

Collaborative Calibration: Extending Shape from Motion Calibration

Richard M. Voyles and Pradeep K. Khosla

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

Abstract

In this paper, we summarize recent research results from a new technique of sensor calibration called Shape from Motion Calibration. We first present the basic technique, which is an eigenspace analysis, and show that it includes the rigor of least squares without the full burden of measuring all the applied inputs. Next we present new research that removes another constraint of the calibration technique and extends the robustness to cover systems with slight nonlinearities.

1 Introduction

Sensors are normally calibrated by applying carefully measured inputs and recording the corresponding outputs. If one wants to calibrate a scale, for example, the scale is zeroed then a known mass is placed on the scale. If the scale is linear, the slope and offset are easily calculated and the calibration is complete.

Admittedly, this is a highly simplified description that ignores noise, nonlinearities, range of operation, and other issues of proper design of experiments. Nonetheless, in many cases, this is sufficient to meet the accuracy requirements of a given task. Even with the difficulties of the real world, such as noise, etc., the calibration problem is usually addressed by taking several measurements, as above, and reducing by least squares. Although there is more data involved, the basic calibration procedure remains the same: apply a known reference and take a measurement.

Voyles, Morrow, and Khosla proposed *Shape from Motion Calibration* as a new approach to force/torque sensor calibration [8] based on work in computer vision by Tomasi and Kanade [5][6]. This technique blurs the distinction between calibration and learning by eliminating the need for constraints on the input space. There still is a model-based constraint (which can be relaxed for other applications, as we will demonstrate), but the need to know the inputs *a priori* is eliminated. Instead, unknown inputs are randomly applied, producing a *slightly more accurate* calibration in *significantly less time* than the standard least squares approach.

In this paper we present a summary of the shape from motion technique including prior research results. We then present new work that removes another constraint from the basic technique to allow near-fully autonomous calibration of pairs of force sensors on dextrous hands or multiple-

manipulator installations. Finally, we summarize work in robot skill extraction applied to both a mobile robot and a PUMA manipulator.

2 Shape from Motion Calibration

The beauty of shape from motion calibration is that it is so simple to perform. Because it is not necessary to record all the applied forces, the sensor can be waved randomly throughout the space with a mass attached to generate applied force vectors. It is only necessary to collect the raw sensor output data, sampled rapidly during excitation. The applied force vectors correspond to the “motion” of the input and are extracted by the algorithm during processing. The motion is not recorded explicitly by the user as in least squares.

The resulting sampled output vectors are stacked in a matrix, \mathbf{Z} . A structured form of principle components analysis (PCA), extracted by singular value decomposition (SVD), provides the *shape* (pseudoinverse of the calibration matrix) that results from the induced *motion* (sequence of direction vectors) of the input force vector. Finally, a few known input/output data pairs are used to orient the calibration matrix with the user-specified reference frame and scale to the desired engineering units. The number of known loads required to establish the user-specified reference frame and scale factor is usually fewer than n for an n -DOF sensor.

3 Algorithmic Details

3.1 The Calibration Problem

A sensor converts an applied load, $\tilde{\mathbf{m}}$, into a measurement vector, \mathbf{z} . The purpose of calibration is to “learn” a transformation that, given a measurement vector, estimates the load which generated it.

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

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

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

3.2 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:

$$\mathbf{z}_i^T = \mathbf{m}_i^T \mathbf{S}, \quad (2)$$

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 as the pseudoinverse of the shape matrix, \mathbf{S} :

$$\mathbf{C} = [\mathbf{S}^T]^+ . \quad (3)$$

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

$$\mathbf{Z} = \mathbf{M} \mathbf{S}, \quad (4)$$

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} .

