

Including Sensor Bias in Shape from Motion Calibration and Sensor Fusion

Richard M. Voyles, J. Daniel Morrow, and Pradeep K. Khosla

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
<http://www.ri.cmu.edu>

Abstract

Shape from Motion data fusion brings a greater degree of autonomy and sensor integration to intelligent systems in which fusion by constant linear transformations is appropriate. To illustrate this, we apply Shape from Motion techniques to applications involving both similar and disparate sensory information vectors. First, nearly autonomous force/torque sensor calibration is demonstrated through fusion of the individual channels of raw strain gauge data. Gathering 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. This calibration example is provided to simply explain the mathematics. Disparate sensory information is fused in a “primordial learning” mobile robot through a similar eigenspace representation. This paper summarizes these shape from motion applications and presents an extension for simultaneously extracting sensor bias.

1 Introduction

The Shape from Motion sensor fusion technique originated from a basic need: more efficient sensor calibration strategies. We initially wanted a quicker method to calibrate a multitude of sensors for multi-fingered robotic and human hands. After developing this novel technique for sensor calibration that eliminates much of the need for known applied loads, we realized it could be used in a more general framework for data fusion involving constant linear transformations. While it is possible to reformulate the Shape from Motion approach for incremental rather than batch processing, it will not be discussed here.

In an abstract sense, calibration involves extracting (or “learning”) the “shape” of the linear mapping between the input and output spaces of groups of sensors. The first sections of this paper describe Shape from Motion calibration and the extensions to sensor bias in this context. Section 9 describes a “primordial learning” application in which a mobile robot uses disparate data fusion to learn to free itself from jams.

2 Motivation for Shape from Motion

Multi-axis force/torque sensor calibration -- the process of recovering the linear calibration matrix mapping raw sensor values to resolved forces and torques -- is a data fusion problem that has traditionally been a very costly

procedure in terms of time. The reason for this is the traditional methods based on reduction by least squares require two sets of data: input and output. Output data is very easy to gather so incurs little cost. Input data, on the other hand, can be quite expensive.

The output data consists of a vector of raw sensor values. For common wrist force/torque sensors based on a maltese cross configuration, such as the JR3 or Lord commercial products, this vector consists of eight strain gauge values. Collecting this vector is a simple matter of data acquisition.

The input data set is quite different. It represents the applied forces and torques that produce each and every output vector. For multi-axis sensors, this is a complex combination of vectors that must be carefully measured with respect to an external reference frame. These careful measurements in 3-space are very time consuming but are of utmost importance in recovering an accurate calibration matrix.

In an ideal world, n separate measurements are required for an n -dimensional sensor. However, there are two primary sources of noise: measurement noise in acquiring the output and inaccuracy in applying the input. In order to recover an accurate calibration matrix, it is desirable to gather as much data as possible; all things being equal, more data results in greater accuracy. Unfortunately, each piece of data comes at high cost due to the input.

Shape from Motion calibration is a new technique that virtually eliminates the need for the input data set. Instead of gathering both inputs and outputs, *outputs only* are gathered. Because collecting outputs is cheap, large numbers of samples can be collected in substantially less time than that required by the least squares technique. This large volume of data is desirable as it provides good immunity to noise, as mentioned above. An additional benefit is the elimination of one source of error: inaccuracy in application of the input.

3 Shape from Motion Procedure

We will develop the shape from motion procedure using the calibration example for simplicity. Later, the fusing of disparate sensory data will be considered.

