In Proceedings of the 1995 IEEE Conf. on Systems, Man,
and Cybernetics, Vancouver, BC, Oct. 1995, vol. 1

Shape from Motion Decomposition as a Learning Approach for Autonomous Agents

Richard M. Voyles, Jr.[†] J. Dan Morrow[†] Pradeep K. Khosla[‡]

[†]Robotics Ph.D. Program
[‡]Dept. of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

This paper explores Shape from Motion Decomposition *as a learning tool for autonomous agents.* Shape from Motion *is a process through which an agent learns the "shape" of some interaction with the world by imparting motion through some subspace of the world. The technique applies singular value decomposition to observations of the motion to extract the eigenvectors. We show how shape from motion applied to a fingertip force sensor "learns" a more precise calibration matrix with less effort than traditional least squares approaches. We also demonstrate primordial learning on a primitive "infant" mobile robot.*

1 Introduction

To be useful in dynamic or unknown environments, autonomous agents must be able to learn about the environment in which they're immersed -- in effect, internally "calibrate" themselves with respect to it. But calibration, in the traditional sense, is anything but autonomous. It usually requires external stimuli from a previously calibrated agent with careful coordination of measurements. This dependence on other agents in order to function reduces the autonomy of the agent and ultimately impacts its robustness.

To *learn autonomously*, the agent must derive a relevant model of some interaction with the environment with minimal reliance on external knowledge. Ideally, a fully autonomous agent would develop models of interaction by exciting variables over which it has control (through actuators) and observing the effect on variables in the world (through sensors), much the way a human infant learns its limbs are part of its body by flailing around.

Yet, calibration and learning are fundamentally similar; they both aim to represent interactions between variables. The key difference is the basis for the representation. Calibration relies on an explicit "shared ontology" -- or common representation -- so the meaning of the output is consistent across all agents. For example, *calibrating* a force sensor involves generating a representation that provides a force vector with respect to some known cartesian reference frame. However, if the ultimate task of the agent is to maintain contact during complex manipulation, for example, it may be more appropriate to learn abstract input/output relationships that achieve this behavior rather than going through the transformations required by the explicit ontology. In fact, most learning problems are cast in such a way that a shared ontology is enforced by the system designer (hence, the equivalence to calibration), although artificial neural networks and principal component analyses generally avoid this problem.

Our technique is derived from a novel approach in computer vision and was first applied to a traditional calibration problem. We describe briefly the computer vision approach in Section 2 because it provides an intuitive basis for discussion. The connection to force sensor calibration is qualitatively discussed in Section 3 while the technique is rigorously developed in Section 4 followed by experimental results in Section 5. Switching from constraint analysis to principal component analysis allows *primordial learning* in Section 6.

2 Shape and Motion in Computer Vision

Shape and motion factorization is a recent and successful technique in computer vision for extracting the motion and shape of objects from a sequence of monocular images. It is successful and robust because it does not rely on first recovering the depth of the objects (which is problematic), nor does it rely on constrained or known motion. Instead, it employs singular value decomposition (SVD) [4] to simultaneously decompose the highly redundant image feature data into two matrices. The rank of these matrices is fixed by the number of degrees of freedom of the system and a fixed geometrical relationship between them. From this known relationship, the two matrices are transformed into the *motion* over time and *shape* in space of the object in the image stream.

In the context of computer vision, the *shape* of the unknown object (the *x-y-z* coordinates of each tracked feature) and the motion of the object relative to the camera are simultaneously recovered without calibration. Tomasi and Kanade [10] first did this by tracking a large number of image features (corresponding to points on the object) through a large number of image frames during unstructured, relative motion between the object and camera. From this highly redundant set of data, SVD produces the singular values and two unique orthonormal matrices from which the relative motion and shape can be extracted with no assumption of smoothness of either motion or shape.

It is important to note that "shape" and "motion" are defined relative to an arbitrary vector space and have concrete meaning only when particular spaces are chosen. In the abstract sense, the matrices just represent a linear mapping of data points; a specification of the structure (shape) and history (motion) of some *interaction*.

3 Calibration by Shape from Motion

The traditional force sensor calibration procedure[7] is rather far from the ideal learning scenario. Many known, externally-calibrated forces are applied to the sensor, then corresponding output vectors are gathered, and the input/output data is reduced by least squares. As we envision the ideal autonomous learning scenario, no externally-supplied calibration forces would need to be supplied. The agent would "properly excite" the environment by flailing around and a calibration matrix would "magically" result.

