In Proceedings of the ASME Dynamic Systems and Control Division,
1995 International Mechanical Engineering Congress and Exposition,
San Francisco, CA, Nov., 1995, DSC-Vol. 57-1

SHAPE FROM MOTION APPROACH TO RAPID AND PRECISE FORCE/TORQUE SENSOR CALIBRATION

Richard M. Voyles, Jr.
J. Daniel Morrow
Robotics Ph.D. Program
Carnegie Mellon University
Pittsburgh, Pennsylvania

Pradeep K. Khosla
Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, Pennsylvania

ABSTRACT

We present a new technique for multi-axis force/torque sensor calibration called *shape from motion*. This technique retains the noise rejection of a highly redundant data set but eliminates the need for explicit knowledge of the redundant applied load vectors, yielding faster, more accurate calibration results. A constant-magnitude force (a mass in a gravity field) is randomly moved through the sensing space while raw data is continuously gathered. Using only the raw sensor signals, the motion of the force vector (the "motion") and the calibration matrix (the "shape") are simultaneously extracted by singular value decomposition. We have applied this technique to several types of force/torque sensors and present experimental results for a 2-DOF fingertip and a 6-DOF wrist sensor with comparisons to the standard least squares approach.

1. INTRODUCTION

Calibration of multi-axis force/torque sensors is a time-consuming ordeal that traditionally requires the precise application of a set of known forces and torques carefully selected to adequately span the space of the sensor. Reduction by least squares, the standard technique of force sensor calibration, requires both the application of a spanning set of precisely known loads and the measurement of associated sensor (e.g. strain gage) values. Although effective, this approach is cumbersome because of the relatively large number of accurate loads that must be applied to the sensor to reduce noise; often 12 to 20 *or more* are required to produce a reasonably accurate calibration matrix for a 6-degree-of-freedom (DOF) sensor, even though six are theoretically sufficient. Despite this, few direct efforts have been made to find alternative calibration methods. A few researchers have attempted to formalize the calibration process as an offshoot of design methods. Uchiyama, et al (1991) have done the most comprehensive treatment along this vein, but Bicchi and Dario (1988) and Nakamura et al (1988) provided early foundations.

Every calibration method benefits from more data. The problem with the least squares technique is that each new piece of data has a high cost. Each piece of data has two parts: the raw output of the sensor and the carefully applied load that produced that output. While the sensor output is easily collected, carefully applying the load is much more time consuming.

We propose a new approach to force sensor calibration based on "shape and motion decomposition" techniques from computer vision (Tomasi and Kanade, 1991). Unlike the least squares technique, *shape from motion* calibration *does not require knowledge of all applied loads.* Instead, calibration is performed with a very large number of easily obtained raw outputs and only a few known loads. This allows much quicker calibration with fewer precise measurements but preserves the robustness and noise-rejection of a large, redundant data set

2. THE CALIBRATION PROBLEM

A sensor converts an applied load, $\tilde{\mathbf{m}}$, into a measurement vector, $\mathbf{z}$ (Figure 1). In the case of a force sensor, for example, the applied load is a vector of forces and torques and the measurement is a vector of strain gage readings. The purpose of the calibration function is to invert this transformation so that, given a measurement vector, we can estimate the load which generated it.

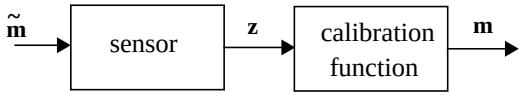

FIGURE 1: SENSOR AND CALIBRATION FUNCTIONS

If the sensor is linear, the calibration function is a constant matrix (and possibly a bias vector) that transforms **z** into **m**.

$$\mathbf{Cz} = \mathbf{m} \; or \; \mathbf{z}^T\mathbf{C}^T = \mathbf{m}^T \tag{1}$$

Our calibration problem is to recover **C**, the calibration matrix in the presence of two types of noise: errors in the applied load vector, $\tilde{\mathbf{m}}$, and measurement noise in the measurement vector, **z**. Because of these sources of random noise, an accurate calibration *cannot* be achieved without redundant data; more redundant data results in a better signal to noise ratio for calibration recovery.

