Is covariance information useful in estimating vision parameters?

Michael J. Brooks, Wojciech Chojnacki, Darren Gawley and Anton van den Hengel

Department of Computer Science, Adelaide University
Adelaide, SA 5005, Australia

ABSTRACT

This paper assesses some of the practical ramifications of recent developments in estimating vision parameters given information characterising the uncertainty of the data. This uncertainty information may sometimes be estimated in association with the observation process, and is usually represented in the form of covariance matrices. An empirical study is made of the conditions under which improved parameter estimates can be obtained from data when covariance information is available. We explore, in the case of fundamental matrix estimation and conic fitting, the extent to which the noise should be anisotropic and inhomogeneous if improvements over traditional methods are to be obtained. Critical in this is the devising of synthetic experiments under which noise conditions can be precisely controlled. Given that covariance information is, in itself, subject to estimation error, tests are also undertaken to determine the impact of imprecise covariance information upon the quality of parameter estimates. We thus investigate the consequences for parameter estimation of inaccuracies in the characterisation of noise that inevitably arise in practical computation.

Keywords: Parameter estimation, covariance matrix, fundamental matrix, epipolar equation, conic

1. INTRODUCTION

Many problems in computer vision can be couched in terms of the estimation of parameters from image-based measurements. Such problems arise in stereo vision, with the estimation of the fundamental matrix,[1–14] in conic fitting, with the estimation of (say) an ellipse's coefficients,[14–27] and in many other areas. Because such problems are typically very sensitive to noise in the data, there has been considerable interest in recent years in the question of how parameter estimation might be improved if additional information is available characterising the uncertainty of the data.[2,6–9,14,18–21,25,28–33] Thus, for example, if information is available indicating the uncertainties of the locations of various corresponding points, it might be possible to estimate a fundamental matrix more accurately. This uncertainty information is usually expressed in terms of *covariance matrices*.

This paper investigates conditions under which the use of covariance matrices enables parameters of improved quality to be obtained. It utilises an implementation of a state-of-the-art estimation method of the authors.[34] Several novel synthetic experiments are carried out under carefully controlled conditions of noise. Prior to describing these, however, we first examine the mathematical nature of covariance matrices.

2. COVARIANCE MATRICES

Suppose that x is a scalar measurement or observation of an underlying true value $\bar{x}$, and that the measurement process exhibits errors conforming to a zero-mean Gaussian distribution. Treating x as a sample value of a random variable, we may write, for some non-negative number σ,

$$x = \bar{x} + \Delta x, \qquad \Delta x \sim \mathrm{N}(0, \sigma^2),$$

where Δx is a random variable representing measurement errors and '$\sim \mathrm{N}(0, \sigma^2)$' means 'distributed with the Gaussian distribution with zero *mean* and *standard deviation* σ (or, equivalently, *variance* σ^2)'. The characteristic parameters of the distribution are defined explicitly by

$$\mathsf{E}\,[\Delta x] = 0, \qquad \mathsf{E}\,\left[(\Delta x)^2\right] = \sigma^2,$$

where $\mathsf{E}[y]$ denotes the expected value of the random variable y.

Authors' e-mail addresses: {`mjb,wojtek,dg,hengel`}`@cs.adelaide.edu.au`

A multi-dimensional analogue of this model is readily developed using covariance matrices. Confining our attention to the two-dimensional case, suppose now that a 2-vector $\boldsymbol{x}$ represents a measurement of a true value $\overline{\boldsymbol{x}}$ via a process that is subject to bivariate zero-mean Gaussian errors. We then have, for some positive definite matrix $\boldsymbol{\Lambda}$,

$$\boldsymbol{x} = \overline{\boldsymbol{x}} + \Delta\boldsymbol{x}, \qquad \Delta\boldsymbol{x} \sim \mathrm{N}(\boldsymbol{0}, \boldsymbol{\Lambda}),$$