The shape from motion calibration procedure is illustrated by the flow chart in Figure 1. Data collection consists of attaching a constant mass to the sensor and

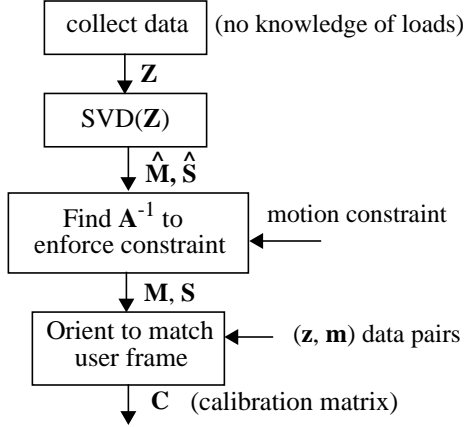


Figure 1 :Shape from Motion Calibration Procedure

moving it randomly through its entire space of measurement, rapidly sampling the output data stream and storing the vectors in matrix, \mathbf{Z} . This reorienting of the sensor/mass moves the applied gravity vector throughout the space. Singular value decomposition (SVD) [3] of the resulting output data matrix produces two matrices that contain all the information of the *motion* of the force vector and the *shape* of the sensor. A hard trigonometric constraint is applied to extract the “motion” and “shape” matrices explicitly, and, finally, a few measured input/output data pairs are used to orient the calibration matrix with the user-specified reference frame. These known input/output data pairs re-introduce some noise due to inaccurate input, but the noise is introduced *after* the shape has been determined, when it is less destructive.

This technique is based on “shape and motion decomposition” techniques pioneered by Tomasi and Kanade in computer vision [10]. The beauty is in the recognition of the analogous “shape” and “motion” matrices as they apply to calibration problems. While this technique is closely related, mathematically speaking, to the least squares approach, it is startlingly different in concept and there is no previous work of which we are aware (other than our own) that demonstrates calibration in this way. A few researchers have attempted to formalize the least squares calibration process in terms of SVD. Uchiyama, et al [11] have done the most comprehensive treatment along this vein, but Bicchi and Dario [1] provided early foundations. These formalizations still require complete input and output data sets, unlike Shape from Motion.

4 Algorithmic Details

4.1 The Calibration Problem

A sensor converts an applied load, $\tilde{\mathbf{m}}$, into a measurement vector, \mathbf{z} (Figure 2). 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.

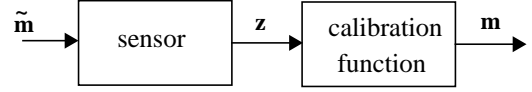


Figure 2 : Sensor and Calibration Functions

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

$$\mathbf{C}\mathbf{z} = \mathbf{m} \text{ or } \mathbf{z}^T \mathbf{C}^T = \mathbf{m}^T \quad (1)$$

The initial calibration problem is to recover \mathbf{C} , the calibration matrix assuming sensor bias (offset) has been removed.

4.2 Least-Squares Calibration Solution

The standard technique for solving calibration problems is the least-squares 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):

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

from which the calibration matrix can be computed using the pseudoinverse of the measurement matrix, \mathbf{Z} :

$$\mathbf{C}^T = \mathbf{Z}^+ \mathbf{M} \quad (3)$$

4.3 Shape from Motion Calibration

With the shape from motion calibration approach *we do not need to know all the individual loads, \mathbf{m}_i* , but only a constraint which relates them. As a result, a large number of loads can be applied cheaply and quickly to improve accuracy. A small number of known loads are required to establish the user-specified reference frame, but usually fewer than n for an n -DOF sensor. (6-DOF is the exception; 6 known loads are required for a 6-DOF sensor.)

4.4 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:

$$\mathbf{z}_i^T = \mathbf{m}_i^T \mathbf{S}, \quad (4)$$

where \mathbf{z}_i^T is a $l \times p$ measurement vector, \mathbf{m}_i^T is a $l \times m$ load vector, and \mathbf{S} is the $m \times p$ 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

$$\mathbf{C} = [\mathbf{S}^T]^+ \quad (5)$$

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

$$\mathbf{Z} = \mathbf{M}\mathbf{S}, \quad (6)$$

where \mathbf{Z} is the $n \times p$ matrix of measurements and \mathbf{M} is the $n \times m$ 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} . The shape from motion 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 SVD on \mathbf{Z} .

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

$$\mathbf{Z} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \quad (7)$$

where \mathbf{U} is an $n \times n$ orthogonal matrix, $\mathbf{\Sigma}$ is an $n \times p$ “diagonal” matrix¹ of the singular values of \mathbf{Z} in descending order, and \mathbf{V} is a $p \times p$ orthogonal matrix.