Shape from Motion calibration is not quite that simple, but all that's required is for the agent to pick up a known mass in a gravity field and then randomly re-orient the sensor while gathering the resulting output vectors from the sensor. No information on the orientation of the sensor in the gravity field is required! Only two known, external loads must be applied to provide a calibrated reference frame for a 3-DOF sensor. Shape from Motion *simultaneously* extracts the calibration matrix and the previously unknown motion of the sensor within the gravity field.

Force sensor calibration is analogous to the computer vision approach described above. Given p sensing elements (analogous to feature points) sampled f times under different *unknown* load conditions, one can produce a similar redundant data set. The key is in recognizing that the movement of the applied calibration forces through the space of measurement corresponds to the object/camera *motion* while the calibration matrix is merely an abstract mathematical description of the *shape* of the sensor. These analogous *shape* and *motion* descriptors can be extracted from the measurement matrix using a similar technique.

"Shape," in the context of force sensors, can be construed both literally and figuratively. In the literal sense, the shape is an encoding of the mechanical structure of the device, the positioning of the sense elements, and the properties of the material from which it is made. (This will be explored with respect to a particular sensor in Section 4.) In the figurative sense, the shape is a mapping from the p-dimensional sensing space to the n-dimensional space of resolved force/torque (where $p \geq n$). The shape can be thought of as a projected hyper-ellipsoid that describes this mapping.

Figure 1: The reconfigurable force sensor shown with both inactive and electrorheological outer surfaces.

4 Theory of Shape from Motion

In this section, we will incrementally develop the shape from motion calibration technique for the 2-degree-of-freedom (DOF) and 3-DOF cases. The full extension to 6-DOF is reported in [11].

The 2-D Force Sensor

As a starting point for proof of concept, we first analyzed a 2-DOF subset of a fingertip sensor of our own design (Figure 1). The intrinsic force sensor is only one component of a complete tactile fingertip that includes a surface array tactile sensor and actuator based on an electrorheological fluid [12]. We chose a 2-D starting point because it very closely matches the shape and motion analysis presented in [10]. It also provides an excellent example for explaining the theoretical basis of the shape from motion approach.

Our fingertip force sensor is a cylindrical cantilever beam with strain gages as the sensing elements. For the 2-D experiments, we used the four axial gages that are approximately equally spaced around the circumference of the beam. Point loads were applied in a plane perpendicular to the cylinder's axis by a string and pulley arrangement (Figure 2) to avoid applied torques.

The voltage, z_i, from the ith sense element is:

$$z_i = g_i \beta_i F \cos\left(\theta - \alpha_i\right) \tag{1}$$

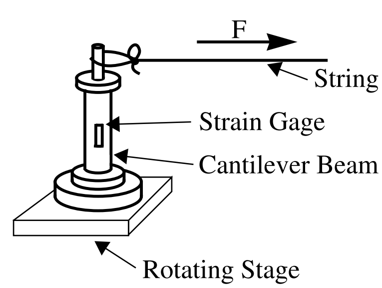

Figure 2: Test setup for the 2-D fingertip sensor.

where g_i is the electrical gain of the system, β_i is the mechanical gain, F is the radial applied force and the angles are as shown in Figure 3.

The two gain variables are lumped parameters. The electrical gain, g_i, includes the amplifier gain, excitation voltage, gage factor, and voltage divider or bridge network. The mechanical gain, β_i, includes the structure of the sensing beams (i.e. cantilever beam, maltese cross, etc.), the radius and length of the beam, the modulus of elasticity of the material, the axial placement of the strain gage along the beam, and the alignment of the strain gage axis with respect to the cylinder's axis. Exact formulae relating these parameters can be found in [1].

Expanding (1) using the double angle formula yields:

$$z_i = g_i \beta_i F \cos(\alpha_i) \cos(\theta) + g_i \beta_i F \sin(\alpha_i) \sin(\theta) \tag{2}$$

which can be simplified to

$$z_i = s_{1i} \cos(\theta) + s_{2i} \sin(\theta) \tag{3}$$

if we keep the magnitude of the force constant (easily achieved by using a constant mass on the string).