where $\mathrm{N}(\boldsymbol{0}, \boldsymbol{\Lambda})$ denotes the bivariate Gaussian distribution with mean $\boldsymbol{0}$ and covariance matrix $\boldsymbol{\Lambda}$. With $\Delta\boldsymbol{x} = (\Delta x_1, \Delta x_2)^T$, the parameters of the distribution are defined explicitly by

$$(\mathsf{E}[\Delta x_1], \mathsf{E}[\Delta x_2])^T = \boldsymbol{0}, \qquad \boldsymbol{\Lambda} = \mathsf{E}\left[(\Delta\boldsymbol{x})(\Delta\boldsymbol{x})^T\right] = \begin{bmatrix} \mathsf{E}\left[(\Delta x_1)^2\right] & \mathsf{E}[\Delta x_1 \Delta x_2] \\ \mathsf{E}[\Delta x_2 \Delta x_1] & \mathsf{E}\left[(\Delta x_2)^2\right] \end{bmatrix}.$$

The simplest form of a covariance matrix is given by

$$\boldsymbol{\Lambda} = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix}.$$

The graph of the corresponding probability density function (pdf) is a rotationally symmetric, bell-shaped surface, with circular level-sets. Perturbations governed by this pdf are equally likely to arise in any direction from the underlying true value. For this reason, the corresponding probability distribution is termed *isotropic*.

A special form of an *anisotropic* distribution has a covariance matrix of the form

$$\boldsymbol{\Lambda} = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix}$$

with $\sigma_1^2 < \sigma_2^2$. Perturbations in the x- and y-directions governed by the corresponding pdf are independent of each other and have variances σ_1^2 and σ_2^2, respectively. The level sets of the graph of the pdf are ellipses with the minor and major axes aligned with the x- and y-axes, respectively. Alternatively, $\boldsymbol{\Lambda}$ can be represented as

$$\boldsymbol{\Lambda} = \alpha \begin{bmatrix} \beta & 0 \\ 0 & 1 - \beta \end{bmatrix},$$

where

$$\alpha = \sigma_1^2 + \sigma_2^2, \qquad \beta = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2}$$

define the *scale* and *eccentricity* of $\boldsymbol{\Lambda}$ and take values in $[0, \infty)$ and $[0, 1/2)$, respectively.

The most general form of a covariance matrix may be expressed as

$$\boldsymbol{\Lambda} = \boldsymbol{R}_\gamma \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{bmatrix} \boldsymbol{R}_\gamma^T,$$

where σ_1^2 and σ_2^2 are non-negative numbers with $\sigma_1^2 \leq \sigma_2^2$, γ is a number in $[0, \pi)$, and

$$\boldsymbol{R}_\gamma = \begin{bmatrix} \cos\gamma & -\sin\gamma \\ \sin\gamma & \cos\gamma \end{bmatrix}$$

is the matrix representing rotation by the angle γ. The level sets of the graph of the corresponding pdf are ellipses as before, but rotated by the angle γ. Alternatively, $\boldsymbol{\Lambda}$ can be represented as

$$\boldsymbol{\Lambda} = \alpha\, \boldsymbol{R}_\gamma \begin{bmatrix} \beta & 0 \\ 0 & 1 - \beta \end{bmatrix} \boldsymbol{R}_\gamma^T.$$

This makes explicit the geometrical factors determining the nature of the elliptical level sets, namely the overall scale, α, the eccentricity, β, and the angle of orientation, γ, the first two parameters being given by

$$\alpha = \operatorname{Tr}\boldsymbol{\Lambda}, \qquad \beta = \frac{1}{2}(1 - [1 - (4\det\boldsymbol{\Lambda})/(\operatorname{Tr}\boldsymbol{\Lambda})^2]^{1/2}).$$

The ranges of α, β, and γ, are $[0, \infty)$, $[0, 1/2]$ and $[0, \pi)$, respectively.