Assuming we know that the “proper” rank² of \mathbf{Z} is r (discussed later), it can be shown [8] that the best projection of \mathbf{Z} onto an r -dimensional space (for $r \leq p$) is

$$\mathbf{Z}^* = \mathbf{U}^*\mathbf{\Sigma}^*\mathbf{V}^{*T}, \quad (8)$$

where \mathbf{U}^* consists of the first r columns of \mathbf{U} , $\mathbf{\Sigma}^*$ is a diagonal matrix of the first r singular values, and \mathbf{V}^{*T} consists of the first r rows of \mathbf{V}^T .

Equation (8) gives us the best possible [3] rank- r representation of \mathbf{Z} , so, combining (6) and (8) yields

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

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 $r \times r$ matrix, \mathbf{A} , (an affine transform)

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

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

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

1. padded with zeroes as needed

2. the rank in the absence of noise

We do this by applying a trigonometric constraint to the individual vectors of the motion matrix and solving for \mathbf{A}^{-1} . Once \mathbf{A} is known, we can solve for \mathbf{S} using (11) and \mathbf{C} using (5). Finally, we introduce a few precise measurements (\mathbf{z}_i , \mathbf{m}_i pairs) in order to orient the calibration matrix with respect to the desired reference frame and to scale the result to the desired engineering units.

5 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.

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} .

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.

6 3-DOF Force Sensor Calibration

The shape from motion calibration technique is illustrated more completely in [12], including experimental comparisons to the least squares technique on both 2-DOF and 6-DOF robotic sensors. We will briefly outline the application of this technique to 3-DOF force-only sensors as a basis for our extension to sensor offsets. For the complete 6-DOF force/torque case, see [12].

As an example of a 3-DOF force sensor, consider a standard Lord wrist sensor and constrain forces to the center of the flange by attaching a compact proof mass. Assuming the moment arm is negligible, we can look at the response of a single strain gauge half-bridge to a load of magnitude, F in accordance with equation (4)

$$z_{ij} = Fs_{1j}\cos\theta_i\sin\psi_i + Fs_{2j}\sin\theta_i\sin\psi_i + Fs_{3j}\cos\psi_i \quad (12)$$

which can be rewritten:

$$\begin{bmatrix} z_{i1} & z_{i2} & \dots & z_{i8} \end{bmatrix} = \begin{bmatrix} \cos\theta_i \sin\psi_i & \sin\theta_i \sin\psi_i & \cos\psi_i \end{bmatrix} * \quad (13)$$

$$\begin{bmatrix} s_{11} & s_{12} & \dots & s_{18} \\ s_{21} & s_{22} & \dots & s_{28} \\ s_{31} & s_{32} & \dots & s_{38} \end{bmatrix}$$

The constant magnitude of the force has been factored out, leaving only $[\cos\theta_i \sin\psi_i \sin\theta_i \sin\psi_i \cos\psi_i]$ in the motion matrix and i is the index over all samples. From this we can see that the rank is 3 (three independent columns) and we also get our motion constraint. The magnitude of a unit vector is always 1 so the constraint sets the sum of the squares of the columns of the motion matrix to 1. If we enforce this constraint, the motion matrix will automatically result as will the complementary shape matrix.

To enforce this constraint we rely on equation (11) and select a symbolic, arbitrary 3×3 matrix to represent \mathbf{A}^{-1} . The columns of \mathbf{A}^{-1} are \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{a}_3 and the i th row of $\hat{\mathbf{M}}$ is \mathbf{m}_i^T so

$$\left(\mathbf{m}_i^T \mathbf{a}_1\right)^2 + \left(\mathbf{m}_i^T \mathbf{a}_2\right)^2 + \left(\mathbf{m}_i^T \mathbf{a}_3\right)^2 = 1 \quad (14)$$

which results in only 6 equations for the 9 unknown elements of the \mathbf{A}^{-1} matrix for each row of the matrix, $\hat{\mathbf{M}}$. With three free parameters, convenient choices include either making the matrix upper diagonal or symmetric. The actual solution is found in the least squares sense over all rows, \mathbf{m}_i^T , to minimize noise.