Equation (3) describes one strain gage, the *i*th gage. Given p strain gages, we can combine the z_i's to produce a p-vector of measurements for one time sample. For the jth time sample, assuming $p = 4$ (for our fingertip sensor), we can rewrite (3) in matrix form:

$$\begin{bmatrix} z_{j1} & z_{j2} & z_{j3} & z_{j4} \end{bmatrix} = \begin{bmatrix} \cos\theta_j & \sin\theta_j \end{bmatrix} \begin{bmatrix} s_{11} & s_{12} & s_{13} & s_{14} \\ s_{21} & s_{22} & s_{23} & s_{24} \end{bmatrix} \tag{4}$$

If we gather f samples (or "frames" as in [9], described at the beginning of Section 3 of this paper), we can construct an f x p matrix, **Z**, from the strain gage measurements; an f x 2 matrix, **M**, that collects the sine and cosine terms from each applied force; and a 2 x p matrix, **S**, that collects all the gains from each strain gage. This is written

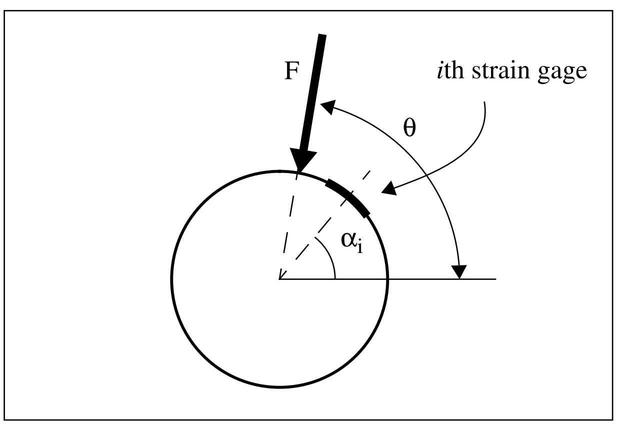

Figure 3: Cross section of the cantilever beam with pure force applied.

$$\mathbf{Z} = \mathbf{MS} \tag{5}$$

and exactly corresponds both to equation (4), above, and the planar images of [6]. We choose the symbol **M** because the sine and cosine of the force vector describe its *motion*. (In this simplified case, we know the origin of the vector and it is confined to the plane.) We choose the symbol **S** because we view these gains as describing the *shape* of the sensor, as mentioned previously. Thus, the gage measurements result from the product of the motion of the force vector and the shape of 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 rather 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 decomposing **Z** through SVD.

SVD produces the unique decomposition of an m x n matrix, **Z**, such that:

$$\mathbf{Z} = \mathbf{U\Sigma V}^T \tag{6}$$

where **U** is an m x m orthonormal matrix, Σ is an m x n "diagonal" matrix[1] of the singular values of **Z** in descending order, and **V** is an n x n orthonormal matrix. We will show how to extract the right-hand side of (5) from the right-hand side of (6).

To do this robustly in the face of noise, it is imperative that we know something about the rank of **Z**. Since **M** and **S** are f x 2 and 2 x p, respectively, neither can have rank greater than two. As a result, *the rank of* **Z** *can be at most two*, in the absence of noise. The degenerate cases, when the rank is less than two, result when all forces are colinear and will be ignored.[2]

Noise injected into a matrix adds dimensionality or additional "structure" to the data that does not exist in reality. The SVD provides an elegant mechanism for eliminating this noise. First of all, the first two singular values of our measurement matrix, **Z**, will be much larger than the others. This is due to the fact that it is fundamentally of rank two. If it were exactly rank two (no noise), all singular values besides the first two would be zero. It can be shown [2] that the best projection of this data onto a two-dimensional space (in other words, the best rank-2 matrix) will be:

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

1. padded with zeroes as needed
2. Bad sensor design can also result in degenerate cases.

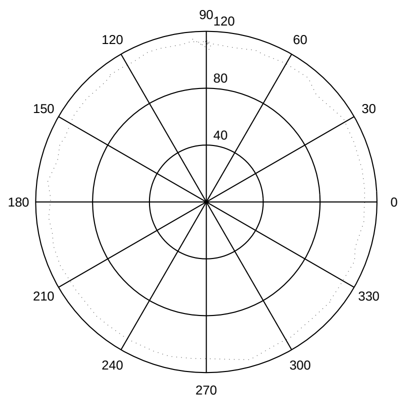

Figure 4: Polar plot of the recovered motion (magnitude and direction) with 111.4g mass.

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

Since equations (5) and (7) both embody the best rank-2 representation of the data, they can be equated:

$$\begin{aligned} \mathbf{MS} &= \mathbf{U}^*\Sigma^*\mathbf{V}^{*T} = \hat{\mathbf{M}}\hat{\mathbf{S}} \\ \hat{\mathbf{M}} &= \mathbf{U}^*(\Sigma^*)^{1/2} \\ \hat{\mathbf{S}} &= (\Sigma^*)^{1/2}\mathbf{V}^{*T} \end{aligned} \tag{8}$$

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 *2* x *2* matrix, **A**:

$$\hat{\mathbf{M}}\hat{\mathbf{S}} = \left(\hat{\mathbf{M}}\mathbf{A}^{-1}\right)(\mathbf{A}\hat{\mathbf{S}}), \tag{9}$$