2.1. Collections of Covariances Matrices

It is useful to review the terminology associated with collections of covariances matrices associated with data. A simple composite model of measurement error is *homogeneous isotropic* noise in which all items of data are assumed to have identical isotropic covariance. This is depicted graphically in Figure 1(a), with the level sets being circles of the same size. In the event that noise is isotropic but that the standard deviation of each Gaussian distribution may change from point to point, we have *inhomogeneous isotropic* noise depicted in Figure 1(b). Here, each covariance matrix may be a different multiple of the identity matrix. Similarly, *homogeneous anisotropic* errors are depicted in Figure 1(c), while Figure 1(d) captures the most general situation we shall deal with, namely *inhomogeneous anisotropic* noise.

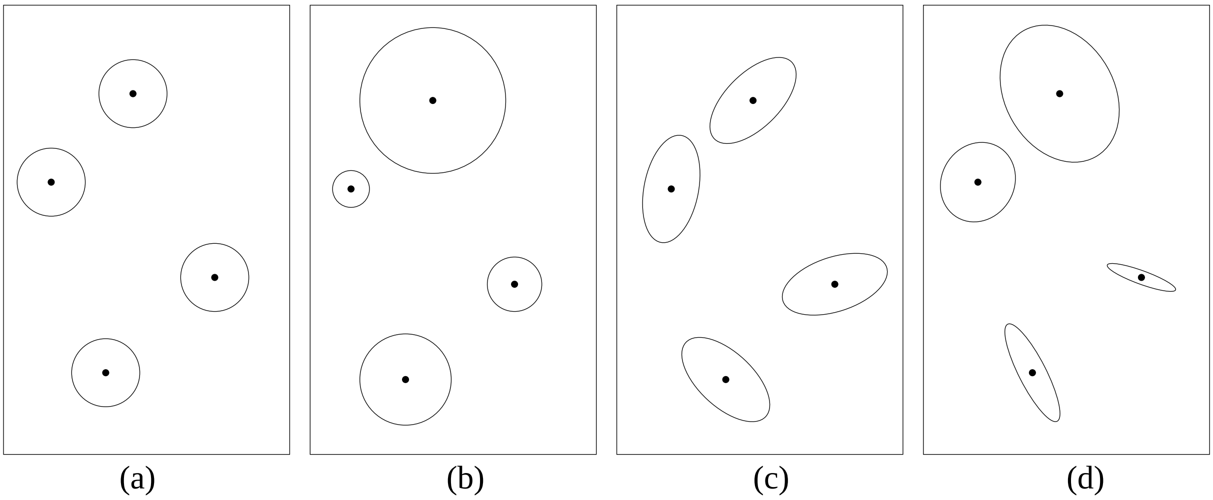

Figure 1. Different types of noise model: (a) isotropic inhomogeneous, (b) isotropic inhomogeneous, (c) anisotropic homogeneous, (d) anisotropic inhomogeneous.

3. EXPERIMENTAL SETUP

Two key estimation problems in computer vision were considered, namely ellipse (conic section) fitting and estimation of the fundamental matrix. Because both of these problems require measurement of image-point locations, there is an opportunity to estimate covariances associated with the measurement process.

Our experiments were synthetic in nature as these permit precise control of the noise conditions. Randomly generated “true data” were first obtained. A corresponding covariance matrix was then generated for each of the true data points. Next, “noisy data” were obtained by perturbing the true data points in a way consistent with their respective covariance matrices.

Three methods for computing the estimate were employed:

- **ALS** = Algebraic least squares scheme—without covariances.
- **FNS** = Fundamental numerical scheme—with full covariances.
- **FNS*** = **FNS**—with identity covariances.

