THE 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*. The novel aspect of this technique is that it does not require explicit knowledge of the redundant applied load vectors, yet it retains the noise rejection of a highly redundant data set and the rigor of least squares. The result is a much faster, slightly more accurate calibration procedure. A constant-magnitude force (produced by 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. Eliminating the need to collect all the applied loads makes collecting large amounts of calibration data nearly effortless. 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 (Watson and Drake, 1975; Shimano and Roth, 1977; Nakamura et al, 1988; Uchiyama et al, 1991), 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 introduced by measurement error; often 12 to 30 *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 (Watson and Drake, 1975; Shimano and Roth, 1977). For example, Lord Corporation allows up to 60 loads for their user calibration function (Lord Corporation, 1986).

Despite this burden, there has been little incentive to find alternative calibration methods. The least squares technique is mathematically sound (excepting its explicit assumption of zero-error load vectors, discussed later); it is accurate, given good experiment design and careful attention to implementation details; and, perhaps, most important to its widespread use, it is intuitive to casual users. Because the relatively high cost of least squares calibration is incurred only occasionally during the life of any given sensor, few researchers have pursued alternative calibration techniques.

Several researchers have tried to alleviate the burden of least squares by designing specialized calibration fixtures and procedures. Watson and Drake (1975) created a calibration table on which the sensor was mounted. The table had two movable pulleys that could be quickly and accurately positioned so that hanging weights applied accurately known forces and torques for least squares analysis. Uchiyama, et al (1991) created a similar but inverted apparatus. Rather than a table, they built a frame from which the sensor was suspended with an accurately adjustable "moment bar." Weights were suspended from the moment bar to create the forces and torques for least squares analysis. Shimano and Roth (1977) used the robot and gripper and known objects in the workspace as calibration fixtures for *in situ* calibration. The wrist was reoriented through a fixed set of discrete poses, some with the gripper empty, some while gripping a known object. The accuracy of this procedure is limited by the accuracy of the kinematics of the arm and grip points and the accuracy of the torques the arm can exert, however.

Assuming the calibration experiment has been properly designed, every calibration method benefits from more data in the presence of noise. The problem with the least squares technique, as previously mentioned, 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.

For a linear sensor, to which we restrict our attention (Bayo and Stubbe, 1989), the calibration function is a constant matrix that transforms $\mathbf{z}$ into $\mathbf{m}$.

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

Our calibration problem is to recover $\mathbf{C}$, the calibration matrix (not to be confused with the compliance matrix of Uchiyama et al (1991) and Nakamura et al (1988)) in the presence of two types of noise: errors in the applied load vector, $\tilde{\mathbf{m}}$, and measurement noise in the measurement vector, $\mathbf{z}$. 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.

Equation (1) assumes zero bias. If the sensor has some constant non-zero bias, this must be determined as part of the experimental procedure and subtracted from every reading. Again, careful experiment design, which is not the subject of this paper, is required to determine a procedure to accurately extract the bias for *any* calibration technique based on (1).

For force/torque sensors with symmetric internal mass distributions, sensor bias can be determined with a simple procedure in which the *z* axis is aligned with the gravity vector. The average of measurements in the parallel and anti-parallel orientations yields the bias vector. However, there are some conditions under which the determination of bias is difficult. An extension to the shape from motion technique to include bias determination is explored in Section 6.

Least Squares Calibration Solution

The standard technique for solving calibration problems is the least squares (LS) or pseudoinverse method. (Least squares and pseudoinverse are equivalent, Shimano and Roth, 1977, Strang, 1988.) 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, $\mathbf{Z}$:

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

Remember there are two types of noise: measurement noise and applied load error. The effect of measurement noise can easily be reduced by taking many measurements for each known load. Applied load error can only be reduced by carefully applying many different known loads.

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 (the "applied load error") must be minimized for an accurate calibration because this error manifests itself directly in the calibration matrix, $\mathbf{C}$. Most statistical analyses of least squares (i.e. Chatterjee and Hadi, 1988) assume $\mathbf{M}$ contains no errors. This is the same as assuming the applied load error is zero and leads to the result that $\mathbf{Z}^+\mathbf{M}$ is an unbiased estimate of $\mathbf{C}$. In fact, applied load errors can lead to a biased estimate of $\mathbf{C}$ (Beaton et al, 1976), which leads to inaccuracy in the calibration result *even if redundant loads are applied*.

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 applied loads 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 acquired*

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 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 (the "proper rank" or rank in the absence of noise, defined later) and applying arbitrary motion to the load.

Shape from Motion Derivation

In this section, we will derive the shape from motion approach in an abstract sense. Later, in Section 3, we will apply it to specific sensors and sensing spaces. Our derivation begins with a representation of the sensor function which maps a load onto a measurement:

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

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

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

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

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

SVD produces the following unique decomposition of any n x p matrix, $\mathbf{Z}$:

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

where $\mathbf{U}$ is an n x n orthogonal matrix, Σ is an n x p "diagonal" matrix (padded with zeroes as needed) of the singular values of $\mathbf{Z}$ in descending order, and $\mathbf{V}$ is a p x p orthogonal matrix.