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

$$\begin{aligned} \mathbf{M} &= \hat{\mathbf{M}}\mathbf{A}^{-1} \\ \mathbf{S} &= \mathbf{A}\hat{\mathbf{S}} \end{aligned} \tag{10}$$

We do this by enforcing a metric constraint on the motion. Since **M** consists of cosθ and sinθ, we can use $(\cos\theta)^2 + (\sin\theta)^2 = 1$ as our constraint. Denoting the elements of $\mathbf{A}^{-1}$ by a_{11}, a_{12}, a_{21}, and a_{22}, and the *j*th row of $\hat{\mathbf{M}}$ by m_{j1} and m_{j2}, the constraint equation becomes:

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

Based on this, we can construct a matrix of the vectors $[m_{j1}^2 \; 2m_{j1}m_{j2} \; m_{j2}^2]$ and solve for the combinations of a_{ji} in (11) using the pseudoinverse.

We use Mathematica code to perform the SVD and to solve the nonlinear equations to extract **A**, **S**, and finally, **C**. Figure 4 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 5.)

3-DOF Force Sensors

For the case of three forces, we took a standard 6-DOF Lord force/torque sensor and mounted a compact mass on the face plate. In reality, there is a small moment arm associated with this arrangement, but we ignore that. Assuming pure forces are acting at the origin of the sensor, superposition applies, just as in equation (12):

$$z_i = g_i\beta_i F\cos\left(q - \alpha_i\right) + g_i''\beta_i''F\cos\psi \tag{12}$$

where ψ represents the angle the force vector makes with respect to the *z*-axis and $g_i''\beta_i''$ represent the electrical and mechanical gains relating force in the *z*-axis to strain.

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. This new constraint is easily derived so we will not repeat it here. A plot of the recovered motion of the Lord force/torque sensor during an actual calibration trial appears in Figure 5. Although it is difficult to judge by eye, it is an accurate sphere. (See [11] for more quantitative results.)

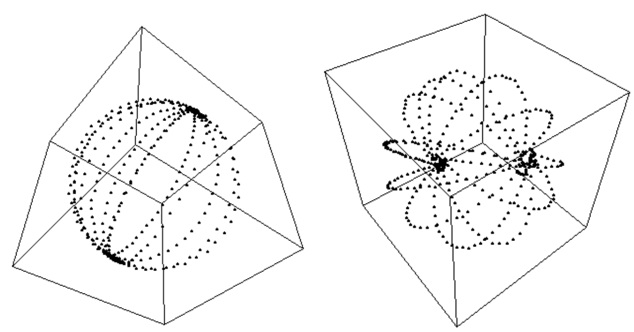

Figure 5: Recovered motion of the Lord F/T sensor.

The structure in Figure 5 can be recovered without applying the constraint equation. In that case, what is lost is a known reference frame. Since the external reference frame is dictated only by the shared ontology, it seems natural that the *shape* can be exploited to extract all the meaningful content of the interaction up to an affine transformation. We hope to use this to learn tactile sensorimotor skills in the future [6][9].

5 Experimental Comparison to Least Squares

The point of calibration, with its shared ontology, is to provide an accurate characterization of the sensor relative to a known frame. But desired accuracy is almost always traded-off with the time required to perform the calibration procedure. Our motion and shape approach is exceptional in that it produces a *more accurate and precise* calibration, for a given accuracy in calibration measurement, in *less time* as compared to the least squares approach.

It takes less time because fewer accurately applied loads are required. For a given accuracy of each measurement, fewer of them results in less time.

The reason for greater accuracy is more subtle. 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 (15), **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 them 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 also allows us to gather more data in less time.

To demonstrate these assertions, we performed both shape from motion and least squares calibrations on our 2-DOF fingertip sensor. For the shape from motion approach, we chose a single mass of 111.4 grams and completed the calibration procedure in 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 least squares calibration masses 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 in favor of shape from motion. 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 noise and determined the standard deviation of the smoothed signal. This assesses the precision of the measurement across all orientations. The extremes 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

6 Primordial Robot Learning

Force sensor calibration is a successful application of shape from motion learning with an explicit external reference. 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 approached by many researchers (e.g.[5][8]). To this end, we have applied the shape from motion 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 the interaction between sensors and actuators to produce "meaningful" external behavior. In this sense, the robot is infant-like because it must learn the most basic input/output relationships from scratch.

