

Propagating Covariance in Computer Vision

Robert M. Haralick
Intelligent Systems Laboratory
Department of Electrical Engineering
University of Washington
Seattle, WA 98195

Abstract

This paper describes how to propagate approximately additive random perturbations through any kind of vision algorithm step in which the appropriate random perturbation model for the estimated quantity produced by the vision step is also an additive random perturbation. We assume that the vision algorithm step can be modeled as a calculation (linear or non-linear) that produces an estimate that minimizes an implicit scalar function of the input quantity and the calculated estimate. The only assumption is that the scalar function have finite second partial derivatives and that the random perturbations are small enough so that the relationship between the scalar function evaluated at the ideal but unknown input and output quantities and the observed input quantity and perturbed output quantity can be approximated sufficiently well by a first order Taylor series expansion.

The paper finally discusses the issues of verifying that the derived statistical behavior agrees with the experimentally observed statistical behavior.

1 Introduction

Each real computer vision problem begins with one or more noisy images and has many algorithmic steps. Development of the best algorithm requires understanding how the uncertainty due to the random perturbation affecting the input image(s) propagates through the different algorithmic steps and results in a perturbation on whatever quantities are finally computed. Perhaps a more accurate statement would be that the quantities finally computed must really be considered to be estimated quantities.

Once we have the perspective that what we compute are estimates, then it becomes clear that even though the different ways of estimating the same quantity typically yield the same result if the input quantities are not affected

by a random perturbation, it is certainly not the case that the different ways of estimating the same quantities yield an estimate with the same distribution when the input is perturbed by a random perturbation. It is clearly the case that the distribution of the estimate depends on the distribution of the input random perturbation and the method or type of estimate.

With this in mind, it is then important to understand how to propagate a random perturbation through any algorithm step in a vision problem. The difficulty is that the steps are not necessarily linear computations, the random perturbations are not necessarily additive, and the appropriate kinds of perturbations change from algorithm step to algorithm step. Nevertheless, there are many computer vision and image analysis algorithm steps in which the appropriate kind of random perturbation is additive or approximately additive. And for these kinds of steps one basic measure of the size of the random perturbation is given by the covariance matrix of the estimate.

In this paper, we describe how to propagate the covariance matrix of an input random perturbation through any kind of a calculation (linear or nonlinear) that extremizes an implicit scalar function, between the input quantity and the calculated estimate. The only assumption is that the scalar function have finite second partial derivatives and that the random perturbations are small enough so that the relationship between the scalar function evaluated at the ideal but unknown input and out quantities and the observed input quantity and perturbed output quantity can be approximated sufficiently well by a first order Taylor series expansion.

As a related case, the given propagation relationships also show how to propagate the covariance of the coefficients of a function for which we wish to find a zero to the covariance of any zero we can find.

The analysis techniques we employ are well-known in statistical and engineering communities, although they are perhaps not so well known in the computer vision community.

The paper concludes with a discussion of how to validate that the software which we use to accomplish the calculation we desire actually works. We argue that this validation can be done by comparing the predicted statistical behavior with the experimentally observed statistical behavior in a set of controlled experiments.

2 The Abstract Model

The abstract model has three kinds of objects. The first kind of object relates to the measurable quantities. There is the unobserved $N \times 1$ vector X of the ideal unperturbed measurable quantities. We assume that each component of X is some real number. Added to this unobserved ideal unperturbed vector is an $N \times 1$ unobserved random vector ΔX of noise. The observed quantity is the randomly perturbed vector $X + \Delta X$.

The second kind of object relates to the unknown parameters. There is the unobserved $K \times 1$ vector Θ . We assume that each component of Θ is some

real number. Added to this ideal unperturbed vector is a $K \times 1$ unobserved vector $\Delta\Theta$ that is the random perturbation on Θ induced by the random perturbation ΔX on X . The calculated quantity is the randomly perturbed parameter vector $\hat{\Theta} = \Theta + \Delta\Theta$.

The third kind of object is a continuous scalar valued function F which relates the vectors X and Θ and which relates the vectors $X + \Delta X$ and $\Theta + \Delta\Theta$. The function F has finite first and second partial derivatives with respect to each component of Θ and X , including all second mixed partial derivatives taken with respect to a component of Θ and with respect to a component of X .