Assuming we know that the "proper rank" of $\mathbf{Z}$ is r, it can be shown (Strang, 1988) that the best projection of $\mathbf{Z}$ onto an r-dimensional space (for $r \leq p$) is

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

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

We have assumed $\mathbf{Z}$ should have rank r (elaborated on below), but measurement noise provides additional independent information. Equation (8) gives us the best possible (Forsythe et al, 1977) rank-r representation of $\mathbf{Z}$ in the presence of that noise, so, combining (6) and (8) yields

$$\mathbf{Z}^* = \mathbf{U}^*\Sigma^*\mathbf{V}^{*\mathrm{T}} = \mathbf{MS} , \tag{9}$$

from which we can estimate $\mathbf{M}$ and $\mathbf{S}$:

$$\hat{\mathbf{M}} = \mathbf{U}^*(\Sigma^*)^{1/2}$$
$$\hat{\mathbf{S}} = (\Sigma^*)^{1/2}\mathbf{V}^{*\mathrm{T}} \tag{10}$$

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

$$\hat{\mathbf{M}}\hat{\mathbf{S}} = (\hat{\mathbf{M}}\mathbf{A}^{-1})(\mathbf{A}\hat{\mathbf{S}}) , \tag{11}$$

so we must find an appropriate matrix, $\mathbf{A}$, such that

$$\mathbf{M} = \hat{\mathbf{M}}\mathbf{A}^{-1}$$
$$\mathbf{S} = \mathbf{A}\hat{\mathbf{S}} \tag{12}$$

We indirectly find $\mathbf{A}$ by applying a geometric constraint to the individual vectors of the motion matrix (which we call the "motion constraint" and describe in Section 3.) to solve for $\mathbf{A}^{-1}$. Knowing $\mathbf{A}^{-1}$, finding $\mathbf{A}$ is trivial and we solve for $\mathbf{S}$ using (12) and $\mathbf{C}$ 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 pro-

cedure in flow chart form. (We present examples in Section 3 for clarification.) There are strong mathematical similarities between LS and shape from motion, but the essence of shape from motion is that we have replaced the LS requirement of knowing all individual loads with the shape from motion requirement of a geometric constraint on $\mathbf{M}$ with no assumptions of smoothness of motion.

Proper Rank

The "proper rank" of the matrix of output vectors is the rank of the matrix in the absence of noise. For a given sensor configuration, we must determine the proper rank before we can apply shape from motion.

We know the rank of the product of two matrices cannot exceed the rank of either individual matrix. (The product is both a subspace of the column space of one and a subspace of the row space of the other.) From equation (6), which does not include noise, we know the rank of $\mathbf{Z}$ is limited by the rank of the "motion" and "shape" matrices. From these we can deduce the proper rank of $\mathbf{Z}$ for any sensor.

We know little about the form of the shape matrix, but we can assume good sensor design will produce maximum rank. If not, the sensor will, in fact, be degenerate. The motion matrix, on the other hand, has a well-defined form. The motion matrix describes the motion of the force vector through euclidean space. Therefore, the rank will be either 2 or 3 for planar and 3-space sensors, respectively. From this, we know the proper rank of $\mathbf{Z}$.

Note that the rank is deduced from forces only. A 3-space wrist sensor can have up to six degrees-of-freedom, but the force vector remains embedded in 3-space. Torques are a linear combination of forces, so they do not increase the rank of the motion matrix. We will see that this complicates the shape from motion procedure, but does not cause failure.

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 similar, in concept, to that of Bicchi and Dario (1988). Force sensing results from the deformation of a single cantilever beam. There are four strain gages that respond to forces in the plane, thus, 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.

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

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

The constant magnitude of the force has been arbitrarily set to one unit, leaving only $\cos\theta$ and $\sin\theta$ in the motion matrix. (Correct engineering units will be selected in the final step.) This allows us to use $\cos^2\theta + \sin^2\theta = 1$ as our motion constraint. Furthermore, the rank of the shape matrix is, at most, 2, so the proper rank of our n x p measurement matrix, $\mathbf{Z}$, is also 2. If we denote the elements of $\mathbf{A}^{-1}$ by a_{11}, a_{12}, a_{21}, and a_{22}, and the ith row of $\hat{\mathbf{M}}$ by m_{i1} and m_{i2}, then the constraint equation becomes

$$\begin{aligned} 1 = {}& m_{i1}^2(a_{11}^2 + a_{12}^2) + 2m_{i1}m_{i2}(a_{11}a_{21} + a_{12}a_{22}) \\ & + m_{i2}^2(a_{21}^2 + a_{22}^2) \end{aligned} \tag{14}$$

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 (14) 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- and 3-DOF cases.) Because there are three equations and four unknowns, one variable is left free. Although any solution is acceptable, forcing A to be upper triangular or symmetric ensures invertibility. (Symmetric is preferred for good conditioning.) Having computed $\mathbf{A}^{-1}$, one can solve for the shape, $\mathbf{S}$, using (12) and the calibration matrix, $\mathbf{C}$, 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 $\mathbf{z}$, $\mathbf{m}$ pair, both vectors known) to rotate and scale it appropriately. To do this, simply apply $\mathbf{C}$ to $\mathbf{z}$ and use the resulting magnitude and angle to rotate and scale the matrix accordingly:

$$\mathbf{C}_o = Rot(\phi)\frac{\|\mathbf{m}\|}{\|\mathbf{Cz}\|}\mathbf{C} , \tag{15}$$

where $\mathbf{C}_o$ is the oriented calibration matrix, ϕ is the angular difference between $\mathbf{m}$ and $\mathbf{Cz}$, and $Rot(\phi)$ is the *2* x *2* rotation matrix.

We use Mathematica to perform the SVD and to solve the nonlinear equations to extract $\mathbf{A}$, $\mathbf{S}$, and finally, $\mathbf{C}$. 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) and we calibrated this 2-DOF sensor with only one known load.

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 [$\cos\theta\sin\psi$ $\sin\theta\sin\psi$ $\cos\psi$] 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 motion constraint (constant magnitude force application). This relation is easily derived so we will not repeat it here. (It is similar in form to (14).) 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.) This 3-DOF sensor requires only two known loads to fully scale and orient the matrix.