Least Squares Calibration Solution

The standard technique for solving calibration problems is the least squares (LS) method. This requires the application of several *known* loads, $\mathbf{m}_i$, and the measurement of the corresponding sensor vectors, $\mathbf{z}_i$. These data form two matrices that plug into (1):

$$\begin{bmatrix} \mathbf{z}_1^T \\ \vdots \\ \mathbf{z}_n^T \end{bmatrix} \mathbf{C}^T = \begin{bmatrix} \mathbf{m}_1^T \\ \vdots \\ \mathbf{m}_n^T \end{bmatrix} \tag{2}$$

from which the calibration matrix can be computed using the pseudoinverse of the measurement matrix, **Z**:

$$\mathbf{C}^T = \mathbf{Z}^+\mathbf{M} \; . \tag{3}$$

The accurate application of all load vectors, $\mathbf{m}_i$, makes the least squares calibration process tedious. The difference between the true applied load and the intended applied load (applied load error) must be minimized for an accurate calibration because this error manifests itself directly in the calibration matrix, **C**. This error can only be minimized by exercising extreme care when applying each and every load used in the LS calibration. Thus, incorporating a large number of redundant data points in the LS calibration procedure is very expensive (in time) due to this applied load error minimization requirement. In practice, this limits the number of redundant loads. In our lab, for a 6-DOF sensor, users typically lose patience after only 6 to 14 redundant loads for a total of 12 to 20 known applied loads (although 30 - 40 are more desirable).

Shape from Motion Calibration

With the shape from motion calibration approach *we do not need to know all the applied loads*, $\mathbf{m}_i$, but only a constraint which relates them. This feature allows us to apply many redundant load vectors and use the resulting raw measurements to determine calibration without knowing exactly what the loads were that caused them. *Literally hundreds of data points can be used in the shape from motion calibration procedure in a fraction of the time that a dozen data points in the LS procedure require.* A small number of known applied loads are required to establish the desired reference frame, *but none of the redundant data requires accurate applied load knowledge for the shape from motion calibration procedure.* The ability to economically collect and apply massive amounts of redundant data in shape from motion calibration accounts for its impressive advantage over least squares.

Why do we call it *shape from motion*? The calibration matrix encodes the mechanical structure of the sensor, including the placement of sensing elements and the properties of the material from which it is made. These are what defines the sensor's intrinsic *shape*. The motion refers to the movement of the applied load around the sensor. Shape from motion refers to the fact that we can recover the shape of the sensor by knowing the theoretical rank of the shape and applying arbitrary motion to the load.

Shape from Motion Derivation

We demonstrate the shape from motion approach beginning with a representation of the sensor function which maps a load onto a measurement:

$$\mathbf{z}_i^T = \mathbf{m}_i^T\mathbf{S}, \tag{4}$$

where $\mathbf{z}_i^T$ is a *1* x *p* measurement vector, $\mathbf{m}_i^T$ is a *1* x *m* load vector, and **S** is the *m* x *p* shape matrix. There are *p* sense elements and *m* DOF. Note, from (1), that the calibration matrix, **C**, is easily computed from the shape matrix, **S**, as

$$\mathbf{C} = [\mathbf{S}^T]^+ . \tag{5}$$

If we apply *n* loads and collect the measurements, we can express (4) as the matrix equation (similar to (2))

$$\mathbf{Z} = \mathbf{MS} , \tag{6}$$

where **Z** is the *n* x *p* matrix of measurements and **M** is the *n* x *m* matrix of applied loads. Note that the shape matrix, **S**, is unchanged. **M** is our motion matrix which encodes the applied loads to the sensor.

In traditional calibration techniques (i.e. least squares), both **Z** and **M** are known. **Z** contains the output signals of the sensor while **M** is constructed from careful external measurements of the applied forces that correspond to each vector in **Z**. These external measurements generally involve scales, bubble levels, and protractors and can be very time consuming. Our technique eliminates the need to know **M** *a priori* by simultaneously determining **M** and **S** given only **Z**. We achieve this by performing a singular value decomposition (SVD) (Klema and Laub, 1980) on **Z.**

SVD produces the following unique decomposition of any *n* x *p* matrix, **Z**:

$$\mathbf{Z} = \mathbf{U}\Sigma\mathbf{V}^T \tag{7}$$

where **U** is an *n* x *n* orthogonal matrix, Σ is an *n* x *p* "diagonal" matrix[1] of the singular values of **Z** in descending order, and **V** is a *p* x *p* orthogonal matrix.

Assuming we know that the "proper rank"[2] of **Z** is *r*, it can be shown (Strang, 1988) that the best projection of **Z** onto an *r*-dimensional space (for $r \le p$) is

$$\mathbf{Z}^* = \mathbf{U}^*\Sigma^*\mathbf{V}^{*T} , \tag{8}$$

where $\mathbf{U}^*$ consists of the first *r* columns of **U**, Σ^* is a diagonal matrix of the first *r* singular values, and $\mathbf{V}^{*T}$ consists of the first *r* rows of $\mathbf{V}^T$. ($\mathbf{Z}^*$ is not the same as $\mathbf{Z}^+$ in equation (3).)