In accordance with the details presented in [8] and [6], SVD produces unique estimates of the motion and shape matrices,

$$\begin{aligned} \mathbf{Z} &= \mathbf{U}^* \Sigma^* \mathbf{V}^{*T} = \hat{\mathbf{M}} \hat{\mathbf{S}} \\ \hat{\mathbf{M}} &= \mathbf{U}^* (\Sigma^*)^{1/2} \\ \hat{\mathbf{S}} &= (\Sigma^*)^{1/2} \mathbf{V}^{*T} \end{aligned} \quad (5)$$

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)

$$\hat{\mathbf{M}} \hat{\mathbf{S}} = (\hat{\mathbf{M}} \mathbf{A}^{-1}) (\mathbf{A} \hat{\mathbf{S}}), \quad (6)$$

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

$$\begin{aligned} \mathbf{M} &= \hat{\mathbf{M}} \mathbf{A}^{-1} \\ \mathbf{S} &= \mathbf{A} \hat{\mathbf{S}} \end{aligned} \quad (7)$$

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 (7) and \mathbf{C} using (3). 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. 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

3.3 Force/Torque Sensor Calibration

The shape from motion calibration technique is illustrated more completely in [8], 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 and summarize the results of the 6-DOF case from [8].

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 :

$$z_{ij} = F s_{1j} \cos \theta_i \sin \psi_i + F s_{2j} \sin \theta_i \sin \psi_i + F s_{3j} \cos \psi_i \quad (8)$$

which can be rewritten for multiple strain gauges in the form of equation (2) as

$$[z_{i1} \ z_{i2} \ \dots \ z_{i8}] = [\cos \theta_i \sin \psi_i \ \sin \theta_i \sin \psi_i \ \cos \psi_i]^* \quad (9)$$

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

The constant magnitude of the force has been factored out, leaving only $[\cos \theta \sin \psi \ \sin \theta \sin \psi \ \cos \psi]$ in each row of the motion matrix. i is the index over all samples. From (9) we can see that the rank is 3 (three independent columns) and from (9) we also get the motion constraint. The motion matrix consists of rows of unit vectors ($\mathbf{m}_i^T = [\cos \theta \sin \psi \ \sin \theta \sin \psi \ \cos \psi]$) so the constraint sets the sum of the squares of the columns of the motion matrix to 1. If we impose this constraint on $\hat{\mathbf{M}}$, the result of the SVD, the true motion matrix, \mathbf{M} , will automatically result as will the complementary shape matrix, \mathbf{S} .

To enforce this constraint we rely on equation (7) 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

$$(\mathbf{m}_i^T \mathbf{a}_1)^2 + (\mathbf{m}_i^T \mathbf{a}_2)^2 + (\mathbf{m}_i^T \mathbf{a}_3)^2 = 1 \quad (10)$$

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 (7) and the calibration matrix, \mathbf{C} , using (3).

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.

Although it takes some insight to apply computer vision principles to the sensor calibration problem, algorithmically, the two applications are the same to this point. However, we have only considered pure forces. Torque is linearly dependent on force components, so the addition of torque into the equation does not increase the rank of the matrix. Hence, the decomposition must be augmented to extract torques.

The “squashing matrix” was introduced in [8] to demonstrate how the additional degrees of freedom introduced by torques become embedded inside the recovered shape matrix, \mathbf{S} :

$$\mathbf{S} = \begin{bmatrix} \mathbf{I} & \mathbf{X}^T \end{bmatrix} \tilde{\mathbf{S}} \quad (11)$$

where $\tilde{\mathbf{S}}$ is the true shape matrix, the squashing matrix is represented by $\begin{bmatrix} \mathbf{I} & \mathbf{X}^T \end{bmatrix}$, and \mathbf{X}^T is the cross-product matrix applied to the moment arm of the load. While the details appear in [8], the true shape is recovered by repeated application of the procedure to the sensor with linearly independent moment arms for each trial. The different trials are combined and the solution found using

$$\tilde{\mathbf{S}} = \begin{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{X}_1 \end{bmatrix} \\ \begin{bmatrix} \mathbf{I} & \mathbf{X}_2 \end{bmatrix} \\ \begin{bmatrix} \mathbf{I} & \mathbf{X}_3 \end{bmatrix} \end{bmatrix}^+ \begin{bmatrix} \mathbf{S}_1 \\ \mathbf{S}_2 \\ \mathbf{S}_3 \end{bmatrix}. \quad (12)$$

Two trials are required for the planar case and three trials, as shown in (12), are required for the full 6-DoF solution.