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

$$\begin{bmatrix} f_x & f_y & \tau \end{bmatrix} = \begin{bmatrix} f_x & f_y \end{bmatrix} \begin{bmatrix} 1 & 0 & -r_y \\ 0 & 1 & r_x \end{bmatrix} \tag{16}$$

Because the torque, τ, 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, as we would hope). This means a planar system with three DOF has a proper rank of 2, the same as the 2-DOF force-only case. Post-multiplying equation (16) by the true shape, $\tilde{\mathbf{S}}$, yields

$$\tilde{\mathbf{M}}\tilde{\mathbf{S}} = \mathbf{M} \begin{bmatrix} 1 & 0 & -r_y \\ 0 & 1 & r_x \end{bmatrix} \tilde{\mathbf{S}} = \mathbf{M}\mathbf{S}, \tag{17}$$

where $\tilde{\mathbf{M}}$ consists of vectors of the form [$f_x f_y$ τ] and **M** consists of vectors of the form [$f_x f_y$] (identical to the 2-DOF case). The moment arm becomes embedded in the shape matrix, **S** (rank 2), we can recover using shape from motion:

$$\mathbf{S} = \begin{bmatrix} \mathbf{I} & \mathbf{X}^{\mathrm{T}} \end{bmatrix} \tilde{\mathbf{S}} \tag{18}$$

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

From the 2-DOF force-only case, we have seen how to recover the $\mathbf{M}$ and $\mathbf{S}$ matrices of equation (17). However, 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 shape from motion calibration twice using two different moment arms and combine the results. As we have pointed out before, the two recovered shape matrices are not aligned with the desired reference frame, nor are they aligned with each other. 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. This is accomplished using the same technique that culminated in equation (15) for the 2-DOF case.

Once the squashed shapes are "co-oriented," we find $\tilde{\mathbf{S}}$ from (18) by conglomerating the shape and squashing matrices from both trials such that

$$
\begin{bmatrix} \mathbf{S}_1 \\ \mathbf{S}_2 \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{I} \ \mathbf{X}_1 \end{bmatrix} \\ \begin{bmatrix} \mathbf{I} \ \mathbf{X}_2 \end{bmatrix} \end{bmatrix} \tilde{\mathbf{S}}
\tag{19}
$$

or

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

where the subscripts 1 and 2 refer to the two calibration trials with different moment arms.

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).

$$\left[f_x \ f_y \ f_z \ \tau_x \ \tau_y \ \tau_z \right] = \left[f_x \ f_y \ f_z \right] \begin{bmatrix} 1 & 0 & 0 & 0 & r_z & -r_y \\ 0 & 1 & 0 & -r_z & 0 & r_x \\ 0 & 0 & 1 & r_y & -r_x & 0 \end{bmatrix} \tag{21}$$

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 (21).

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 (20) using co-oriented shape matrices:

$$\tilde{\mathbf{S}} = \begin{bmatrix} \left[\mathbf{I} \ \mathbf{X_1} \right] \\ \left[\mathbf{I} \ \mathbf{X_2} \right] \\ \left[\mathbf{I} \ \mathbf{X_3} \right] \end{bmatrix}^{+} \begin{bmatrix} \mathbf{S_1} \\ \mathbf{S_2} \\ \mathbf{S_3} \end{bmatrix}, \tag{22}$$

where $[\mathbf{I} \ \mathbf{X_i}]$ is as it appears in (21). Note that a total of 6 precise $(\mathbf{z}, \mathbf{m})$ 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," as defined by the expected forces resulting from the tasks for which the sensor was designed. We picked the upper extreme to improve signal-to-noise, but we stayed within the "reasonable range" to reduce any effects of minor nonlinearities. This is the extent to which we considered experiment design because we assumed the sensor is highly linear. Subsequent testing showed this assumption is valid. 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 to reduce the effects of measurement noise. We chose a large number of loads to reduce the effects of applied load error. 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 (including redundancy at each data point).

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 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 high frequency noise and determined the standard deviation of the smoothed signal. This quantifies the "imprecision" of the measurement across all orientations for a constant magnitude force. The best and worst of these data are tabulated in Table 1.

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 to collect data and compute moment vectors. (Without the calibration fixture, using our own version of the "torque bar" as in Uchiyama, et al, 1991, took even longer.) This compares to only 34 minutes for shape from motion to collect up to 500 measurement vectors using the single 33.71 N weight and compute the moment vectors. 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. Also, these times are aided by the use of the calibration fixture, which makes applying precise loads easier, primarily benefitting least squares.

Gathering the raw data for the shape from motion trials involved sampling the sensor at 2 Hz while it slowly moved up and down the longitudinal lines of an imaginary sphere. Each trial varied slightly in length so the total volume of data was different each time, but the raw data collection time was a small percentage (10-15%) of the total. The majority of time was consumed in gathering the precise loads used to orient the calibration matrix with the desired reference frame. This was done to the same precision as the least squares loads.

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 imprecision (standard deviation) is always better.

Figure 8 shows the magnitude of the linear forces of the one-bar experiment from Table 2. Ideally, this should be a straight line with a value of 25.24N including internal mass. We chose to illustrate this data set not because it makes shape from motion look best (in fact, other data sets show superior performance of shape from motion compared to least squares), 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 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.

5. SIMULATION STUDIES

To verify the above real experiments we ran dozens of simulations comparing least squares to shape from motion. For these simulations, we assumed a decoupled sensor with known calibration matrix:

$$\begin{bmatrix} 0.03 & 0 & 0 & 0 & -0.03 & 0 & 0 & 0 \\ 0 & 0 & 0.03 & 0 & 0 & 0 & -0.03 & 0 \\ 0 & 0.02 & 0 & 0.02 & 0 & 0.02 & 0 & 0.02 \end{bmatrix}$$

Simulated data was generated with various levels of independent, uniformly distributed noise in the sensor outputs and in the components of the load vector input. When both techniques were presented with the same large data sets of 210 loads (least squares using both load inputs and raw outputs, shape from motion using only the outputs), they performed nearly equivalently with a slight advantage to shape from motion. We expect the performance to be nearly the same under identical conditions because the underlying mathematics are closely related (both compute a least squares fit).