1. padded with zeroes as needed
2. the rank in the absence of noise

Equation (8) gives us the best possible (Forsythe et al, 1977) rank-*r* representation of **Z**, so, combining (6) and (8) yields

$$\begin{aligned} \mathbf{MS} &= \mathbf{U}^*\Sigma^*\mathbf{V}^{*\mathrm{T}} = \hat{\mathbf{M}}\hat{\mathbf{S}} \\ \hat{\mathbf{M}} &= \mathbf{U}^*(\Sigma^*)^{1/2} \\ \hat{\mathbf{S}} &= (\Sigma^*)^{1/2}\mathbf{V}^{*\mathrm{T}} \end{aligned} \tag{9}$$

Unfortunately, $\hat{\mathbf{M}}$ and $\hat{\mathbf{S}}$ are not yet the true *motion* and *shape*. They are indeterminate by an affine transformation. Given *any* invertible *r* x *r* matrix, **A**, (an affine transform)

$$\hat{\mathbf{M}}\hat{\mathbf{S}} = (\hat{\mathbf{M}}\mathbf{A}^{-1})(\mathbf{A}\hat{\mathbf{S}}), \tag{10}$$

so we must find an appropriate matrix, **A**, such that

$$\begin{aligned} \mathbf{M} &= \hat{\mathbf{M}}\mathbf{A}^{-1} \\ \mathbf{S} &= \mathbf{A}\hat{\mathbf{S}} \end{aligned} \tag{11}$$

We do this by applying a geometric constraint to the individual vectors of the motion matrix to solve for $\mathbf{A}^{-1}$. Once **A** is known, we can solve for **S** using (11) and **C** using (5). Finally, we introduce a few precise measurements ($\mathbf{z}_i$, $\mathbf{m}_i$ both known) in order to orient the calibration matrix with respect to the desired reference frame and to scale the result to the desired engineering units. Figure 2 illustrates this procedure in flow chart form. (We present examples in Section 3 for clarification.) The gist of this technique is that we have replaced the least squares requirement of knowing all individual loads with the shape from motion requirement of a geometric constraint on **M** with no assumptions of smoothness of motion.

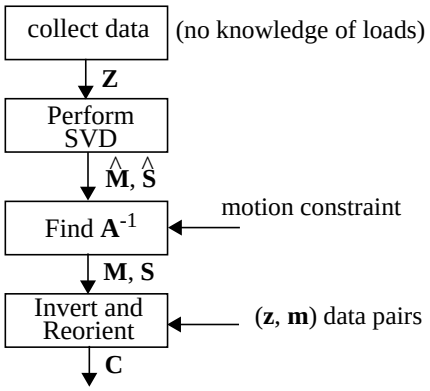

FIGURE 2: SHAPE FROM MOTION CALIBRATION PROCEDURE

3. APPLICATIONS OF SHAPE FROM MOTION

Force-Only Sensors

2-DOF Calibration. As an example of the shape from motion calibration procedure, we first consider a 2-DOF subset (Figure 3) of a fingertip sensor for a dextrous hand (Voyles et al, 1989). Force sensing results from the deformation of a single cantilever beam. There are four strain gages that respond to forces in the plane, thus,

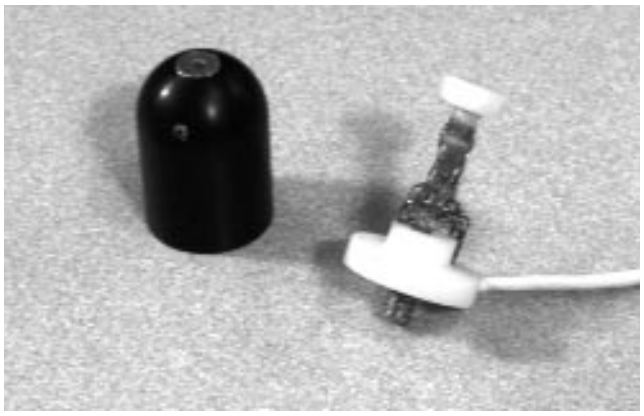

FIGURE 3: THE FINGERTIP SENSOR

the measurement vector for this experiment has four elements and the load vector has two elements (*x* and *y* directions). To collect measurements, a *constant-magnitude* force is applied in random directions in the plane using the setup of Figure 4.

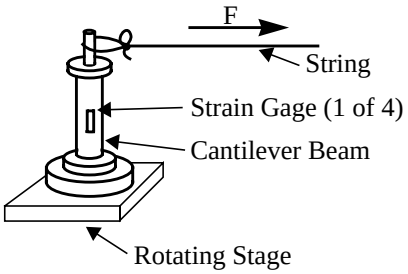

FIGURE 4: TEST JIG FOR THE 2-D FINGERTIP SENSOR

The sensor equation can be written in the form of (4):

$$\begin{bmatrix} z_{i1} & z_{i2} & z_{i3} & z_{i4} \end{bmatrix} = \begin{bmatrix} \cos\theta_i & \sin\theta_i \end{bmatrix} \begin{bmatrix} s_{11} & s_{12} & s_{13} & s_{14} \\ s_{21} & s_{22} & s_{23} & s_{24} \end{bmatrix} . \tag{12}$$

