7.87 cm
Exp. 3	Daily (Tbl. 5)	DS2A	DS2A	2.04 cm
Exp. 3	Daily (Tbl. 5)	DS2A	DS2B	2.35 cm

Table 2. 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
θ_j	0	0	0	0	0	-90	0	0	-60[1]	-90
d_j	30.48[2]	0	36.83	0	36.83	0	-1.27	28.575[3]	0	2.92
α_j	-90	90	-90	90	-90	90	0	90	90	0
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

Table 3. Robonaut D-H Parameters after full calibration using DS1. 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
θ_j	-8.444	-1.535	-2.013	-3.780	3.414	-95.58	4.227	-1.057	-59.969	-90.0
d_j	31.856	0	35.498	0	35.498	0	-0.053	28.292	0	2.537
α_j	-90	90	-90	90	-90	90	0	90	90	0
a_j	-5.056	5.056	-0.99	.99	0	0	11.358	-6.773	-12.667	0

Table 4. Robonaut D-H Parameters after daily calibration using DS2. Angles are in degrees.

	Shoulder Roll	Shoulder Pitch	Elbow Roll	Elbow Pitch	Wrist Roll	Wrist Pitch	Wrist Yaw	Neck Yaw	Neck Pitch	Neck Roll
θ_j	-1.70	0.243	-0.149	2.323	3.810	-90.634	9.340	0.601	-61.249	-93.078

Wrench-Grasping Experiment

As an example of the types of tasks that the DRL is demanding of Robonaut, this section presents the contribution of 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 5 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 [17] is used to attempt to grasp the wrench in this location.

In this experiment, described in more detail in [18], a 6DOF Cartesian-space vision-workspace correction was initially implemented. This workspace correction was computed by measuring the position of the wrench as computed by the arm kinematics and as measured by the vision system at several locations using teleoperator data. The correction used was then a linear combination of the recorded corrections. This workspace correction reduced vision/kinematic mismatches at the novel targets, but not enough to enable Robonaut to grasp the wrench.

Later, the updated DHPs shown in Table 3 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.

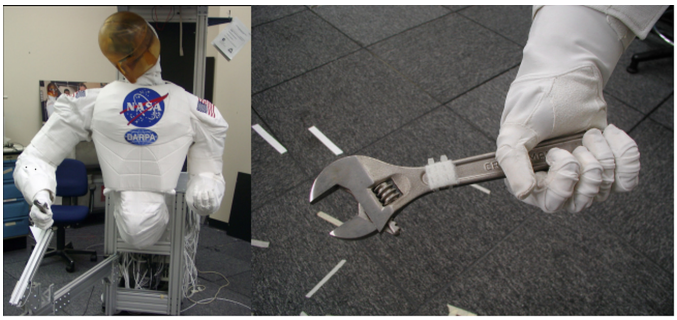

Figure 5. Vanderbilt wrench-grasping experiment

Table 5. Robonaut D-H Parameters after daily calibration on DS2A. Angles are in degrees.

	Shoulder Roll	Shoulder Pitch	Elbow Roll	Elbow Pitch	Wrist Roll	Wrist Pitch	Wrist Yaw	Neck Yaw	Neck Pitch	Neck Roll
θ_j	0.186	0.153	0.404	3.198	2.716	-88.780	8.946	-0.121	-62.990	-93.620