However, we ran the greatest number of simulations around an operating point similar to the real sensor experiments: 12 load vectors for least squares and 210 load vectors for shape from motion. Under these conditions, shape from motion demonstrated a clear advantage. With peak sensor noise at 1 percent of peak sensor output and an applied load error ball with a peak radius of $\sqrt{3}$ degrees, the variation in resolved force magnitude from the two calibration procedures is displayed in Figure 10. These plots are analogous to those with real data in Figure 8. The first line of Table 4 lists the mean and standard deviation data analogous to Table 2 with real data. The extra column for "load error" is the peak error for each component of the load vector used in the least squares technique. (Measurement noise was injected in all trials.)

These data were generated by simulating 110 raw sensor vectors with no noise and then applying the calibration matrices resulting from each calibration method to the raw vectors. Like the simulated data used to generate the calibration matrices and the real data of Section 4, this data was generated by moving the force vector along the longitudinal lines of an imaginary sphere. Although these data are from only two simulation trials, they are representative of all simulation trials we ran.

The second line of Table 4 shows a simulation run in which no errors were injected into applied loads for the least squares procedure. The error in the measurements stayed the same at 1 percent. In this case, shape from motion is still superior because of the large number of data points used (210 compared to 12). More data reduces the effects of noise in the measurements, as expected.

6. SENSOR BIAS

Sensor bias consists of a constant vector of offsets that is added to each and every sample. From equation (1) we see that both the least squares method and shape from motion as described both assume zero-bias.

In practice, this does not cause a problem because sensor bias is usually easy to determine experimentally. For a robotic wrist force/torque sensor, it is usually sufficient to take readings with the sensor pointing straight up and straight down and average them (i.e. Shimano and Roth, 1977).

In certain situations, for example an asymmetrical mass distribution (a gripper) that is inconvenient to remove, it would be convenient to determine sensor bias and the calibration matrix simultaneously. We show how LS can be augmented to include bias identification, and then present such an extension to the shape from motion technique that is similar to Tomasi and Kanade's original computer vision derivation [5].

Bias and Least Squares

Rewriting equation (1) in terms of the sensor function of Figure 1 yields:

$$\mathbf{z} = \mathbf{C}^+\mathbf{m} + \mathbf{z}_0,\tag{23}$$

where $\mathbf{z}_0$ is the constant bias vector and $\mathbf{C}^+$ is the pseudoinverse of $\mathbf{C}$. For every applied force, $\mathbf{m}$, the sensor response, $\mathbf{z}$, includes the additive offset, $\mathbf{z}_0$. By augmenting $\mathbf{C}^+$ and $\mathbf{m}$, we can write (23) as a strictly linear function:

$$\mathbf{z} = \begin{bmatrix} \mathbf{C}^+ & \mathbf{z}_0 \end{bmatrix}\begin{bmatrix} \mathbf{m} \\ 1 \end{bmatrix}\tag{24}$$

This requires two LS (pseudoinverse) procedures, however. The first extracts $\mathbf{C}^+$ and $\mathbf{z}_0$ while the second finds $\mathbf{C}$.

Bias and Shape from Motion

The idea of adding bias extraction to the shape from motion technique is analogous to the least squares case. For the single strain gage as in equation (12), adding a non-zero offset becomes:

$$\tag{25}$$
$$z_{ij} = Fs_{1j}\cos\theta_i\sin\psi_i + Fs_{2j}\sin\theta_i\sin\psi_i + Fs_{3j}\cos\psi_i + z_{0j}$$

which can be rewritten in accordance with equations (13) and (24):

$$\begin{bmatrix} z_{i1} & z_{i2} & \cdots & z_{i8} \end{bmatrix} = \begin{bmatrix} \cos\theta_i\sin\psi_i & \sin\theta_i\sin\psi_i & \cos\psi_i & 1 \end{bmatrix}\begin{bmatrix} s_{11} & s_{12} & \cdots & s_{18} \\ s_{21} & s_{22} & \cdots & s_{28} \\ s_{31} & s_{32} & \cdots & s_{38} \\ z_{01} & z_{02} & \cdots & z_{08} \end{bmatrix}\tag{26}$$

The effect of this is the increase of the rank of the motion matrix by one and the subsequent increase of the "proper" rank of $\mathbf{Z}$ by one. It is true that, except for specific pathological cases (such as all offsets equal to zero), increasing the columns of any rank deficient matrix by constant offsets will increase the rank of the matrix by exactly 1.

Here we have the first caveat: $\mathbf{Z}$ must be rank deficient. In other words, there must be a minimum of $n+1$ sense elements for an n-space sensor. In general, this does not present a problem because the critical n is not the degrees-of-freedom of the sensor but the dimensionality of the pure force vector. For example, the three-beam "modified maltese cross" commercial products from ATI (Little, 1992) which have only 6 sense elements for a 6-DOF sensor will work fine because the dimensionality of the force vector is 3 (not 6). In fact, *any* force/torque sensor will have the required number of sense elements. Only pure force sensors can potentially fall short of this requirement.

Armed with the new proper rank of $\mathbf{Z}$, the shape from motion procedure progresses in similar fashion as reported in Section 5 until the determination of $\mathbf{A^{-1}}$ from the motion constraint. Clearly, the motion constraint has changed, but also the size of $\mathbf{A^{-1}}$ has changed; it is now *4* x *4*.

Using the same nomenclature for $\mathbf{A^{-1}}$ and $\hat{\mathbf{M}}$ as in Section 3, we find we have two, decoupled constraints:

$$(\mathbf{m}_i^T \mathbf{a}_1)^2 + (\mathbf{m}_i^T \mathbf{a}_2)^2 + (\mathbf{m}_i^T \mathbf{a}_3)^2 = 1$$
$$\mathbf{m}_i^T \mathbf{a}_4 = 1$$

(27)

The second constraint is just the description of a plane and is easy to solve. The first constraint is trickier, though, and can be rewritten

$$\mathbf{m}^T \mathbf{B}^T \mathbf{B} \mathbf{m} = 1 \tag{28}$$

where $\mathbf{B^T}$ is a *4* x *3* submatrix consisting of the first three columns of $\mathbf{A^{-1}}$ and $\mathbf{m^T}$ is a row of $\hat{\mathbf{M}}$. This is a difficult nonlinear problem to solve that involves the fitting of data to a cylinder. We solve this problem by successively refining the solution numerically.