The basic problem is: given $\hat{X} = X + \Delta X$, determine a $\hat{\Theta} = \Theta + \Delta\Theta$ to minimize $F(\hat{X}, \hat{\Theta})$ given the fact that Θ minimizes $F(X, \Theta)$.

Of course, if $\hat{\Theta}$ is computed by an explicit function h , so that $\hat{\Theta} = h(\hat{X})$, the function F is just given by $f(X, \Theta) = (\Theta - h(X))'(\Theta - h(X))$.

3 Example Computer Vision Problems

There is a rich variety of computer vision problems which fit the form of the abstract model. In this section we outline a few of them, specifically: curve fitting, coordinated curve fitting, local feature extraction, exterior orientation, and relative orientation. Other kinds of calculations in computer vision such as calculation of curvature, invariants, vanishing points, or points at which two or more curves intersect, or problems such as motion recovery are all examples of problems which can be put in the abstract form as given above.

3.1 Curve Fitting

In the general curve fitting scenario, there is the unknown free parameter vector, ψ , of the curve and the set of unknown ideal points on the curve $\{x_1, \dots, x_N\}$. Each of the ideal points is then perturbed. If Δx_n is the random noise perturbation of the n^{th} point, then the observed point n^{th} point is $\hat{x}_n = x_n + \Delta x_n$. The form of the curve is given by a known function f which relates a point on the curve to the parameters of the curve. That is, for each ideal point x_n we have $f(x_n, \psi) = 0$. We also assume that the parameters of the curve satisfy its own constraint $h(\psi) = 0$. The curve fitting problem is then to find an estimate $\hat{\psi}$ to minimize $\sum_{n=1}^N f^2(\hat{x}_n, \hat{\psi})$ subject to $h(\hat{\psi}) = 0$.

To put this problem in the form of the abstract problem we let

$$\begin{aligned} X &= (x_1, \dots, x_N) \\ \hat{X} &= (x_1 + \Delta x_1, \dots, x_N + \Delta x_N) \\ \Theta &= (\psi, \lambda), \quad \hat{\Theta} = (\hat{\psi}, \lambda) \\ F(X, \Theta) &= \sum_{n=1}^N f^2(x_n, \psi) + \lambda h(\psi) \end{aligned}$$

Then the curve fitting problem is to find $\hat{\Theta}$ to minimize $F(\hat{X}, \hat{\Theta})$ where $F(X, \Theta) = 0$.

3.2 Coordinated Curve Fitting

In the coordinated curve fitting problem, multiple curves have to be fit on independent data, but the fitted curves have to satisfy some joint constraint. We illustrate the discussion in this section with a coordinated fitting of two curves and a constraint that the two curves must have some common point at which they are tangent.

Let (x_1, \dots, x_I) be the ideal points which are associated with the first curve whose parameters are ψ_1 and whose constraint is $h_1(\psi_1) = 0$. Each point x_i satisfies $f_1(x_i, \psi_1) = 0$, $i = 1, \dots, I$.

Likewise, let (y_1, \dots, y_J) be the ideal points which are associated with the second curve whose parameters are ψ_2 and whose constraint is $h_2(\psi_2) = 0$. Each point y_j satisfies $f_2(y_j, \psi_2) = 0$, $j = 1, \dots, J$.

The coordinated constraint is that for some unknown z ,

$$\begin{aligned} f_1(z, \psi_1) &= 0 \\ f_2(z, \psi_2) &= 0 \\ \frac{\partial f_1}{\partial z}(z, \psi_1) &= \frac{\partial f_2}{\partial z}(z, \psi_2) \end{aligned}$$

The observed points \hat{x}_i and \hat{y}_j are related to the corresponding ideal points by

$$\begin{aligned} \hat{x}_i &= x_i + \Delta x_i \\ \hat{y}_j &= y_j + \Delta y_j \end{aligned}$$