To briefly summarize the experimental results presented in [8], two sensors were calibrated: a planar 2-D sensor and a 6-DOF Lord force/torque sensor. The Shape from Motion technique resulted in a ten-times reduction in the time necessary to fully calibrate the sensor with slightly greater accuracy. The 6-DoF case was not quite as dramatic, taking half as long with Shape from Motion but still producing a slightly more accurate result. Table 1 tabulates the results of attaching a known load to the 6-DoF 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 magnitude of the load. “Acc” is the standard deviation of the smoothed force measurement, which provides a measure of

shape accuracy (zero is ideal). “Time” is the time required to execute each procedure.

Table 1: LORD F/T SENSOR COMPARISON

load (N)	Least Squares				Shape from Motion			
	mean (N)	error (%)	acc	time (min)	mean (N)	error (%)	acc	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	

4 Including Sensor Bias

The extension of shape from motion to include sensor bias is described completely in [9]. The resulting equations are similar to Tomasi and Kanade’s original computer vision derivation outlined in [5].

For the single strain gauge as in equation (8), adding a non-zero offset becomes:

$$z_{ij} = F s_{1j} \cos \theta_i \sin \psi_i + F s_{2j} \sin \theta_i \sin \psi_i + F s_{3j} \cos \psi_i + z_{0j} \quad (13)$$

which can be rewritten in accordance with equation (9):

$$\begin{bmatrix} z_{i1} & z_{i2} & \dots & z_{i8} \end{bmatrix} = \begin{bmatrix} \cos \theta_i \sin \psi_i & \sin \theta_i \sin \psi_i & \cos \psi_i & 1 \end{bmatrix} * \quad (14)$$

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

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.

Armed with the new proper rank of \mathbf{Z} , the shape from motion procedure progresses in similar fashion as reported in Section 3.3 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 instead of 3×3 .

The short story is that two decoupled constraints result:

$$\begin{aligned} (\mathbf{m}_i^T \mathbf{a}_1)^2 + (\mathbf{m}_i^T \mathbf{a}_2)^2 + (\mathbf{m}_i^T \mathbf{a}_3)^2 &= 1 \\ \mathbf{m}_i^T \mathbf{a}_4 &= 1 \end{aligned} \quad (15)$$

the latter of which is just the description of a plane while the former describes a cylinder. Fitting data to a cylinder is a

tricky nonlinear problem. The solution proposed by Tomasi is to approximate it as an ellipse and refine the solution numerically, which works well in the calibration case and is described in [9].

Concatenating the solution from the plane constraint to the solution from the cylindrical constraint yields a 4×4 matrix, \mathbf{A}^{-1} . Using equation (7), the augmented motion and shape matrices are reconstructed and stripping off the offsets indicated in (14) yields the calibration matrix.

5 Collaborative Calibration

The Shape from Motion Calibration technique outlined thus far is clearly a different approach to calibration because it eliminates the need for a reference for each and every measurement. However, we would like to improve it in two areas: greater robustness to nonlinearities and greater autonomy.

Nonlinearities can be damaging because the entire calibration procedure is performed at a single force magnitude. If that magnitude is not chosen correctly the linearized result may be suboptimal. The autonomy, on the other hand, is already dramatically better than the least squares technique because the requirements for externally-applied loads is reduced. Still, a constant mass must be attached externally.

This section details how systems with multiple sensors, such as multi-robot systems with wrist sensors or dextrous hands with fingertip sensors, can collaboratively calibrate their sensors, eliminating the need for an external mass and allowing an infinite range of force magnitudes to contribute to the linearization.