Having computed \mathbf{A}^{-1} , one can solve for the shape, \mathbf{S} , using (11) and the calibration matrix, \mathbf{C} , using (5).

Unfortunately, the resulting calibration matrix is not oriented with respect to any particular reference frame and the weight of the proof mass has been arbitrarily set to 1 unit. To align it with our desired reference frame (the sensor body), we introduce two precise measurements (\mathbf{z} , \mathbf{m} pairs), to rotate and scale the calibration matrix appropriately.

To briefly summarize the experimental results presented in [12], calibrating a 6-DOF Lord force/torque sensor took half as long with Shape from Motion and produced a slightly more accurate result. Table 1 tabulates the results of attaching a known load to the sensor and slowly reorienting it throughout space comparing both the Least Squares and Shape from Motion calibration matrices. "Error" is the mean measurement compared to the actual load. "Smooth" is the standard deviation of the smoothed force measurement, which provides a measure of consistency.

7 Sensor Bias

Sensor bias consists of a constant vector of offsets that is added to each and every sample. From equation (1) we see

Table 1: LORD F/T SENSOR COMPARISON

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

that both the least squares method and shape from motion as described both assume zero-bias. The reason is that a system with bias is not strictly linear, but affine.

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.

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 present such an extension to the shape from motion technique that is similar to Tomasi and Kanade's original computer vision derivation [9].

7.1 A Single Strain Gauge

For the single strain gauge as in equation (12), adding a non-zero offset becomes:

$$z_{ij} = F s_{1j} \cos\theta_i \sin\psi_i + F s_{2j} \sin\theta_i \sin\psi_i + F s_{3j} \cos\psi_i + z_{0j} \quad (15)$$

which can be rewritten in accordance with equation (13):

$$\begin{bmatrix} z_{i1} & z_{i2} & \dots & z_{i8} \end{bmatrix} = \begin{bmatrix} \cos\theta_i \sin\psi_i & \sin\theta_i \sin\psi_i & \cos\psi_i & 1 \end{bmatrix} * \quad (16)$$

$$\begin{bmatrix} s_{11} & s_{12} & \dots & s_{18} \\ s_{21} & s_{22} & \dots & s_{28} \\ s_{31} & s_{32} & \dots & s_{38} \\ z_{01} & z_{02} & \dots & z_{08} \end{bmatrix}$$

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 force

vector. For example, the new “modified maltese cross” commercial products from ATI which have only 6 sense elements for a 6-DOF sensor will work fine because the dimensionality of the force vector is 3. 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 6 until the determination of \mathbf{A}^{-1} from the motion constraint. Clearly, the motion constraint has changed, but also the shape of \mathbf{A}^{-1} has changed; it is now 4×4 .

Using the same nomenclature for \mathbf{A}^{-1} and $\hat{\mathbf{M}}$ as in Section 6, we find we have two, decoupled constraints:

$$\left(\mathbf{m}_i^T \mathbf{a}_1\right)^2 + \left(\mathbf{m}_i^T \mathbf{a}_2\right)^2 + \left(\mathbf{m}_i^T \mathbf{a}_3\right)^2 = 1 \quad (17)$$

$$\mathbf{m}_i^T \mathbf{a}_4 = 1$$

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

$$\mathbf{m}^T \mathbf{B}^T \mathbf{B} \mathbf{m} = 1 \quad (18)$$

where \mathbf{B}^T is a 4×3 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 (18) and the equation for an ellipse:

$$\mathbf{m}^T \mathbf{Q} \mathbf{m} = 1 \quad (19)$$

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 4×3 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 to the solution from the cylindrical constraint yields a 4×4 matrix, \mathbf{A}^{-1} , that we can verify is invertible and well conditioned. With equation (11), we can reconstruct the augmented motion and shape matrices from equation (16) and strip off the offset vector from the shape matrix before taking the pseudoinverse that yields the calibration matrix.