To put this problem in the framework of the abstract model, we take

$$\begin{aligned} \hat{X} &= (\hat{x}_1, \dots, \hat{x}_I, \hat{y}_1, \dots, \hat{y}_J) \\ \hat{\Theta} &= (\hat{\psi}_1, \hat{\psi}_2, \hat{z}, \lambda_1, \lambda_2, \lambda_3, \lambda_4, \lambda_5) \end{aligned}$$

and define

$$\begin{aligned} F(\hat{X}, \hat{\Theta}) &= \sum_{i=1}^I f_1^2(\hat{x}_i, \hat{\psi}_1) + \sum_{j=1}^J f_2^2(\hat{y}_j, \hat{\psi}_2) + \lambda_1 h_1(\hat{\psi}_1) + \lambda_2 h_2(\hat{\psi}_2) \\ &+ \lambda_3 f_1(\hat{z}, \hat{\psi}_1) + \lambda_4 f_2(\hat{z}, \hat{\psi}_2) + \lambda_5 \left[\frac{\partial f_1}{\partial z}(\hat{z}, \hat{\psi}_1) - \frac{\partial f_2}{\partial z}(\hat{z}, \hat{\psi}_2) \right] \end{aligned}$$

The coordinated curve fitting problem is then to determine a $\hat{\Theta}$ to minimize $F(\hat{X}, \hat{\Theta})$, where the perturbed $\hat{\Theta}$ is considered are related to the ideal Θ by $\hat{\Theta} = \Theta + \Delta\Theta$.

3.3 Local Feature Extraction

There are a variety of local features that can be extracted from an image. Examples include edges, corners, ridges, valleys, flats, saddles, slopes, hillsides, saddle hillsides, etc. Each local feature involves the calculation of some quantities assuming that the neighborhood has the feature and then a detection is

performed based on the calculated quantities. For example, in the simple gradient edge feature, the quantity calculated is the gradient magnitude and the edge feature is detected if the calculated gradient magnitude is high enough. Here we concentrate on the calculation of the quantities associated with the feature and not the detection of the feature itself.

To put this problem in the setting of the abstract problem, we let Θ be the vector of unknown free parameters of the feature and X be the unobserved neighborhood array of noiseless brightness values. We let \hat{X} be the perturbed observed neighborhood array of brightness values, $\hat{X} = X + \Delta X$, and $\hat{\Theta}$ be the calculation of the required quantities from the perturbed brightness values \hat{X} . The form of the feature is given by the known function f which satisfies that $f(X, \Theta) = 0$. The feature extraction problem is then to find the estimate $\hat{\Theta}$ to minimize $F(\hat{X}, \Theta) = f^2(\hat{X}, \hat{\Theta})$.

3.4 Exterior Orientation

In the exterior orientation problem, there is a known 3D object model having points $(x_n, y_n, z_n), n = 1, \dots, N$. The unobserved noiseless perspective projection of the point (x_n, y_n, z_n) is given by (u_n, v_n) . The relationship between a 3D model point and its corresponding perspective projection is given by a rotation and translation of the object model point, to put it in the reference frame of the camera, followed by a perspective projection. So if ψ represents the triple of tilt angle, pan angle, and swing angle of the rotation, t represents the x-y-z-translation vector, and k represents the camera constant (the focal length of the camera lens), we can write:

$$\begin{aligned}(u_n, v_n)' &= \frac{k}{r_n}(p_n, q_n)' \text{ where} \\ (p_n, q_n, r_n)' &= R(\psi)(x_n, y_n, z_n)' + t\end{aligned}$$

and where $R(\psi)$ is the 3×3 rotation matrix corresponding to the rotation angle vector ψ .

The function to be minimized can then be written as:

$$\begin{aligned}f_n(u_n, v_n, \psi, t) &= f(u_n, v_n, x_n, y_n, z_n, \psi, t) \text{ where} \\ f(u_n, v_n, x_n, y_n, z_n, \psi, t) &= \left[u_n - k \frac{(1, 0, 0)(R(\psi)(x_n, y_n, z_n)' + t)}{(0, 0, 1)(R(\psi)(x_n, y_n, z_n)' + t)} \right]^2 \\ &+ \left[v_n - k \frac{(0, 1, 0)(R(\psi)(x_n, y_n, z_n)' + t)}{(0, 0, 1)(R(\psi)(x_n, y_n, z_n)' + t)} \right]^2\end{aligned}$$

To put this problem in the form of the abstract description we take

$$\begin{aligned}X &= (u_1, v_1, \dots, u_n, v_n) \\ \hat{X} &= (\hat{u}_1, \hat{v}_1, \dots, \hat{u}_n, \hat{v}_n) \\ \Theta &= (\psi, t)\end{aligned}$$

$$\hat{\Theta} = (\hat{\psi}, \hat{t})$$

and define

$$F(\hat{X}, \hat{\psi}) = \sum_{n=1}^N f_n^2(\hat{u}_n, \hat{v}_n, \hat{\Theta})$$

The exterior orientation problem is then to find a $\hat{\Theta}$ to minimize $F(\hat{X}, \hat{\Theta})$, given that $F(X, \Theta) = 0$. And because F is non-negative it must be that Θ minimizes $F(X, \Theta)$.