The constant magnitude of the force has been factored out, leaving only cosθ and sinθ in the motion matrix. This allows us to use $\cos^2\theta + \sin^2\theta = 1$ as our constraint. Furthermore, the rank of the shape matrix is, at most, 2, so the proper rank of our *n* x *p* measurement matrix, **Z**, is also 2. If we denote the elements of $\mathbf{A}^{-1}$ by a_{11}, a_{12}, a_{21}, and a_{22}, and the *i*th row of $\hat{\mathbf{M}}$ by m_{i1} and m_{i2}, then the constraint equation becomes

$$\begin{aligned} 1 = {} & m_{i1}^2\left(a_{11}^2 + a_{12}^2\right) + 2m_{i1}m_{i2}(a_{11}a_{21} + a_{12}a_{22}) \\ & + m_{i2}^2\left(a_{21}^2 + a_{22}^2\right) \end{aligned} \tag{13}$$

We solve for $(a_{11}^2 + a_{12}^2)$, $(a_{11}a_{21} + a_{12}a_{22})$, and $(a_{21}^2 + a_{22}^2)$ in (13) in the least squares sense and then numerically solve for the individual a_{ij} values. (In fact, an analytic solution exists for the a_{ij}'s in the 2-DOF case.) Having computed $\mathbf{A}^{-1}$, one can solve for the shape, **S**, using (11) and the calibration matrix, **C**, using (5).

Unfortunately, the resulting calibration matrix is not oriented in any particular direction. To align it with our desired reference frame, we introduce *one* precise load (a **z**, **m** pair, both vectors known) to rotate and scale it appropriately.

We use Mathematica to perform the SVD and to solve the non-linear equations to extract **A**, **S**, and finally, **C**. Figure 5 shows a polar plot of the recovered motion of one calibration trial. Despite significant noise, the motion displays excellent average circularity. (Quantitative assessments of precision appear in Section 4.)

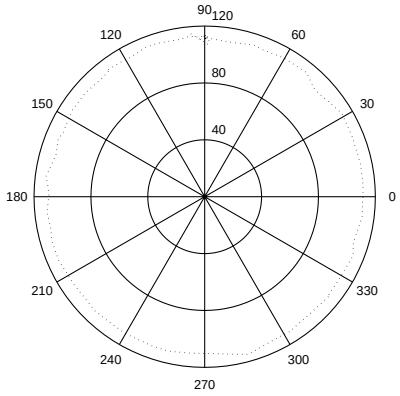

FIGURE 5: POLAR PLOT OF RECOVERED MOTION (MAGNITUDE AND DIRECTION) WITH 111.4G MASS

3-DOF Force Sensor. An extension of the 2-D fingertip force sensor to 3-D was performed using a standard 6-DOF Lord force/torque sensor with a compact mass mounted on the face plate. In reality there is a small moment arm associated with this arrangement, but we will ignore that for now.

The resulting motion matrix consists of vectors of the form [cosθsinψ sinθsinψ cosψ] so **M** is now rank 3. Likewise, **S** acquires another row, becoming 3 x *p*. The derivation of motion and shape is the same as in the 2-DOF case with the necessary modifications to the metric constraint (constant magnitude force application). This relation is easily derived so we will not repeat it here. (It is similar in form to (13).) Plots of the recovered motion of the Lord force/torque sensor during an actual calibration trial appear in Figure 6. Although it is difficult to judge by eye, it is an accurate sphere. (Again, quantitative results appear in Section 4.)

Force/Torque Sensors

The previous section showed the application of the shape from motion calibration method to 2-DOF and 3-DOF force-only sensors. This section will develop the shape from motion technique for force/torque sensors. We will see that the introduction of torque measurement introduces a new problem: the moment arm of the applied force is embedded in the shape matrix.

2-D Force/Toque Sensor: 2 Forces, 1 Torque. Consider a 2-D force/torque sensor which responds to forces *and* torques in the plane (3-DOF). Our motion vector has 3 elements: two forces and one torque, so the true shape, $\tilde{\mathbf{S}}$, must have 3 rows. However,

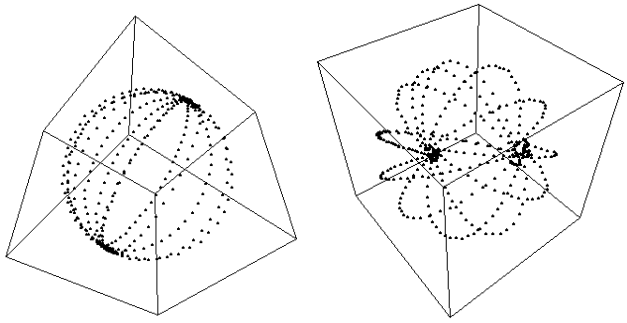

FIGURE 6: RECOVERED MOTION OF LORD F/T SENSOR

if we generate the load by applying a force to moment arm, **r**, then the torque value is linearly related to the force values.