8 Simulations

We ran simulations to test the extended shape from motion algorithm on data sets with varying levels of noise

and bias. We used an idealized 3-DOF force sensor with a uniform calibration matrix:

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

to generate 176 strain gauge data points. To this we added different levels of uniformly distributed noise and various random but constant offsets to each channel. To evaluate each run, we look at the offsets extracted and examine the mean of the magnitude of the motion and also the standard deviation of the smoothed magnitude. We smooth the magnitude before taking the standard deviation to accentuate any inaccuracies in the shape.

Trial 1:

Offsets

Actual: 3.0 2.5 -2.5 2.0 4.0 -2.5 5.0 -3.0
Recovered: 2.98 2.50 -2.59 2.00 4.02 -2.50 5.09 -3.00
Error (%): 2.7 0 3.6 0 0.5 0 1.8 0

Peak Noise = 0 (0%)

Std. Dev.

SFM with offset calc.: 0.00026
SFM: 0.011

Mean

SFM with offset calc.: 0.9998
SFM: 0.9997

Trial 2:

Offsets

Actual: -10.0 12.5 12.5 14.5 17.5 -12.5 7.5 -17.5
Recovered: -12.7 12.3 12.9 14.4 20.1 -12.7 7.1 -17.5
Error (%): 27.0 1.6 3.2 0.7 14.9 1.6 5.3 0

Peak Noise = 3 (2%)

Std. Dev.

SFM with offset calc.: 0.0077
SFM: 0.039

Mean

SFM with offset calc.: 0.992
SFM: 0.995

Trial 3:

Offsets

Actual: -30.0 25.0 25.0 20.0 35.0 -25.0 1.0 -39.0
Recovered: -48.7 26.7 34.3 20.1 54.2 -25.1 -8.3 -38.1
Error (%): 62.3 6.8 37.2 0.5 54.9 0.4 2.3

Peak Noise = 15 (10%)

Std. Dev.

SFM with offset calc.: 0.052
SFM: 0.101

Mean

SFM with offset calc.: 0.964
SFM: 0.968

Trial 4:

Offsets

Actual: -30.0 25.0 25.0 20.0 35.0 -25.0 1.0 -39.0

Recovered: -66.7 23.3 41.1 18.5 70.8 -27.5 -16.7 -42.1
Error (%): 122 6.8 64.4 7.5102 10.0 7.9
Peak Noise 15 15 25 15 15 15 25 15
Std. Dev.

SFM with offset calc.: 0.092
SFM: 0.093

Mean

SFM with offset calc.: 0.927
SFM: 0.970

To put the noise and offsets in perspective, the maximum signal for all illustrated trials was about 150, so the noise and offset values illustrated are in the 0 - 20% range.

Trials 3 and 4 were chosen for inclusion to illustrate a potential problem with noise sensitivity for relatively high levels of noise. Trial 3 holds noise at about 10% of signal across all sense elements with offsets set at just less than twice that level. The only thing that changes in Trial 4 is the noise level on only two sense elements is raised to about 17% of signal. Offset recovery is affected rather strongly on some sense elements.

9 Disparate Sensory Data

Force sensor calibration is a successful application of shape from motion sensor fusion with an explicit external reference and similar sensory data. We don't think this exhibits the full power of the paradigm, however. The above example is mired in the need to transform the intrinsic shape to an explicit vector space. As we hinted, for a machine to "learn," it only needs to develop its own internal representation of the "shape" (in the abstract sense) of the interaction.

Primordial robot learning -- learning fundamental interactions with no prior knowledge -- has been suggested for mobile robot behavior specifically (e.g.[6][7]) and is similar to non-parametric learning in general. To this end, we have applied the shape from motion fusion technique to a small mobile robot that has no explicit knowledge of its limited set of actuators and sensors. All the robot is "aware of" is the streams of data that come from or go to the sensors and actuators. The shape from motion technique allows the robot to develop an internal representation of teleoperated interaction between sensors and actuators to produce "meaningful" externalized behavior. In this sense, the robot is primordial, or infant-like, because it must learn the most basic input/output relationships from a teacher, fuse them as required to mimic the behavior, and ignore superfluous data.