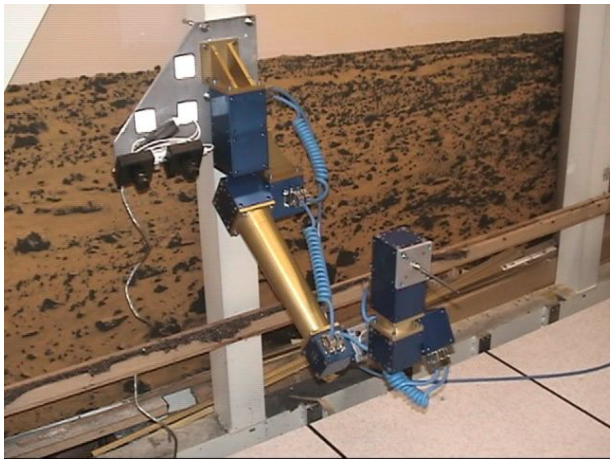

Figure 6. The Five Degree of Freedom Modular manipulator

5. Results - Modular Manipulator

The calibration procedure described in this paper has also been implemented at NASA's Jet Propulsion Laboratory, in a comparison study of vision-guided manipulation algorithms described in detail in [24], on one of the manipulators in JPL's Modular Robotic Testbed, where robotic arms of various kinematic configurations can be quickly and easily constructed from AMTEC PowerCubes and steel piping. This system, shown in Figure 6, has the kinematic structure and camera locations similar to the Mars Exploration Rovers [25], [26].

Baseline manipulation with this manipulator is similar to that of Robonaut, as described above, being slightly more accurate due in large part to the shorter manipulator and fixed cameras. The nominal camera and arm models yield approximately 1.53 cm of positioning error, or difference between the position of the fiducial predicted by the arm model and that reported by the camera system, over a test set of 50 points.

Using the Full Calibration procedure described above in Section 2, the DH parameters for the modular manipulator were adjusted based on training sets of 50, 25, 20, 15, 10, 5, and 1 poses. These poses were randomly sampled from a set of training data. After the Full Calibration, 50 points not used for training were used to test the calibration. The average difference between the position of the fiducial predicted by the (modified) arm model and that reported by the vision system was reduced severalfold to between 3 and 5.5 mm, as shown in Table 6.

Table 6. Mean Residual between Kinematics and Vision for 50 novel test points

Training Sample Size	Mean Error (mm)
(Baseline)	15.26
1	5.26
5	3.22
10	3.37
15	3.27
20	3.44
25	3.62
50	3.38

6. Discussion and Conclusions

Closed-loop self-calibration of the combined kinematic and visual systems for Robonaut Unit A has been performed. This calibration 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. In particular, the DH parameters derived are not claimed to be the ones that most accurately reflect the "true" structure of the manipulator, but rather are a set of parameters[5] that cause the closed-loop vision/manipulator system perform optimally. 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.

In separate work, the same system was tested on a mock-up of a planetary manipulator at the Jet Propulsion Laboratory, reducing vision-kinematic mismatch from 15mm to 3-5.5mm.

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 paper 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.

[5] With a redundant manipulator, there may be many sets of parameters, none of which are the "true" value, that perform equally well.

ACKNOWLEDGMENT

The authors would like to thank NASA, DARPA, and the Center for Intelligent Systems at Vanderbilt University for their contributions to this work.

Some of this work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The work involves important contributions from many colleagues at both JPL and collaborating institutions. We gratefully acknowledge these interactions and note many of the specific developments in references that follow.

REFERENCES

[1] R. Ambrose, H. Aldridge, R. Askew, R. Burridge, W. Bluethmann, M. Diftler, C. Lovchik, D. Magruder, and F. Rehnmark, "Robonaut: NASA's space humanoid," *Humanoid Robotics*, vol. 15, no. 4, pp. 57–62, July 2000.

[2] E. T. Baumgartner *et al.*, "The Mars Exploration Rover instrument positioning system," in *Proc. IEEE Aerospace Conference*, Big Sky, MT, Mar. 2005.

[3] H. Zou and L. Notash, "Discussions on the camera-aided calibration of parallel manipulators," in *Proc CCTOMM Symposium on Mechanism, Machines, and Mechatronics*, 2001.

[4] M. Meggiolaro, G. Scriffignano, and S. Dubowsky, "Manipulator calibration using a single endpoint contact constraint," in *Proc. ASME Design Engineering Technical Conference*, Sept. 2000.

[5] T. Huntsberger, Y. Cheng, E. Baumgartner, M. Robinson, and P. Schenker, "Sensory fusion for planetary surface robotic navigation, rendezvous, and manipulation operations," in *Proc. IEEE International Conference on Advanced Robotics*, 2003, pp. 1417–1424.

[6] 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 e. G.Hirzinger, Eds. Springer, 1996, pp. 558–568.