$$\begin{bmatrix} f_x & f_y & \tau \end{bmatrix} = \begin{bmatrix} f_x & f_y \end{bmatrix} \begin{bmatrix} 1 & 0 & -r_y \\ 0 & 1 & r_x \end{bmatrix} \tag{14}$$

Because the torque is a linear combination of the first two columns of the motion matrix, it does not increase the rank. Therefore, given a constant moment arm, **r**, the motion matrix has rank 2 (not rank 3).

$$\tilde{\mathbf{M}}\tilde{\mathbf{S}} = \mathbf{M} \begin{bmatrix} 1 & 0 & -r_y \\ 0 & 1 & r_x \end{bmatrix} \tilde{\mathbf{S}} = \mathbf{M}\mathbf{S} \tag{15}$$

where $\tilde{\mathbf{M}}$ consists of vectors of the form $[f_x f_y \tau]$ and **M** consists of vectors of the form $[f_x f_y]$. The moment arm becomes embedded in the recovered shape matrix, **S** (rank 2):

$$\mathbf{S} = \begin{bmatrix} \mathbf{I} & \mathbf{X}^{\mathrm{T}} \end{bmatrix} \tilde{\mathbf{S}} \tag{16}$$

where **I** is the *2* x *2* identity matrix and **X** is the planar equivalent of the cross product matrix, $[-r_y\ r_x]$. (This vector produces the magnitude of the cross product of **r** and any other vector in the *x-y* plane.) We call the matrix $[\mathbf{I}\ \mathbf{X}^{\mathrm{T}}]$ the *rank squashing matrix* (or *squashing matrix*) because the constant moment arm "squashes" down the rank of **S**.

We are interested in recovering the true shape matrix, $\tilde{\mathbf{S}}$, which is the same for any load with any moment arm. We cannot simply pseudo-invert the squashing matrix because it has rank of only 2 and so will not completely specify the true shape matrix, which is rank 3. To overcome this, we need to perform the calibration using two different moment arms and combine the results. But there is a problem: the two recovered shape matrices are *arbitrarily oriented.* So we must introduce a precise data point to orient each squashed shape matrix to a consistent frame of reference before we can combine them to extract the true shape matrix.

Once the squashed shapes are "co-oriented," we find $\tilde{\mathbf{S}}$ from (16) by conglomerating the shape and squashing matrices from both trials such that

$$\begin{bmatrix}\mathbf{S}_1\\ \mathbf{S}_2\end{bmatrix} = \begin{bmatrix}[\mathbf{I}\ \mathbf{X}_1]\\ [\mathbf{I}\ \mathbf{X}_2]\end{bmatrix}\tilde{\mathbf{S}} \tag{17}$$

or

$$\tilde{\mathbf{S}} = \begin{bmatrix}[\mathbf{I}\ \mathbf{X}_1]\\ [\mathbf{I}\ \mathbf{X}_2]\end{bmatrix}^{+}\begin{bmatrix}\mathbf{S}_1\\ \mathbf{S}_2\end{bmatrix}. \tag{18}$$

In summary, for the 3-DOF planar force/torque sensor, we run the calibration procedure twice with the same constant force but different moment arms and then combine the results. This provides three independent columns of information so $\tilde{\mathbf{S}}$ (which is rank 3) can be recovered.

3-D Force/Torque Sensor: 3 forces and 3 torques. The 6-DOF force/torque sensor shape from motion calibration procedure follows a similar development as the 3-DOF force/torque procedure. Now our load vector has six elements: three forces and three torques. Again, if we apply a constant magnitude force with a particular moment arm, **r**, the torques are linear combinations of the force so the rank of the motion matrix is 3 (not 6).

$$\begin{bmatrix}f_x & f_y & f_z & \tau_x & \tau_y & \tau_z\end{bmatrix} = \begin{bmatrix}f_x & f_y & f_z\end{bmatrix}\begin{bmatrix}1 & 0 & 0 & 0 & r_z & -r_y\\ 0 & 1 & 0 & -r_z & 0 & r_x\\ 0 & 0 & 1 & r_y & -r_x & 0\end{bmatrix} \tag{19}$$

Again, the squashing matrix consists of the identity matrix and the transpose of the cross product matrix, but now these submatrices are represented in 3-space (19).