The shape from motion technique is modified to match the principal component analyses of Hancock and Thorpe [4] and Pierce [7]. The standard principal components analysis provides a substitute for the geometric constraint used in equation (14). For primordial learning, we do not wish to constrain the problem. Although some constraints can be useful to achieve more sophisticated results (as in

[7]), they impose the potential limitations of a pre-supposed shared ontology which may not be expressive enough for the situation at hand.

The subject of this experiment is the MK-V, a three-wheeled, non-holonomic mobile robot with a single drive wheel and an unpowered steering wheel. Sensors include wheel encoders, a compass, bump sensors and motor current sensors. The drive and steering motors have three-valued -- forward-reverse-off -- commands with no closed-loop controllers on velocity or position. There are also no redundant groups of very similar and correlated sense elements such as a visual retina or sonar ring.

With no algorithmic structure imposed at all on the fusion process, the result learned is dependent on what the robot is allowed to observe. For training, we teleoperated the robot in a cluttered environment allowing it to wander while bumping into obstacles.

The training data was batch-processed using SVD to extract the eigenvectors and selected the largest n eigenvectors using the largest ratio of adjacent singular values as the threshold for n . For run-time operation (sensor fusion in real-time), the new sensor image is projected onto the eigenspace as described in [4].

Using this technique, the real robot achieved both wandering behavior and escaped stalls and collisions without highly redundant or correlated sensors and with no algorithmic structure imposed on the learned result. Figure 3 shows a teleoperated run of the MK-V with comparisons of what the human did and what the MK-V would have done on its own with eigenvectors learned on a prior training set. The drive motor commands of the Shape from Motion technique closely match the commands of the human. Furthermore, the steering motor learned *when* to steer, but the steering *direction*, which is random and superfluous to getting un-stuck, does not match. This fused result is appropriate; superfluous information and, in fact, entire

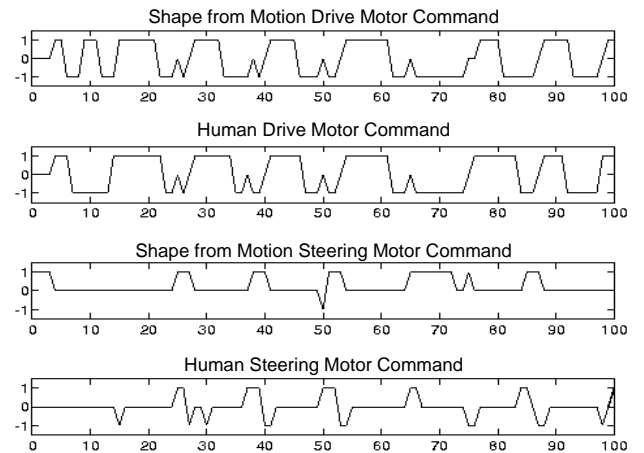


Figure 3: Comparison of Human and Shape from Motion Commands on a Teleoperated Dataset.

superfluous sensors were ignored to achieve the trained behavior.

10 Discussion

This paper first summarized the shape from motion calibration approach for force/torque sensors reported in [12]. Calibration is a simple form of data fusion and shape from motion is a novel approach to sensor calibration that dramatically reduces the need for accurately applied load vectors. This, in turn, dramatically speeds up the calibration process, allowing the collection of much larger sets of data resulting in slightly more accurate calibration. (More experimental evidence appears in [12].)

We have not discussed the effects of experimental design on calibration accuracy. This topic is beyond the scope of this work and is better left to the classical references in the field such as [2]. Suffice it to say, no volume of additional data can make up for poor experiment design and shape from motion calibration is no exception. The success of the technique is dependent on proper excitation of the measurement space.

Regarding the extension of shape from motion fusion to include sensor bias, there are a few caveats. First is the dependence on $n+1$ sense elements for an n -space sensor. In practice, this is not a significant problem because n is the dimensionality of the force vector, not the degrees-of-freedom of the sensor. All force/torque sensors have at least $n+1$ sense elements.