First, we note the similarity between (28) and the equation for an ellipse:

$$\mathbf{m}^T \mathbf{Q} \mathbf{m} = 1 \tag{29}$$

where $\mathbf{Q}$ is symmetric. Intuitively, an ellipse should give us a good estimate for the cylinder and we can use it as an initial starting point for the numerical refinement of the cylindrical fit.

Finding $\mathbf{Q}$ is a linear problem that we must solve in the least squares sense over all the rows of the matrix, $\hat{\mathbf{M}}$. Given $\mathbf{Q}$, we decompose it with SVD and set the smallest singular value to zero. This gives us a *4* x *3* matrix representing a cylinder that best fits the ellipse. Using this as a starting point, we numerically improve the solution using gradient descent over the mean squared error.

Concatenating the solution from the plane constraint and the solution from the cylindrical constraint yields a *4* x *4* matrix, $\mathbf{A^{-1}}$, that we can verify is invertible and well conditioned. With equation (12), we can reconstruct the augmented motion and shape matrices from equation (26) and strip off the offset vector from the shape matrix before taking the pseudoinverse that yields the calibration matrix.

7. CONCLUSIONS

We have described an improved calibration technique for linear force/torque sensors based on *shape from motion* decomposition that provides a slightly more accurate calibration matrix with much 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 at least as precise and accurate as 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. While it is possible to achieve equivalent precision and accuracy using least squares, the additional time required can be prohibitive for sensor manufacturers or researchers that must calibrate sensors often.

Although we determined bias vectors as part of the experimental procedure, we show that it is possible to augment the technique to automatically determine these as well, provided there are at least *n+1* sensing elements for a sensor in *n*-space (not to be confused with *n*-DOF).

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. (Although future work will relax this requirement. See below.) 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 that the *shape* of the sensor is completely extracted without the application of any known loads. The final step of applying a few known loads is only for scaling and orienting the shape matrix with respect to an external reference frame. The raw shape 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). Our ongoing research efforts

are more in this direction, toward fully autonomous calibration. We have also laid the theoretical groundwork for eliminating

the constant magnitude requirement by pressing two uncalibrated sensors together as in a multi-robot system or on the fingertips

of a dextrous hand.

8. ACKNOWLEDGEMENTS

9. REFERENCES

Bayo, E. and Stubbe, J.R., 1989, "Six-Axis Force Sensor Evaluation and a New Type of Optimal frame Truss Design for Robotic Applications," *Journal of Robotic Systems*, v.6, n.2, pp. 191-208.

Beaton, A.E., D.B. Rubin, and J.L. Barone, 1976, "The Acceptability of Regression Solutions: Another Look at Computational Accuracy," *Journal of the American Statistical Association*, v. 71, pp. 158-168.

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.

Chatterjee, S. and A.S. Hadi, 1988, *Sensitivity Analysis in Linear Regression*, John Wiley & Sons, New York, NY.

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.

Little, R., 1992, "Force/Torque Sensing in Robotic Manufacturing," *Sensors*, v.9, n.11.

Lord Corporation, 1986, *Installation and Operations Manual for F/T Series Force/Torque Sensing Systems*, Lord Corporation, Cary, NC, Appendix A.

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.

Shimano, B. and Roth, B., 1977, "On Force Sensing Information and its Use in Controlling Manipulators," *Proceedings of the IFAC International Symposium on information-Control Problems in Manufacturing Technology*, pp. 119 - 126.

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.

Van Brussel, H., H. Belien, and H. Thielemans, 1986, "Force Sensing for Advanced Robot Control," *Robotics*, v.2, n.2, pp. 139-148.

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*, v. 1, pp 407-412.

Watson, P.C. and S.H. Drake, 1975, "Pedestal and Wrist Force Sensors for Automatic Assembly," *Proceedings of the 5th International Symposium on Industrial Robots*, pp. 501-511.

List of Figure Captions

FIGURE 1: SENSOR AND CALIBRATION FUNCTIONS

FIGURE 2: SHAPE FROM MOTION CALIBRATION PROCEDURE

FIGURE 3: THE FINGERTIP SENSOR

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

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

FIGURE 6: RECOVERED MOTION OF LORD F/T SENSOR

FIGURE 7: CALIBRATION FIXTURE ON SENSOR FROM ASSURANCE TECHNOLOGIES

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

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

FIGURE 10: SIMULATED COMPARISON OF SHAPE FROM MOTION (TOP) AND LEAST SQUARES (BOTTOM) AS ORIENTATION CHANGES WITH UNIT-FORCE MASS

$\widetilde{\mathbf{m}}$ → | sensor | → $\mathbf{z}$ → | calibration function | → $\mathbf{m}$

Figure 1

collect lots of raw data (no knowledge of loads)

$\mathbf{Z}$

SVD($\mathbf{Z}$)

$\hat{\mathbf{M}}$, $\hat{\mathbf{S}}$

Find $\mathbf{A}^{-1}$ to enforce constraint ← motion constraint

$\mathbf{M}$, $\mathbf{S}$

Orient result to match user-specified frame ← ($\mathbf{z}$, $\mathbf{m}$) data pairs