A full description of these methods is provided in Ref. 34. **ALS** is employed as a simple one-step least squares scheme that uses singular-value decomposition and does not employ covariance information. **FNS** is used as a sophisticated iterative method able fully to utilise covariance information. Its accuracy has been shown empirically to be almost identical with that of a Levenberg-Marquardt implementation. Finally, **FNS*** is the **FNS** method supplied with (default) identity covariances. It may be regarded as a top-performing non-covariance method, and as such provides a benchmark for our tests. Clearly, the extent to which **FNS** improves upon **FNS*** is a direct indication of the usefulness of the covariances employed.

3.1. Generating True Data and Perturbed Data

In the case of ellipse fitting, a randomly oriented ellipse was generated such that the ratio of its major to minor axes was in the range $[2, 3]$, and its major axis was approximately 200 pixels in length. About one third of the ellipse's boundary was chosen as the base curve, and this included the point of maximum curvature of the ellipse. A set of true points $\overline{\boldsymbol{x}}$ was then randomly selected from a distribution uniform along the length of the base curve.

For estimating the fundamental matrix, a realistic stereo camera configuration was first selected with non co-planar optical axes, and slightly differing left and right camera intrinsic parameters. Randomly chosen 3D points were then projected onto the images so as to generate many pairs of corresponding points $(\overline{\boldsymbol{x}}, \overline{\boldsymbol{x}}')$.

Assuming a particular *average level of noise*, σ, for each true point $\overline{\boldsymbol{x}}$, a covariance matrix $\boldsymbol{\Lambda}_{\overline{\boldsymbol{x}}}$ was generated by drawing α, β and γ randomly from the distributions $\mathrm{U}(0, 2\sigma)$, $\mathrm{U}(0, \frac{1}{2})$, $\mathrm{U}(0, 2\pi)$, respectively. Here, $\mathrm{U}(a, b)$ denotes the uniform distribution over $[a, b]$. Since $\mathrm{Tr}\, \boldsymbol{\Lambda}_{\overline{\boldsymbol{x}}} = \alpha$ and $\mathsf{E}\,[\alpha] = \sigma$, it follows that $\mathsf{E}\,[\mathrm{Tr}\, \boldsymbol{\Lambda}_{\overline{\boldsymbol{x}}}] = \sigma$, giving the statistical interpretation of the assumed level of noise.

Given a true point $\overline{\boldsymbol{x}}$ and an associated covariance matrix $\boldsymbol{\Lambda}_{\overline{\boldsymbol{x}}}$, a noisy point $\boldsymbol{x}$ consistent with $\boldsymbol{\Lambda}_{\overline{\boldsymbol{x}}}$ was obtained by adding a vector $\Delta\boldsymbol{x}$ to $\overline{\boldsymbol{x}}$. The vector $\Delta\boldsymbol{x}$ was generated with use of the following algorithm:

1. Find a matrix $\boldsymbol{U}$ such that $\boldsymbol{\Lambda}_{\boldsymbol{x}} = \boldsymbol{U}\,\boldsymbol{U}^T$; this can be done by performing, say, Cholesky decomposition of $\boldsymbol{\Lambda}_{\boldsymbol{x}}$.[35]
2. Generate a random vector $\boldsymbol{p} = [p_1, p_2]^T$, where each element is drawn independently from the Gaussian distribution with zero mean and unit standard deviation.
3. Set $\Delta\boldsymbol{x} = \boldsymbol{U}\boldsymbol{p}$.

Statistically, $\Delta\boldsymbol{x}$ has mean $\mathbf{0}$ and covariance $\boldsymbol{\Lambda}$, since

$$\mathsf{E}\,[\Delta\boldsymbol{x}] = \mathsf{E}\,[\boldsymbol{U}\boldsymbol{p}] = \boldsymbol{U}\mathsf{E}\,[\boldsymbol{p}] = \mathbf{0},$$

$$\mathsf{E}\left[(\Delta\boldsymbol{x})(\Delta\boldsymbol{x})^T\right] = \mathsf{E}\left[(\boldsymbol{U}\boldsymbol{p})(\boldsymbol{U}\boldsymbol{p})^T\right] = \mathsf{E}\left[\boldsymbol{U}(\boldsymbol{p}\boldsymbol{p}^T)\boldsymbol{U}^T\right] = \boldsymbol{U}\mathsf{E}\left[\boldsymbol{p}\boldsymbol{p}^T\right]\boldsymbol{U}^T = \boldsymbol{U}\,\boldsymbol{U}^T = \boldsymbol{\Lambda}.$$