A much bigger problem is the susceptibility of the process to noise at reasonably high levels. If the offset is small with respect to noise, the $n+1$ -th singular value will contain a large noise component. This can rapidly deteriorate the accuracy of the calibration result. In fact, the robustness of determining the offsets by a separate experimental procedure specifically tailored for extracting the offsets will almost always be superior to simultaneous extraction using shape from motion.

We also noted from the simulation results that even when the offsets are greater than the noise, significant levels can infiltrate the motion and shape matrices. Fortunately these effects were observed with noise in the realm of 10 - 20% of signal and should not be much of a concern for most robotic applications. Most commercial sensors we have tested have noise well below the 1% level.

This paper also summarized the results of an experiment to fuse disparate sensory information in the context of mobile robot learning. This approach replaced the geometric constraint with a more traditional principal components analysis to derive the fusion coefficients. Comparisons of the results of the Shape from Motion weighting matrix to human effort demonstrate successful fusion of the sensory data to replicate the intended behavior.

In closing, the Shape from Motion fusion approach is effective for finding constant coefficients for fusing similar or disparate data by linear transformation. While it is possible to re-formulate the Shape from Motion procedure in an incremental manner which would allow dynamic fusion coefficients for sensory information, to date, we have not explored this option.

References

- [1] Bicchi, A. and P. Dario, "Intrinsic Tactile Sensing for Artificial Hands," in *Proc. of the International Symposium on Robotics Research*, R.C. Bolles and B. Roth, editors, MIT Press, Cambridge, MA, 1988, pp. 83-90.
- [2] Diamond, W., *Practical Experiment Designs for Engineers and Scientists*, Lifetime Learning Pub., 1981.
- [3] Forsythe, G.E., M. Malcolm, and C.B. Moler, *Computer Methods for Mathematical Computations*, Prentice Hall, Englewood Cliffs, NJ, 1977.
- [4] Hancock, J.A. and C.E. Thorpe, "ELVIS: Eigenvectors for Land Vehicle Image System," Tech. report CMU-RI-TR-94-43, Robotics Institute, Carnegie Mellon Univ., Pitt., PA, Dec. 1994.
- [5] Klema, V.C. and A.J. Laub, "The Singular Value Decomposition: Its Computation and Some Applications," *IEEE Transactions on Automatic Control*, v. 25, n. 2, pp. 164-176, April, 1980.
- [6] Lin, L-J, "Scaling Up Reinforcement Learning for Robot Control," 8th International Workshop on Machine Learning, Evanston, IL, 1991.
- [7] Pierce, D., "Learning a Set of Primitive Actions with an Uninterpreted Sensorimotor Apparatus," 8th International Workshop on Machine Learning, Evanston, IL, 1991, pp 338 - 342.
- [8] Strang, G., *Linear Algebra and Its Applications*, third edition, Harcourt Brace Jovanovich Publishers, San Diego, CA, 1988, appendix A.
- [9] Tomasi, C. and T. Kanade, "Shape and Motion from Image Streams: a Factorization Method, Part1, Planar Motion" Tech. Report CMU-CS-90-166, Comp. Sci., Carnegie Mellon University, Pittsburgh, PA, Sept. 1990.
- [10] Tomasi, C. and T. Kanade, "Shape and Motion from Image Streams: a Factorization Method" Ph.D. Thesis, CMU-CS-91-172, Computer Science, Carnegie Mellon University, Pittsburgh, PA, Sept. 1991.
- [11] Uchiyama, M., E. Bayo, and E. Palma-Villalon, "A Systematic Design Procedure to Minimize a Performance Index for Robot Force Sensors," *Journal of Dynamic Systems, Measurement, and Control*, vol. 113, no. 3, pp. 388-394, Sept. 1991.
- [12] Voyles, R.M., J.D. Morrow, and P.K. Khosla, "Shape from Motion Approach to Rapid and Precise Force/Torque Sensor Calibration," ASME International Mechanical Engineering Congress and Exposition, San Francisco, CA, Nov. 1995.