The shape from motion technique is modified to match the principal component analyses of Hancock and Thorpe [3] and Pierce [8]. The standard principal components analysis provides a substitute for the geometric constraint used in equation (11). For primordial learning, we do not wish to constrain the problem. Although some constraints can be useful to achieve more sophisticated results (as in [8]), they impose the potential limitations of a shared ontology.

We have two small robots that are very simple in construction with minimal, but disparate, sensing and actuation. The subject of this experiment is the MK-V, a three-wheeled, non-holonomic structure 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 learning 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.

We batch-processed the training data 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, the new sensor image is projected onto the eigenspace as described in [3].

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 6 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 matches the human quite nicely as this is more deterministic. The steering motor, on the other hand, does not match steering *direction*, which is random, but does know *when* to steer.

A benefit of this technique is that the mapping onto the eigenspace produces not only the command vector for the actuators, but a reconstruction of the sensor input, as well. This reconstruction can be compared to the actual sensor vector to derive a measure of the confidence in the correctness of the command values for the given input.

7 Discussion

We have described an improved calibration technique for linear force/torque sensors based on *shape from motion* decomposition that learns 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. The random motion and the pseudoinverse of the calibration matrix are simultaneously recovered from the raw sensor values. Two known loads are sufficient for an external reference frame for the matrix in 3-D. In the sense that little external information is required to "discover" the matrix, it is an appropriate learning paradigm for autonomous agents.

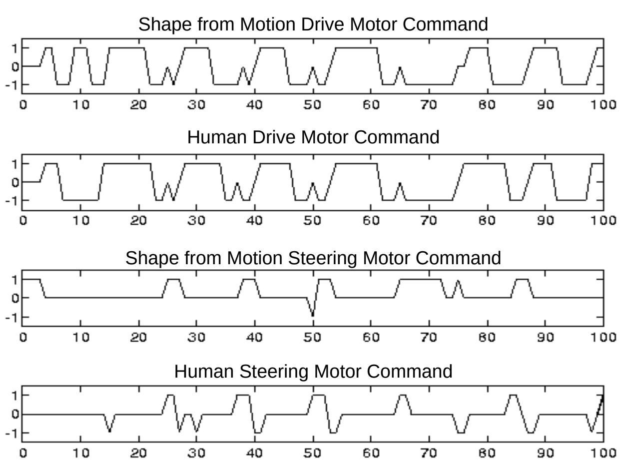

Figure 6: Comparison of Human and Shape from Motion Commands on a Teleoperated Dataset.

Without an external reference, the shape can still be recovered up to an affine transform. Although the internal reference frame is unknown, the matrix still holds meaningful information on the "shape" of the interaction. To that end, our first experiments with the MK-V mobile robot used eigenspace projection. We hope to more fully exploit internally-referenced shape matrices in learning tactile manipulation skills.

8 Acknowledgments

Mr. Voyles was supported in part by a National Defense Science and Engineering Graduate Fellowship, a DOE Integrated Manufacturing Predoctoral Fellowship and by Sandia National Labs contract NAG1-1075. Mr. Morrow was supported in part by a DOE Computational Science Graduate Fellowship.

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

[3] 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.

[4] 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.

[5] Lin, L-J, "Scaling Up Reinforcement Learning for Robot Control," 8th International Workshop on Machine Learning, Evanston, IL, 1991.

[6] Morrow, J.D. and P.K. Khosla, "Sensorimotor Primitives for Robotic Assembly Skills," in *Proceedings of 1995 IEEE International Conf. on Robotics and Automation*, Nagoya, Japan, May, 1995.

[7] Nakamura, Y., T. Yoshikawa, and I. Futamata, "Design and Signal Processing of Six-Axis Force Sensors," in *Robotics Research: The 4th International Symposium*, R. Bolles and B. Roth, eds., 1988, pp 75-81.

[8] 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.

[9] Rucci, M. and P. Dario, "Autonomous Learning of Tactile-Motor Coordination in Robotics," in *Proceedings of 1994 IEEE International Conf. on Robotics and Automation*, San Diego, CA, May, 1994, v. 4, pp 3230 - 3236.

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

[11] Voyles, R.M., J.D. Morrow, and P.K. Khosla, "The Shape from Motion Approach to Rapid and Precise Force/Torque Sensor Calibration," to appear in *Proceedings of the ASME International Mechanical Engineering Congress and Exposition*, San Francisco, CA, Nov. 1995.

[12] Voyles, R.M., B.L. Stavnheim, and B. Yap, "Practical Electrorheological Fluid-Based Fingertips for Robotic Applications," *Proc. of IASTED International Symposium on Robotics and Manufacturing*, pp. 280-283, Santa Barbara, CA, Nov. 1989.