3.5 Relative Orientation

The relative orientation problem can be put into the form of the abstract problem in a similar way to the exterior orientation problem. We let the perspective projection of the n^{th} point on the left image be (u_{nL}, v_{nL}) and the perspective projection of the n^{th} point on the right image be (u_{nR}, v_{nR}) . Then we can write that

$$\begin{aligned} (u_{nL}, v_{nL})' &= \frac{k}{z_n} (x_n, y_n)' \text{ and that} \\ (u_{nR}, v_{nR})' &= \frac{k}{r_n} (p_n, q_n) \end{aligned}$$

where (p_n, q_n, r_n) is the rotated and translated model point as given in the description of the exterior orientation problem.

The observed perspective projection of the n^{th} model point is noisy and represented as $(\hat{u}_n, \hat{v}_n) = (u_n + \Delta u_n, v_n + \Delta v_n)$. Then taking

$$\begin{aligned} X &= (u_{1L}, v_{1L}, u_{1R}, v_{1R}, \dots, u_{NL}, v_{NL}, u_{NR}, v_{NR}) \\ \hat{X} &= (\hat{u}_{1L}, \hat{v}_{1L}, \hat{u}_{1R}, \hat{v}_{1R}, \dots, \hat{u}_{NL}, \hat{v}_{NL}, \hat{u}_{NR}, \hat{v}_{NR}) \\ \Theta &= (x_1, y_1, z_1, \dots, x_N, y_N, z_N, \psi, t) \\ \hat{\Theta} &= (\hat{x}_1, \hat{y}_1, \hat{z}_1, \dots, \hat{x}_N, \hat{y}_N, \hat{z}_N, \hat{\psi}, \hat{t}) \end{aligned}$$

the relative orientation problem is to find $\hat{\Theta}$ to minimize

$$F(\hat{X}, \hat{\Theta}) = \sum_{n=1}^N f(u_{nR}, v_{nR}, x_n, y_n, z_n, \psi, t) + f(u_{nL}, v_{nL}, x_n, y_n, z_n, 0, 0)$$

4 Zero Finding

Zero finding such as finding the zero of a polynomial in one or more variables occurs in a number of vision problems. Two examples are the three point perspective resection problem and some of the techniques for motion recovery. The zero finding problem is precisely in the form required for computing the covariance matrix $\Sigma_{\Delta\Theta}$ as described in the solution section. Let X be the ideal input vector and \hat{X} be the observed perturbed input vector. Let Θ be a $K \times 1$ vector zeroing the $K \times 1$ function $g(X, \Theta)$; that is, $g(X, \Theta) = 0$. Finally, let $\hat{\Theta}$ be the computed vector zeroing $g(\hat{X}, \hat{\Theta})$; that is, $g(\hat{X}, \hat{\Theta}) = 0$.

5 Solution

For the purpose of covariance determination of the computed $\hat{\Theta} = \Theta + \Delta\Theta$, the technique used to solve the extremization problem is not important, provided that there are no singularities or near singularities in the numerical computation procedure itself.

To understand how the random perturbation ΔX acting on the measurable vector X propagates to the random perturbation $\Delta\Theta$ on the parameter vector Θ , we can take partial derivatives of F with respect to each of the K components of Θ forming the gradient g of f . The gradient g is a $K \times 1$ vector function.

$$g(X, \Theta) = \frac{\partial F}{\partial \Theta}(X, \Theta)$$

The solution $\hat{\Theta} = \Theta + \Delta\Theta$ extremizing $F(X + \Delta X, \hat{\Theta})$, however it is calculated, must be a zero of $g(X + \Delta X, \hat{\Theta})$. Now taking a Taylor series expansion of g around (X, Θ) we obtain to a first order approximation:

$$g^{K \times 1}(X + \Delta X, \Theta + \Delta\Theta) = g^{K \times 1}(X, \Theta) + \frac{\partial g^{K \times N}}{\partial X}(X, \Theta)\Delta X^{N \times 1} + \frac{\partial g^{K \times K}}{\partial \Theta}(X, \Theta)\Delta\Theta^{K \times 1}$$