To reconstruct the full 6-DOF shape matrix (which is rank 6), *three* trials are required with three different moment arms. Three trials are required because it is impossible to choose two squashing matrices that yield rank greater than 5 when combined. We solve for $\tilde{\mathbf{S}}$ in a manner similar to (18) using co-oriented shape matrices:

$$\tilde{\mathbf{S}} = \begin{bmatrix}[\mathbf{I}\ \mathbf{X}_1]\\ [\mathbf{I}\ \mathbf{X}_2]\\ [\mathbf{I}\ \mathbf{X}_3]\end{bmatrix}^{+}\begin{bmatrix}\mathbf{S}_1\\ \mathbf{S}_2\\ \mathbf{S}_3\end{bmatrix}, \tag{20}$$

where $[\mathbf{I}\ \mathbf{X}_i]$ is as it appears in (19). Note that a total of 6 precise (**z**, **m**) points are required - two per moment arm - in order to consistently orient the $\mathbf{S}_i$'s.

4. EXPERIMENTAL COMPARISON TO LEAST SQUARES

The point of any sensor calibration is to provide an accurate characterization of the sensor. But desired accuracy is almost always traded-off with the time required to perform the calibration procedure. Usually, more accurate calibration measurements will produce a more accurate calibration, but they will take more time. Our motion and shape approach is exceptional in that it allows the collection of much more data in *less time* as compared to the least squares approach. This produces a *more accurate and precise* calibration, for a given accuracy in calibration measurement.

The reason it takes less time is obvious -- fewer accurately applied loads are required. For a given accuracy of each measurement, fewer of them results in less time, despite the fact we collect many more data points.

The reasons for greater accuracy are threefold. First of all, more data can be gathered in less time. Second, the applied loads are unknown so *there is no applied load error.* Finally, the least squares approach takes all the data available and attempts to minimize the total error. It does this at the expense of the shape because, from equation (3), **C** is the only free variable. As such, both the measurement errors *and* the applied load errors distort the shape. On the other hand, the shape from motion approach has two free variables: the shape and the motion. Errors are distributed between the two because shape from motion relies on a known physical constraint of the motion and minimizes error with respect to that constraint to maintain the shape. It is important to realize that this is a true constraint based on geometry, not a constraint based on some sensor model that may be wrong. Greater accuracy results from this hard physical constraint.

To demonstrate these assertions, we performed both shape from motion and least squares calibrations on our 2-DOF fingertip sensor and a 6-DOF Lord force/torque sensor. In both cases, the least squares method was aided by averaging each force measurement over time to minimize noise.

Fingertip Results

For the shape from motion approach, we chose a mass in the linear region of the sensor, near the upper extreme of what we consider "reasonable."[1] We picked the upper extreme to improve signal-to-noise, but we stayed within the "reasonable range" to reduce any effects of minor nonlinearities. The chosen test mass weighed in at 111.4 grams and the calibration procedure took only 1.5 minutes while collecting a total of 96 unknown data points and 1 known load.

For the least squares approach, we chose five different masses, each of which was applied at 40-degree increments around the circle for a total of 45 accurate loads. For each load we gathered several measurements and averaged them. The same criteria were used in selecting the calibration masses which were: 60.4, 78.3, 95.2, 111.4, and 122.8 grams. The procedure consisted of adjusting the angle of the rotating stage to within 0.2 degrees, sequentially applying the five masses, and then incrementing the angle. The entire procedure took 17 minutes to collect all 45 known data points.

Because of the nearly exhaustive nature of the least squares calibration, the time difference is more than 10:1. But even cutting the procedure down to 15 measurements results in nearly a 4:1 time advantage for shape from motion, not to mention the loss in accuracy for least squares due to lesser noise rejection. The large number of