$\mathbf{C}$
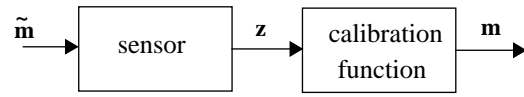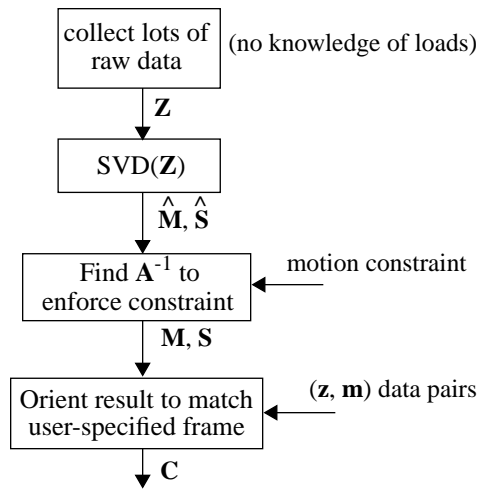
Figure 2

Figure 3 USE Photograph previously provided
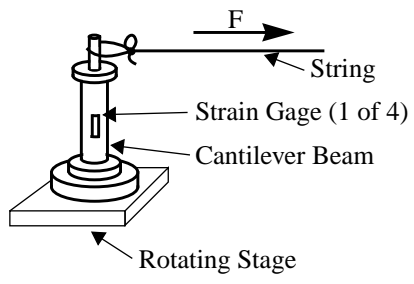
Figure 4
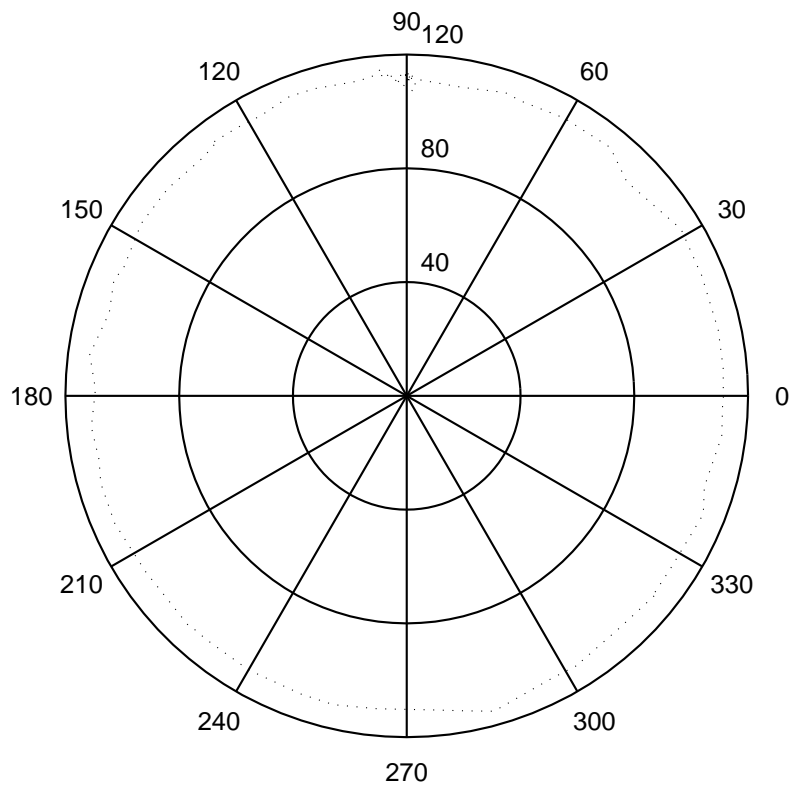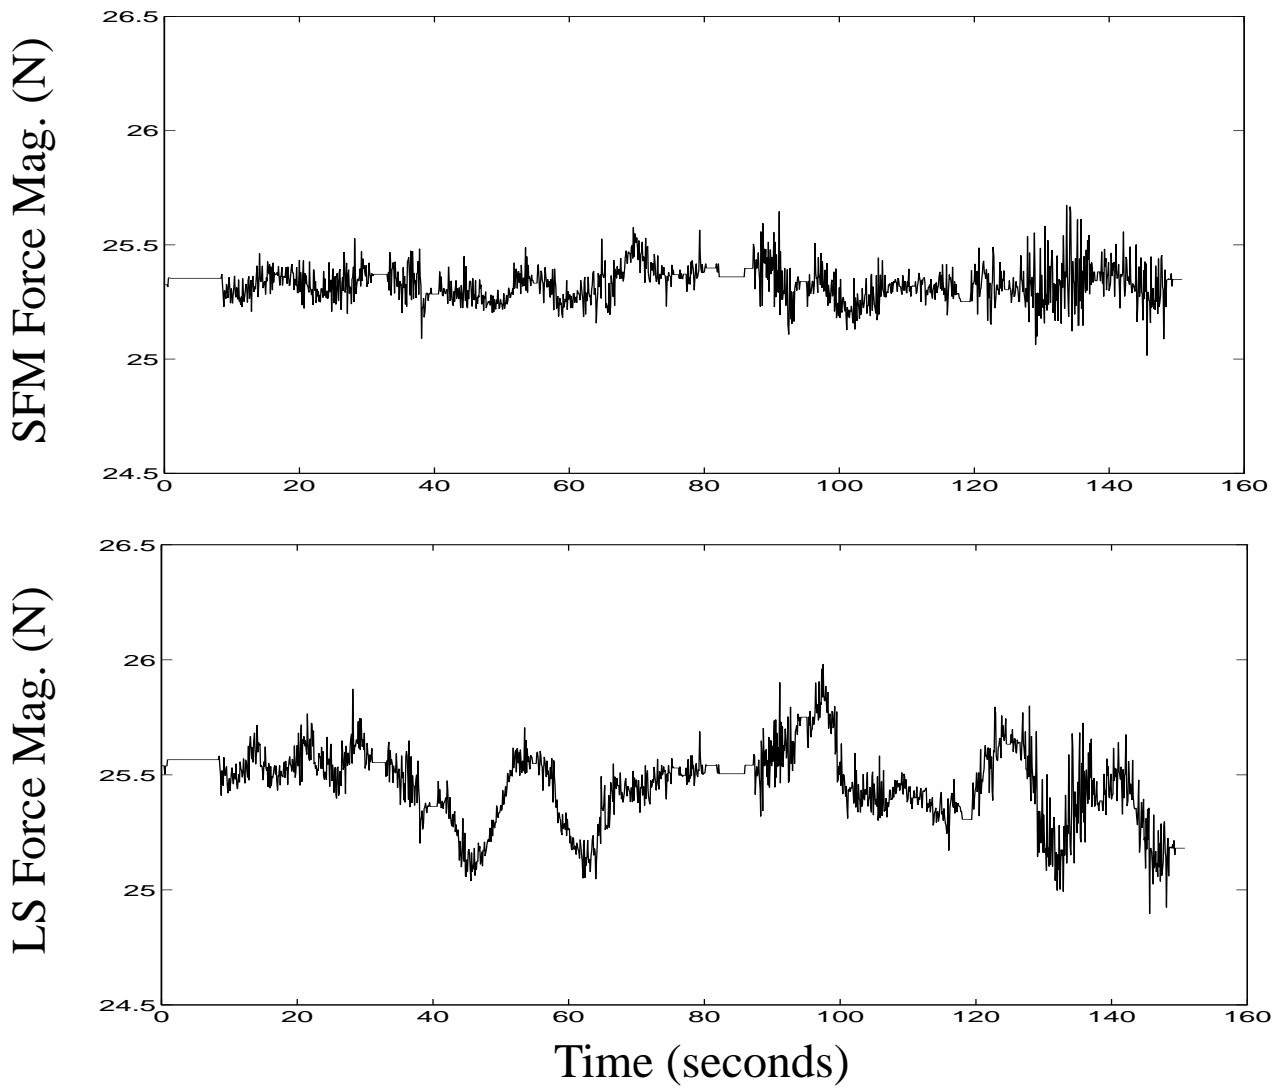