But since $\Theta + \Delta\Theta$ extremizes $F(X + \Delta X, \Theta + \Delta\Theta)$, $g(X + \Delta X, \Theta + \Delta\Theta) = 0$. Also, since Θ extremizes $F(X, \Theta)$, $g(X, \Theta) = 0$. Thus to a first order approximation,

$$0 = \frac{\partial g}{\partial X}(X, \Theta)\Delta X + \frac{\partial g}{\partial \Theta}\Delta\Theta$$

Since the relative extremum of F is a relative minimum, the $K \times K$ matrix

$$\frac{\partial g}{\partial X}(X, \Theta) = \frac{\partial^2 f}{\partial^2 \Theta}(X, \Theta)$$

must be positive definite for all (X, Θ) . This implies that

$$\frac{\partial g}{\partial X}(X, \Theta) = \frac{\partial^2 f}{\partial^2 \Theta}(X, \Theta)$$

is non-singular. Hence

$$\left(\frac{\partial g}{\partial \Theta}\right)^{-1}$$

exists and we can write:

$$\Delta\Theta = -\left\{\frac{\partial g}{\partial \Theta}(X, \Theta)\right\}^{-1} \frac{\partial g}{\partial X}(X, \Theta)\Delta X$$

This relation states how the random perturbation ΔX on X propagates to the random perturbation $\Delta\Theta$ on Θ . If the expected value of ΔX , $E[\Delta X]$, is zero, then from this relation we see the $E[\Delta\Theta]$ will also be zero, to a first order approximation.

This relation also permits us to calculate the covariance of the random perturbation $\Delta\Theta$.

$$\begin{aligned}
\Sigma_{\Delta\theta} &= E[\Delta\theta\Delta\theta'] \\
&= E[-(\frac{\partial g}{\partial\theta})^{-1}\frac{\partial g}{\partial X}\Delta X(-(\frac{\partial g}{\partial\theta})^{-1}\frac{\partial g}{\partial X}\Delta X)'] \\
&= (\frac{\partial g}{\partial\theta})^{-1}\frac{\partial g}{\partial X}E[\Delta X\Delta X'](\frac{\partial g}{\partial X})'(\frac{\partial g}{\partial\theta})^{-1} \\
&= (\frac{\partial g}{\partial\theta})^{-1}\frac{\partial g}{\partial X}\Sigma_{\Delta X}(\frac{\partial g}{\partial X})'(\frac{\partial g}{\partial\theta})^{-1}
\end{aligned}$$

Thus to the extent that the first order approximation is good, (i.e. $E[\Delta\theta] = 0$), then

$$\Sigma_{\hat{\theta}} = \Sigma_{\Delta\theta}$$

The way in which we have derived the covariance matrix for $\Delta\theta$ based on the covariance matrix for ΔX requires that the matrices

$$\frac{\partial g}{\partial\theta}(X, \theta) \text{ and } \frac{\partial g}{\partial X}(X, \theta)$$

be known. But X and θ are not observed. $X + \Delta X$ is observed and by some means $\theta + \Delta\theta$ is then calculated. So if we want to determine an estimate $\hat{\Sigma}_{\hat{\theta}}$ for the covariance matrix $\Sigma_{\hat{\theta}}$, we can proceed by expanding $g(X, \theta)$ around $g(X + \Delta X, \theta + \Delta\theta)$.

$$g(X, \theta) = g(X + \Delta X, \theta + \Delta\theta) - \frac{\partial g}{\partial X}(X + \Delta X, \theta + \Delta\theta)\Delta X - \frac{\partial g}{\partial\theta}(X + \Delta X, \theta + \Delta\theta)\Delta\theta$$

Here we find in a similar manner,

$$\Delta\theta = -(\frac{\partial g}{\partial\theta}(X + \Delta X, \theta + \Delta\theta))^{-1}\frac{\partial g}{\partial X}(X + \Delta X, \theta + \Delta\theta)\Delta\theta$$