1. Reasonable in relation to the tasks we intend to tackle.

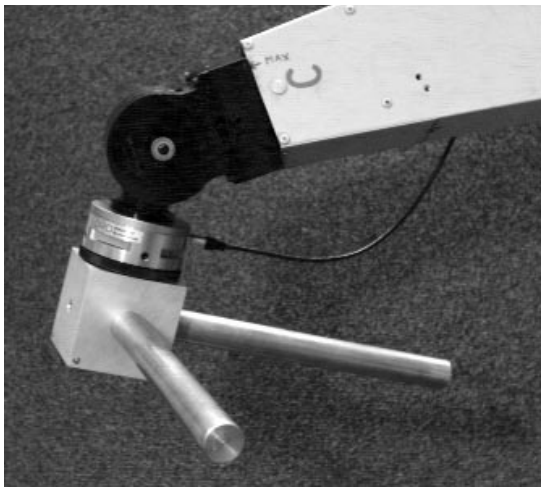

FIGURE 7: CALIBRATION FIXTURE ON LORD SENSOR

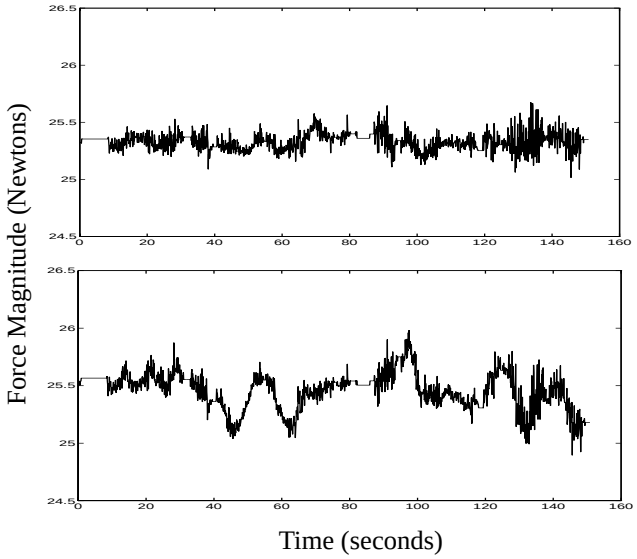

FIGURE 8: COMPARISON OF SHAPE FROM MOTION (TOP) AND LEAST SQUARES (BOTTOM) AS ORIENTATION CHANGES

least squares measurements gathered should produce nearly the best possible least squares calibration.

To assess the calibration results quantitatively, we hung each of the five masses used for the least squares calibration on the string and gathered raw data as the sensor was rotated 360 degrees. We then calculated the force magnitude at each sample point and performed two calculations. The first determined the average magnitude to assess absolute accuracy. The second smoothed the magnitude signal to eliminate noise and determined the standard deviation of the smoothed signal. This assesses the precision of the measurement across all orientations. The best and worst of these data are tabulated in Table 1.

Table 1: FINGERTIP SENSOR COMPARISON

	Least Squares			Shape from Motion		
load (g)	mean (g)	error (%)	std dev	mean (g)	error (%)	std dev
60.4	60.1	0.5	1.20	60.5	0.2	1.13
111.4	110.2	1.1	0.83	111.6	0.2	0.60

Lord F/T Results

For the 6-DOF sensor we built a special calibration fixture (Figure 7) to help us make quick, accurate measurements. It consists of a 3-inch aluminum cube (approximately 1 kg) and two 1-inch diameter brass bars (approximately 1 kg each). The cube has one threaded hole in the center of each face into which the bars can screw. This provides 16 unique force/torque combinations that can be quickly and precisely selected during calibration. The simple design is easy to machine yet it dramatically reduces the time required for the least squares approach because the flat faces provide an accurate reference surface for leveling the device in various orientations. It is quite convenient for the shape from motion approach as well.

Despite using the calibration fixture and collecting only 13 load vectors (force magnitudes of 12.09N, 22.9N, and 33.71N with different moments), the least squares approach required 69 minutes of data collection time compared to only 34 for shape from motion[1] (force magnitude of 33.71 N). Although the time difference is not as dramatic as the 2-DOF case, it is significant and gets more so as the number of least squares force vectors increases.

To compare precision, we performed the same experiments as in the previous subsection but randomly moved the sensor in 3-D rather than just the plane. The extremes of these data are tabulated in Table 2. Again, shape from motion accuracy (mean) is as good or better while precision (standard deviation) is always better.

Table 2: LORD F/T SENSOR COMPARISON

	Least Squares			Shape from Motion		
load (N)	mean (N)	error (%)	std dev	mean (N)	error (%)	std dev
9.26	9.09	1.8	.084	9.09	1.8	.058
25.24	25.53	1.7	.157	25.41	0.7	.062

Figure 7 shows the magnitude of the linear forces of the one-bar experiment from Table 2 (25.24N including internal mass). We chose to illustrate this data set not because it's the best (it's not), but because the mass was *different* from that used during the shape from motion approach but the *same* as in four of the least squares vectors. This avoids giving the shape from motion approach an unfair advantage. The two plots result from applying the calibration matrices from both techniques to the *same batch of data*. It is clear that the shape from motion plot has less variation, which indicates

1. both include the time to compute applied force vectors

greater precision. This is confirmed by the smoothed standard deviation measurements in Table 2.

To verify the absolute accuracy of both techniques, we assembled the calibration fixture in two configurations that were not used in either calibration procedure and gathered data in ten different orientations. The accuracy of each applied load in orientation and force is 1 milliradian and 0.02 Newtons, respectively. Table 3 shows average errors in force magnitude and direction across all ten trials. Figure 9 shows the angular error of both techniques for all ten trials.