[7] D. J. Bennett, M. M. Hollerbach, and D. Geiger, "Autonomous robot calibration for hand-eye coordination," *Intl. J. Robot. Res.*, vol. 10, no. 5, pp. 550–559, Oct. 1991.

[8] D. J. Bennett and J. M. Hollerbach, "Self-calibration of single-loop, closed kinematic chains formed by dual or redundant manipulators," in *Proceedings of the Conference on Decision and Control*, Dec. 1988, pp. 627–629.

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

[10] M. Robinson, E. Baumgartner, T. Litwin, and K. Nickels, "Hybrid image plane/stereo (HIPS) manipulation for robotic space applications," , Under Review.

[11] S. B. Skaar, W. H. Brockman, and W. S. Jang, "Three-dimensional camera space manipulation," *Intl. J. Robot. Res.*, vol. 9, no. 4, pp. 22–39, 1990.

[12] M. Bajracharya, M. DiCicco, and P. Backes, "Vision-based end-effector position error compensation," in *Proc. IEEE Aerospace Conference*, Big Sky, MT, Mar. 2006.

[13] R. I. Hartley, "Self-calibration of stationary cameras," *Intl. J. Comp. Vis.*, vol. 22, no. 1, pp. 5–23, 1997.

[14] E. Hayman, K. Knight, and D. W. Murray, "Self-alignment of an active head from observations of rotation matrices," in *Proc. 15 Int. Conf. on Pattern Recog.*, Sept. 2000.

[15] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, no. 11, pp. 1330–1334, 2000.

[16] 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.

[17] M. E. Cambron and R. A. Peters II, "Learning sensory motor coordination for grasping by a humanoid robot," in *Proceedings of the 2000 IEEE International Conference on Systems, Man and Cybernetics*, vol. 5, Oct. 2000, pp. 870–875.

[18] M. Diftler, R. Platt, C. Culbert, R. Ambrose, and W. Bluethmann, "Evolution of the NASA/DARPA robonaut control system," in *Proc. 2003 IEEE Intl. Conf. Robot. Automat.*, Sept. 2003.

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

[20] M. W. Spong and M. Vidyasagar, *Robot dynamics and control*. New York, N.Y.: Wiley, 1991.

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

[22] M. Wohlfart, "Hough transform applications in computer graphics (with focus on medical visualization)," unpublished.

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

[24] K. Nickels, M. DiCicco, M. Bajracharya, and P. Backes, "Vision guided manipulation for planetary robotics - position control," 2007, in Preparation.

[25] S. Squyres *et al.*, "Athena Mars rover science investigation," *Journal of Geophysical Research*, vol. 108, no. E12, p. 8062, Nov. 2003.

[26] J. N. Maki *et al.*, "Mars Exploration Rover engineering

cameras," *Journal of Geophysical Research*, vol. 108, no. E12, p. 8071, Nov. 2003.

Kevin Nickels *is an associate professor in the Department of Engineering Science at Trinity University. He received the B.S. degree in Computer and Electrical Engineering from Purdue University (1993), and received the M.S. degree (1996) and the Ph. D. (1998) in Electrical Engineering from The University of Illinois at Urbana- Champaign. He is currently working in the areas of computer vision, pattern recognition, and robotics. Dr. Nickels spent summer 2003 as a Faculty Fellow in Aeronautics and Space Research at NASA's Johnson Space Center, concentrating on hand-eye calibration, and the 2005-2006 academic year as a Senior Engineer in Avionics Systems and Space Technology at NASAs Jet Propulsion Laboratory, concentrating on Vision Guided Manipulation.*

Eric Huber *works for the Mitre Corporation at NASA's Johnson Space Center. He has been involved in real-time vision for over 15 years, and is currently the vision lead software engineer for Robonaut.*

Matthew DiCicco *is an associate member of technical staff at the NASA Jet Propulsion Laboratory. He received a masters degree from MIT where he studied high power robotic manipulators for the Navy. At JPL he has continued his work on robotic manipulators and is currently performing research tasks for the Mars Science Lab rover mission as a member of the Mobility and Manipulation Group.*