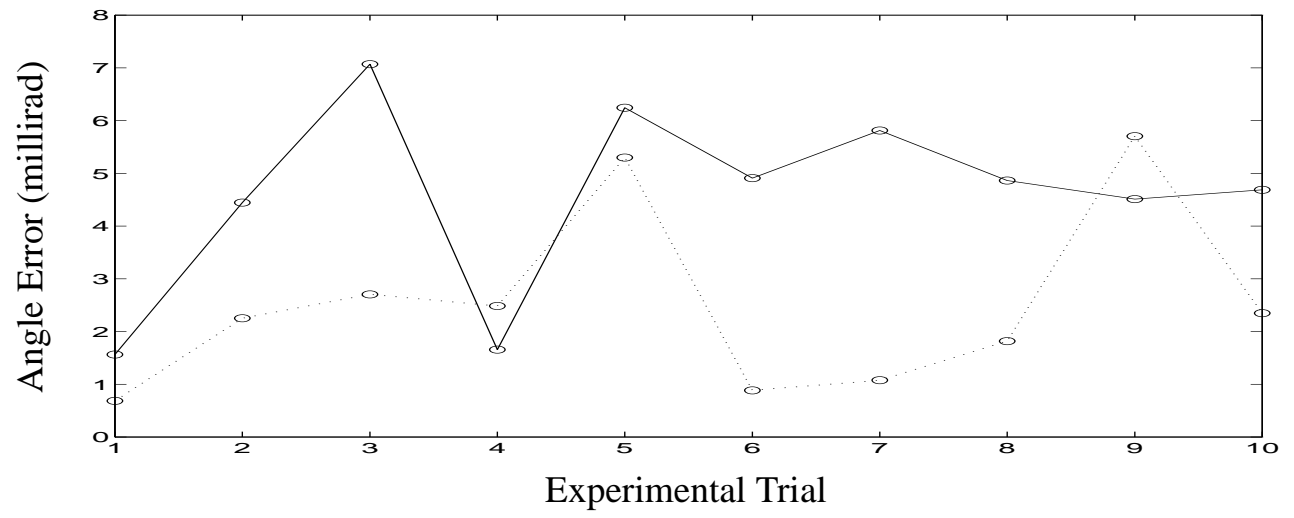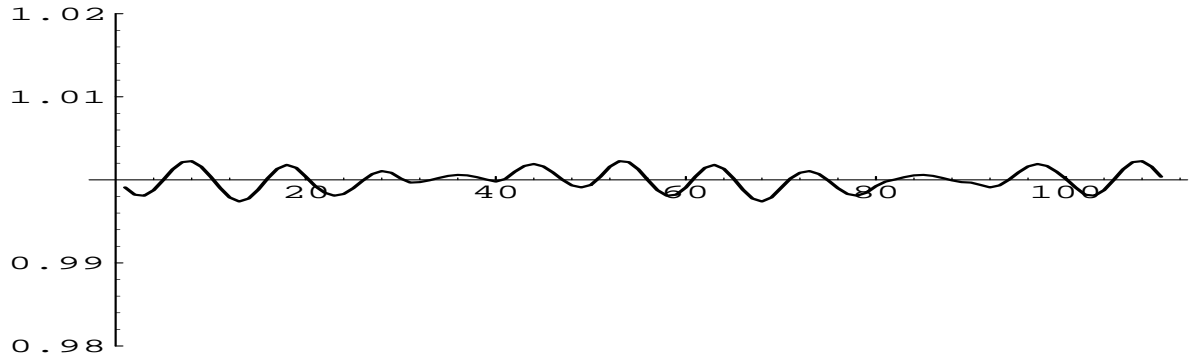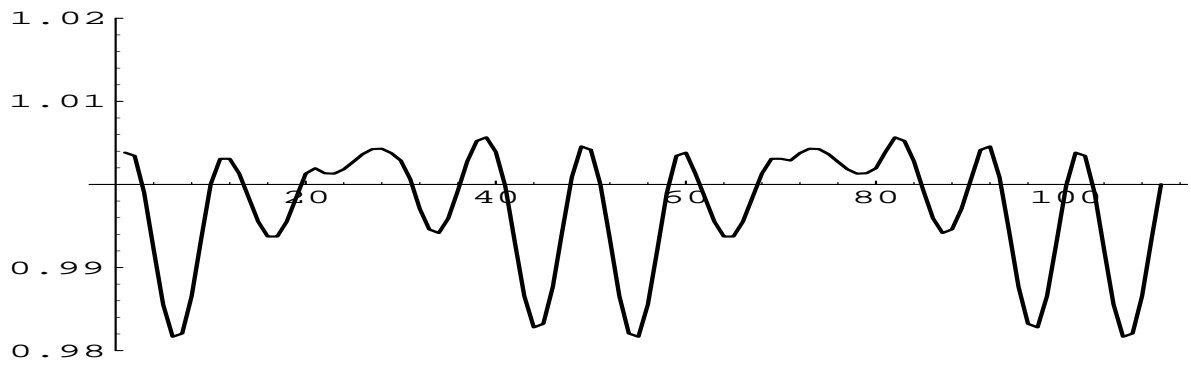
Figure 5
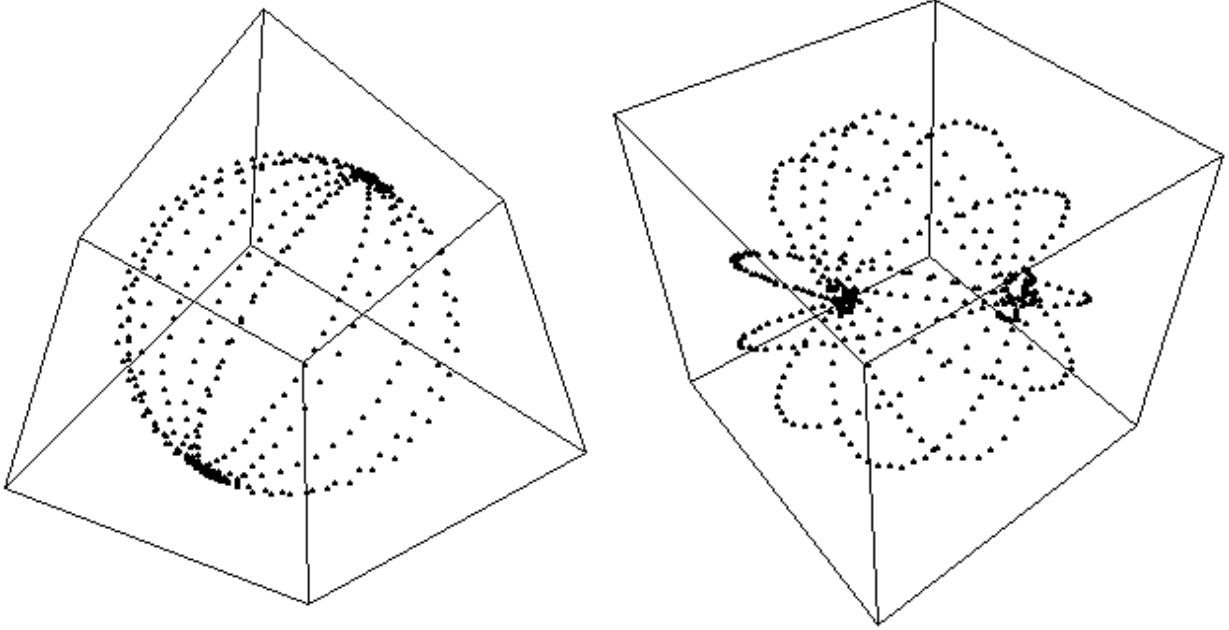
Figure 6
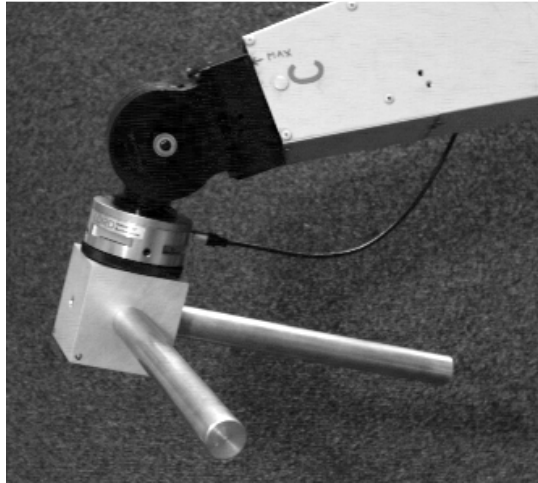
Figure 7 USE Photograph previously provided