This motivates the estimator $\hat{\Sigma}_{\Delta\theta}$ for $\Sigma_{\Delta\theta}$ defined by

$$\begin{aligned}
\hat{\Sigma}_{\Delta\theta} &= (\frac{\partial g}{\partial\theta}(X + \Delta X, \theta + \Delta\theta))^{-1}\frac{\partial g}{\partial X}(X + \Delta X, \theta + \Delta\theta)\Sigma_{\Delta X} \\
&\quad \times \frac{\partial g}{\partial X}(X + \Delta X, \theta + \Delta\theta)'(\frac{\partial g}{\partial\theta}(X + \Delta X, \theta + \Delta\theta))^{-1}
\end{aligned}$$

So to the extent that the first order approximation is good, $\hat{\Sigma}_{\hat{\theta}} = \hat{\Sigma}_{\Delta\theta}$.

The relation giving the estimate $\hat{\Sigma}_{\hat{\theta}}$ in terms of the computable

$$\frac{\partial g}{\partial\theta}(X + \Delta X, \hat{\theta}) \text{ and } \frac{\partial g}{\partial X}(X + \Delta X, \hat{\theta})$$

means that a covariance matrix for the computed $\hat{\theta} = \theta + \Delta\theta$ can also be calculated at the same time that the estimate $\hat{\theta}$ of θ is calculated.

6 Validation

Software for performing the optimization required to compute the estimate $\hat{\Theta}$ is often complicated and it is easy for there to be errors that are not immediately observable (like optimization software that produces correct answers on a few known examples but fails in a significant fraction of more difficult cases). So a key issue in testing is whether the software is producing the right answers in the hard cases and whether the statistical properties of the answers it produces are similar to the statistical properties we expect. Both these kinds of issues can be handled by doing a significant number of experiments whose results can be used in a hypothesis test that everything is as it should be.

Consider what happens in a hypothesis test: a significance level, α , is selected. When the test is run, a test statistic, say $\hat{\phi}$, is computed. The test statistic is typically designed so that in the case that the hypothesis is true, the test statistic will tend to have its values distributed around zero, in accordance with a known distribution. If the test statistic has a value say higher than a given ϕ_0 , we reject the hypothesis. Else we do not reject, in effect tentatively accepting the hypothesis. The value of ϕ_0 is chosen so that the probability that we reject the hypothesis, given that the hypothesis is true is less than the significance level α .

In the case that we do not know the distribution of the test statistic, but we do know its mean and variance, we can still do the hypothesis test by independently generating the test statistic a number of times. So suppose that we repeat the test M times. The m th repetition generates a test statistic ϕ_m . So after M repetitions we have observed the M independent and identically distributed statistics ϕ_1, \dots, ϕ_M . Let the mean and variance of the each of the test statistics be denoted by μ and σ^2 , respectively. Calculate the mean test statistic $\hat{\phi} = \frac{1}{M} \sum_{m=1}^M \phi_m$. Let t_0 be chosen so that the probability that a Normally distributed random variable will have a value greater than t_0 is less than the significance level α . With M large enough so that the mean test statistic approaches being a Normal, under the hypothesis, the probability that $\hat{\phi}$ will be less than $\frac{\sigma}{\sqrt{M}} t_0 + \mu$ is equal to α .

To test the hypothesis that $\hat{\Theta}$ is unbiased and that the covariance matrix $\Sigma_{\Delta\Theta}$ for the $\hat{\Theta}$ is calculated

$$\Sigma_{\hat{\Theta}} = \left(\frac{\partial g}{\partial \Theta}\right)^{-1} \frac{\partial g}{\partial X} \Sigma_{\Delta X} \left(\frac{\partial g}{\partial X}\right)' \left(\frac{\partial g}{\partial \Theta}\right)^{-1}$$

we can fix a value for α , the significance level of the hypothesis test. We can fix σ , the noise standard deviation. We can fix M , the number of different tests and we can fix J , the number of repetitions of each test. Each test consists of randomly choosing a configuration $X^{N \times 1}$ and the associated Θ . Each repetition of the test independently randomly generates a perturbation ΔX which is added to X thereby forming the observed $X + \Delta X$. From the observed $X + \Delta X$, the estimate $\hat{\Theta}$ is calculated by the optimization technique or whatever computational technique. The distribution of $\hat{\Theta}$ is not known.