Finally, each $\boldsymbol{x}$ was equipped with a covariance matrix $\boldsymbol{\Lambda}_{\boldsymbol{x}}$ which was taken to be $\boldsymbol{\Lambda}_{\overline{\boldsymbol{x}}}$.

4. RESULTS USING EXACT COVARIANCES

4.1. General Utility of Covariances

The estimates computed by the three methods **ALS**, **FNS***, and **FNS** were compared using an appropriate error function. Specifically, for each estimated ellipse, an error measure was computed as the sum of the shortest pixel distances of each *true point* (lying on the true ellipse) from the estimated ellipse. For each fundamental matrix estimated, an error measure was computed as the sum of the distances of the underlying *true points* to the epipolar lines derived from the estimated fundamental matrix, averaged for the left and right images.

Each test involved 60 data points generated independently in each repetition (with new true data points and covariance matrices generated at each iteration). The average value of the error for each of the estimation methods was computed over 2000 runs. The entire process was performed for an average noise level σ ranging from 0.5 to 10.0 in steps of 0.5. The results presented in Figure 2 show that employing covariance information in estimation can be advantageous. In the case of ellipse fitting and fundamental matrix estimation, errors obtained by **FNS** are around 40% and 70%, respectively, of those obtained via **FNS***. **ALS** lags further behind **FNS***.

4.2. Varying the Spread of Covariances

A further experiment was carried out in which the level of anisotropicity and inhomogeneity was controlled. This was done by introducing parameter $0 \leq \rho \leq 1$ to capture the 'spread' of the skew and scale parameters. Fixing the level of noise, possible variation of skew was governed by

$$\beta \sim \mathrm{U}(\frac{1}{2}(1-\rho), \frac{1}{2}(1+\rho))$$

and variation of scale by

$$\alpha \sim \mathrm{U}(\sigma(1-\rho), \sigma(1+\rho)).$$

The following tests were performed for ellipse fitting and fundamental matrix estimation:

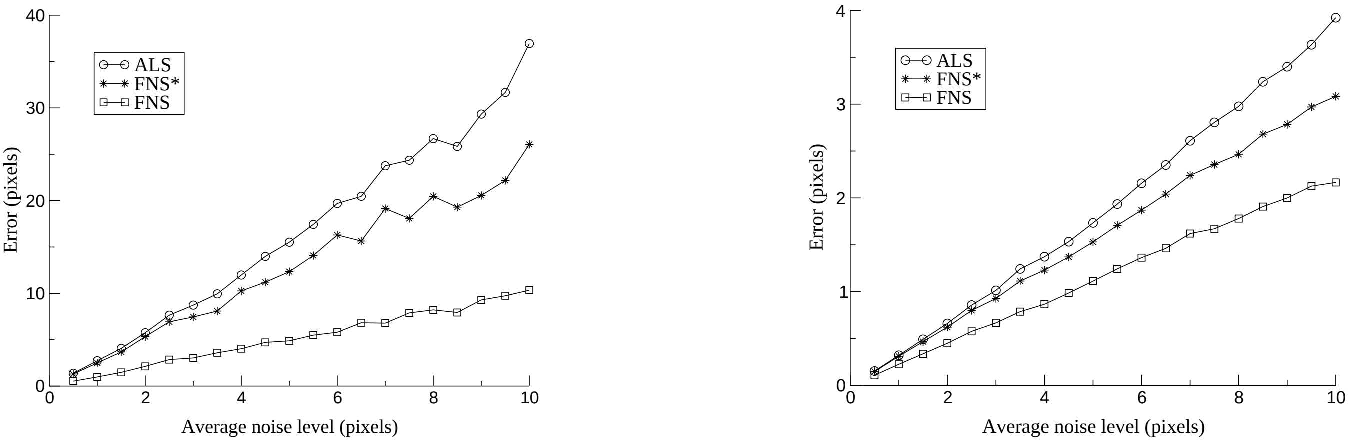

Figure 2. Error for ellipses fitting (left), and fundamental matrix estimation (right).