Figure 8

Figure 9

Figure 10

List of Table Captions

Table 1: FINGERTIP SENSOR COMPARISON
Table 2: LORD F/T SENSOR COMPARISON
Table 3: AVERAGE LORD FORCE ERRORS
Table 4: SIMULATION COMPARISON

Table 1

load (g)	Least Squares				Shape from Motion			
	mean (g)	error (%)	imprecision (std dev)	time (min)	mean (g)	error (%)	imprecision (std dev)	time (min)
60.4	60.1	0.5	1.20	17	60.5	0.2	1.13	1.5
111.4	110.2	1.1	0.83		111.6	0.2	0.60	

Table 2

	Least Squares				Shape from Motion			
load (N)	mean (N)	error (%)	imprecision (std dev)	time (min)	mean (N)	error (%)	imprecision (std dev)	time (min)
9.26	9.09	1.8	.084	69	9.09	1.8	.058	34
25.24	25.53	1.7	.157		25.41	0.7	.062	

Table 3

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

Table 4

load		Least Squares			Shape from Motion		
mag. (N)	error (N)	mean (N)	error (%)	imprecision (std dev)	mean (N)	error (%)	imprecision (std dev)
1.0	.017	.9971	.29	.0071	1.00001	.001	.0013
1.0	0.0	1.0003	.03	.0041	.9998	.02	.0016