But to a first order approximation, the expected value of each $\hat{\Theta}$ is known to be Θ and the covariance of each $\hat{\Theta}$ is known to be $\Sigma_{\hat{\Theta}}$. Hence, after J repetitions, the experimentally determined $\bar{\Theta}$ should have a multivariate Normal distribution with mean Θ and covariance $\Sigma_{\hat{\Theta}}/J$. Therefore, the test statistic

$$J(\bar{\Theta} - \Theta)' \Sigma_{\hat{\Theta}}^{-1} (\bar{\Theta} - \Theta)$$

will have a χ -squared distribution with N degrees of freedom. So for each of the M different tests we generate a test statistic which is χ -squared distribution with N degrees of freedom. Since each of these tests is independent, if we sum the test together we obtain a test statistic, say T , which is χ -squared distribution with $M \times N$ degrees of freedom. So to perform the test, we determine the value t_0 satisfying that the probability of a χ -squared distributed random variable with $N \times M$ degrees of freedom taking a value greater than t_0 is equal to the significance level α . Then if the test statistic T is greater than t_0 , we reject the hypothesis. If the test statistic T is less than or equal to t_0 , we tentatively accept the hypothesis.

Once we have accepted the hypothesis that the expected value of the random vector $\hat{\Theta}$ is Θ and that $\hat{\Theta}$ has covariance $\Sigma_{\hat{\Theta}}$, where

$$\Sigma_{\hat{\Theta}} = \left(\frac{\partial g}{\partial \Theta} \right)^{-1} \frac{\partial g}{\partial X} \Sigma_{\Delta X} \left(\frac{\partial g}{\partial X} \right)' \left(\frac{\partial g}{\partial \Theta} \right)^{-1}$$

we can test the hypothesis that the estimated covariance matrix $\hat{\Sigma}_{\hat{\Theta}}$ is indeed an estimate for $\Sigma_{\hat{\Theta}}$. We can perform M different tests. For each test, we will randomly generate a configuration X and corresponding Θ . Each test will have J repetitions in which a perturbation ΔX is added to X and the corresponding $\hat{\Theta}$ and $\hat{\Sigma}_{\hat{\Theta}}$ is calculated. The resulting J estimated covariance matrices can be averaged together forming the experimental average covariance matrix $\bar{\Sigma}_{\hat{\Theta}}$ and the squared difference between the diagonal and upper half entries of $\Sigma_{\hat{\Theta}}$ and $\bar{\Sigma}_{\hat{\Theta}}$ should then be sufficiently small. If enough of the M tests result in squared differences which are sufficiently small then we accept the hypothesis.

Another hypothesis we can test is whether the experimentally determined covariances are sufficiently close to the theoretically predicted covariances. Here for each configuration X and associated Θ , we perform J repetitions calculating the estimate $\hat{\Theta}_j$ on the j^{th} repetition. The experimentally determined covariance $\bar{\Sigma}_{\hat{\Theta}}$ is given by

$$\bar{\Sigma}_{\hat{\Theta}} = \frac{1}{J} \sum_{j=1}^J (\hat{\Theta}_j - \Theta)(\hat{\Theta}_j - \Theta)'$$

As J gets large we should observe that the difference between the analytically derived covariance $\Sigma_{\hat{\Theta}}$ and the experimentally determined covariance $\bar{\Sigma}_{\hat{\Theta}}$ goes to zero.

7 Conclusion

Making a successful vision system for any particular application typically requires many steps, the optimal choice of which is not always apparent. To understand how to do the optimal design, a synthesis problem, requires that we first understand how to solve the analysis problem: given the steps of a particular algorithm, determine how to propagate the parameters of the perturbation process from the input to the parameters describing the perturbation process of the computed output. The first basic case of this sort of uncertainty propagation is the propagation of the covariance matrix of the input to the covariance matrix of the output. This is what this paper has described.

This work does not come near to solving what is required for the general problem, because the general problem involves perturbations which are not additive. That is, in mid and high-level vision, the appropriate kinds of perturbations are perturbations of structures. Now, we are in the process of understanding some of the issues with these kinds of perturbations and expect to soon have some results in this area.