Table 3: AVERAGE LORD FORCE ERRORS

method	magnitude error (N)	angle error (rad)
Shape from Motion	0.080	0.0025
Least Squares	0.144	0.0046

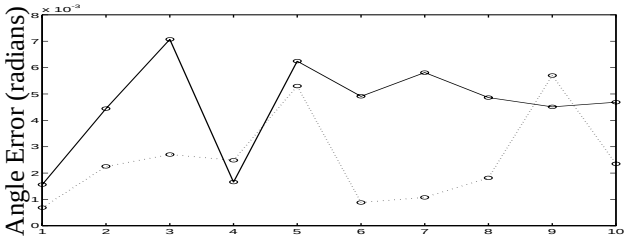

FIGURE 9: ANGULAR ERROR OF LEAST SQUARES (SOLID) AND SHAPE FROM MOTION (DOTTED) AT SPECIFIC KNOWN LOADS.

5. CONCLUSIONS

We have described an improved calibration technique for linear force/torque sensors based on *shape from motion* decomposition that provides a more accurate calibration matrix with less effort from the user. The intrinsic "shape" of the sensor is extracted by randomly moving the calibration force through a spanning range of motion. Both the motion and the pseudoinverse of the calibration matrix are simultaneously recovered from the raw sensor values.

The shape from motion approach yields a calibration matrix that is more precise and more accurate than the corresponding least squares technique that takes several times as long because more data can be collected with fewer error sources. It has been successfully applied to force-only and force/torque sensors from 2 to 6 degrees of freedom with a single batch file written for Mathematica.

Although we determined bias vectors as part of the experimental procedure, it is possible to augment the technique to automatically determine these as well, provided there are at least *n+1* sensing elements for an *n*-DOF sensor.

Of course, this technique has limitations. It linearizes the response around a single force magnitude whereas the least squares technique finds the best linear representation across a range of force magnitudes. In practice, we have found our force sensors have sufficient linearity that shape from motion still prevails. However, this raises the important issue of design of experiments. We briefly discussed our rationale for selection of applied loads in Section 4, but we leave the important topic of experiment design to the classical textbooks in the field (e.g. Diamond, 1981).

Another potential experimental flaw is induced by gathering the raw measurements while the sensor is in motion. This permits the quickest data collection with very dense coverage of the sensing space, but is subject to dynamic effects. Care must be taken to move the sensor with minimal acceleration so measurements are quasistatic. Of course, static "move-stop-measure" data collection works, too, and remains faster than least squares.

In closing, we note the elimination of the final step, orienting the shape with respect to an external reference frame, still provides useful information about the sensor. Because known loads are not required to derive this, it is possible to use shape from motion for *primordial learning* in autonomous agents (Voyles et al, 1995).

6. ACKNOWLEDGEMENTS

Mr. Voyles was supported in part by a National Defense Science and Engineering Graduate Fellowship. Mr. Morrow was supported in part by a DOE Computational Science Graduate Fellowship.

7. REFERENCES

Bicchi, A. and P. Dario, 1988, "Intrinsic Tactile Sensing for Artificial Hands," *Robotics Research: The 4th International Symposium*, R.C. Bolles and B. Roth, editors, MIT Press, Cambridge, MA, pp. 83-90.

Diamond, W., 1981, *Practical Experiment Designs for Engineers and Scientists*, Lifetime Learning Pub., Belmont, CA.

Forsythe, G.E., M. Malcolm, and C.B. Moler, 1977, *Computer Methods for Mathematical Computations*, Prentice Hall, Englewood Cliffs, NJ.

Klema, V.C. and A.J. Laub, 1980, "The Singular Value Decomposition: Its Computation and Some Applications," *IEEE Trans. on Automatic Control*, v.25, n.2, pp. 164-176.

Nakamura, Y., T. Yoshikawa, and I. Futamata, 1988, "Design and Signal Processing of Six-Axis Force Sensors," *Robotics Research: The 4th International Symposium*, R.C. Bolles and B. Roth, editors, MIT Press, Cambridge, MA, pp 75-81.

Strang, G., 1988, *Linear Algebra and Its Applications*, third edition, Harcourt Brace Jovanovich Publishers, San Diego, CA, appendix A.

Tomasi, C. and T. Kanade, 1991, "Shape and Motion from Image Streams: a Factorization Method" Tech. Report CMU-CS-91-172, Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Uchiyama, M., E. Bayo, and E. Palma-Villalon, 1991, "A Systematic Design Procedure to Minimize a Performance Index for Robot Force Sensors," *Journal of Dynamic Systems, Measurement, and Control*, v.113, n.3, pp. 388-394.

Voyles, R.M., J.D. Morrow, and P.K. Khosla, 1995, "Shape from Motion Decomposition as a Learning Approach for Autonomous Agents," *Proceedings of the IEEE International Conf. on Systems, Man, and Cybernetics*, Vancouver, BC.

Voyles, R.M., B.L. Stavnheim, and B. Yap, 1989, "Practical Electrorheological Fluid-Based Fingertips for Robotic Applications," *Proc. of IASTED International Symp. on Robotics and Manufacturing*, pp. 280-283.