5.1 An Alternative Constraint

The constant magnitude adds a constraint to the solution that is required to make the problem solvable. It manifests itself in equation (10), where the left-hand side is always equal to a constant (assumed to be 1 unit). To solve (10) it is not necessary that the right-hand side be constant, only that it be constrained. Given two force sensors pressing against each other, the two magnitudes are constrained to be equal to each other. This equality constraint can be used to replace the constant equality constraint of equation (10):

$$(\mathbf{m}_{1i}^T \mathbf{a}_{11})^2 + (\mathbf{m}_{1i}^T \mathbf{a}_{12})^2 + (\mathbf{m}_{1i}^T \mathbf{a}_{13})^2 = (\mathbf{m}_{2i}^T \mathbf{a}_{21})^2 + (\mathbf{m}_{2i}^T \mathbf{a}_{22})^2 + (\mathbf{m}_{2i}^T \mathbf{a}_{23})^2 \quad (16)$$

The vectors \mathbf{m}_{1i} and \mathbf{m}_{2i} result from two separate but concurrent executions of the shape from motion procedure outlined in Section 3.2, resulting in:

$$\begin{aligned} \mathbf{M}_1 &= \hat{\mathbf{M}}_1 \mathbf{A}_1^{-1} \\ \mathbf{S}_1 &= \mathbf{A}_1 \hat{\mathbf{S}}_1 \\ \mathbf{M}_2 &= \hat{\mathbf{M}}_2 \mathbf{A}_2^{-1} \\ \mathbf{S}_2 &= \mathbf{A}_2 \hat{\mathbf{S}}_2 \end{aligned} \quad (17)$$

The bilinear problem of equation (16) can be solved iteratively by guessing a solution to the right-hand side, holding it fixed, then solving for the left-hand side. Next, using the solution just found, hold the left-hand side constant and solve for the right-hand side, and so on. Because zeroes on both sides is the only exact solution, one must take care to re-normalize the results periodically to prevent relaxing to the trivial answer.

5.2 Experimental Results

To test this new result in shape from motion calibration, we physically attached two 6-DoF force/torque sensors together with a universal joint. The universal joint is not absolutely necessary. The only requirement is that the moment arms for each sensor stay constant. This can be achieved by exerting radial forces on a spherical surface, such as the tip of a robotic finger. We used the U-joint for convenience.

Holding the two sensors, one in each hand, the experimenter merely pushed and pulled on the two sensors manually to exert a wide range of force magnitudes that sufficiently excited the sensing spaces of both sensors. (Sufficiently exciting the sensors is an issue of design of experiments and will not be discussed here. As a simple example, if the applied vector is constrained to a plane for a 3-D sensor, the result will be poor. Refer to any standard text on the design of experiments, such as [1].)

The raw output data from both sensors was sampled synchronously at 10 Hz. After the processing, the extracted motion matrices are plotted in Figure 1 and Figure 2. Obviously, the applied force was highly random but corresponding segments of the motion through space are easy to pick out. After applying two known forces to scale and orient the force sensor reference frame, analysis of the accuracy of the calibration matrix as performed in [8] produced results similar to Table 1, as expected.

The benefits of this extension to the basic shape from motion technique are twofold. First, it allows fully autonomous calibration of sensors when an external reference frame is not pre-specified, hence there is no need for applying *any* known forces. (These cases will be explored in the next section.) Second, and more important from a theoretical standpoint, it allows calibration across a




Figure 1: Motion of the force vector around one sensor origin during collaborative calibration.




Figure 2: Motion of the force vector around the other sensor origin during collaborative calibration. Note the inverted, mirror image.

wide range of force magnitudes. This is an important aspect of the design of experiments because it allows the best linear fit for sensors that have slight nonlinearities.

6 Primordial Learning

In several instances we have referred to hypothetical cases in which we may not need to establish a known reference to make use of the calibration information. One such area of study is skill acquisition. Humans learn to

calibrate their sensors through practicing skill-based tasks, yet they generally don't apply known references (with the obvious exception of the gravity vector).

Because skill acquisition seems well off the subject of sensor calibration we will not pursue details here, but we have applied a variant of Shape from Motion Calibration to learning various skills for a mobile robot and a PUMA manipulator. The mobile robot learned an effective wandering behavior that escaped stalls and jams [7] and the PUMA learned several sensorimotor primitives for guarded moves and force-based accommodation [10]. In these applications the intrinsic structure of the data resulted from human teleoperation and the "skills" were extracted by SVD.

7 Discussion

In this paper we first summarized a novel approach to force sensor calibration that eliminates the need to measure all applied forces. This significantly reduces the burden of data collection and results in a slightly more accurate procedure by eliminating the input measurements as a source of noise. To our knowledge, calibration has never before been posed as an output-only problem, but has always been posed as an input/output problem.

We also presented an important new extension to the shape from motion paradigm: collaborative calibration. This allows two sensors to mutually calibrate without the need for an attached mass. The two sensors can just press against each other randomly as long as the moment arm to each sensor origin stays constant. Although we used a universal joint to connect the sensors physically, that is unnecessary in general. In fact, we plan to apply the technique to fingertip tactile sensors on a Utah/MIT hand in the near future. Hemispherical domes on the tips will allow contact anywhere on the surface while maintaining a zero moment arm. (Although we must deal appropriately with tangential forces of friction.)

Another benefit of collaborative calibration is that it allows the applied force magnitude to vary over a wide range. This is important for the proper design of experiments in that it better accounts for slight nonlinearities in the sensor and varying operating points of the task.

Finally, we briefly mentioned "shape from motion primordial learning" as an approach to skill acquisition. Although no details were presented, we have successfully extracted simple sensorimotor primitives from human demonstration on both a mobile robot and a PUMA manipulator and were able to identify those primitives in subsequent demonstrations.

Of course, shape from motion primordial learning has some important shortcomings. Because it relies on a linear representation, it can not be used to calibrate sensors of

unknown nonlinear behavior. The same applies to skill acquisition. Careful segmentation must be performed to extract piecewise linear primitives of complex skills and many classes of skills will be entirely unlearnable. For example, the manipulation primitives learned by our PUMA are equivalent to various modes of damping force control. Still, it is a convenient representation that provides not just acquisition, but also recognition and transformation of skills (or primitives), which are capabilities central to the goal of intelligent programming by human demonstration.

8 References

- [1] Diamond, W., *Practical Experiment Designs for Engineers and Scientists*, Lifetime Learning Pub., 1981.
- [2] Forsythe, G.E., M. Malcolm, and C.B. Moler, *Computer Methods for Mathematical Computations*, Prentice Hall, Englewood Cliffs, NJ, 1977.
- [3] 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.
- [4] Strang, G., *Linear Algebra and Its Applications*, third edition, Harcourt Brace Jovanovich Publishers, San Diego, CA, 1988, appendix A.
- [5] 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.
- [6] Tomasi, C., "Shape and Motion from Image Streams: a Factorization Method," Ph.D. Thesis, CMU-CS-91-172, Comp. Sci., Carnegie Mellon University, Pittsburgh, PA, Sept. 1991.
- [7] Voyles, R.M., J.D. Morrow, and P.K. Khosla, "Shape from Motion Decomposition as a Learning Approach for Autonomous Agents," in *Proceedings of 1995 IEEE Conf. on Systems, Man, and Cybernetics*, Vancouver, B.C., Oct. 1995.
- [8] 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.
- [9] Voyles, R.M., J.D. Morrow, and P.K. Khosla, "Including Bias in Shape from Motion Calibration and Sensor Fusion," to appear in *Proceedings of IEEE Multisensor Fusion and Integration Conf.*, Washington, DC, Dec 1996.
- [10] Voyles, R.M., J.D. Morrow, and P.K. Khosla, "Gesture-Based Programming, Part 2: Primordial Learning," to appear in *Smart Engineering Systems Design: Neural Networks, Fuzzy Logic, and Evolutionary Programming*, Dagli, Akay, Chen, Fernandez, and Ghosh, eds., ASME Press, 1996.