- Skew was selected as above, and scale was kept constant. This corresponds to homogenous anisotropic noise.
- Scale was selected as above, and skew was set to $1/2$. This corresponds to inhomogenous isotropic noise.
- Both parameters were selected as above. This corresponds in general to inhomogenous anisotropic noise.

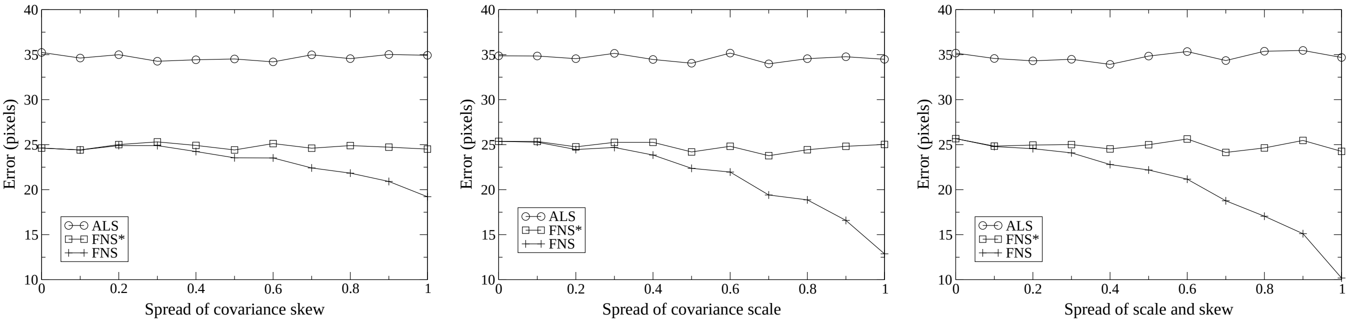

Figure 3. Increasing spread (ρ) for ellipse estimation.

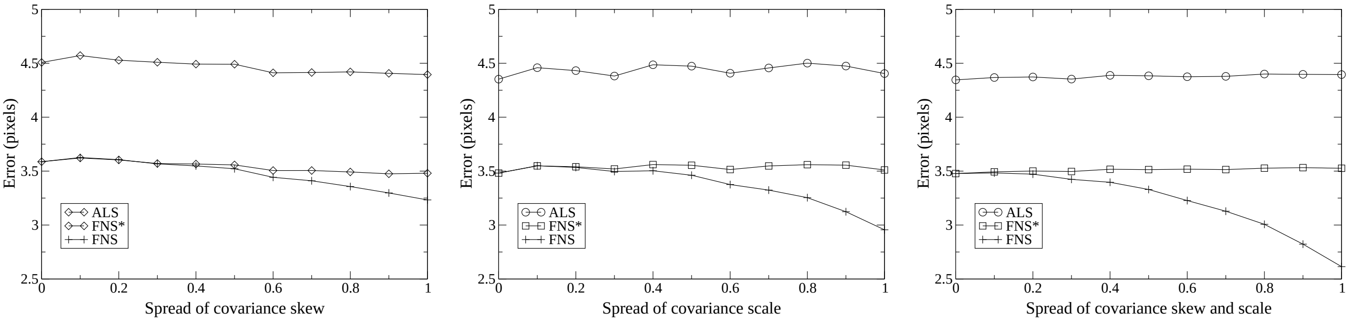

Figure 4. Increasing spread (ρ) for fundamental matrix estimation.

The results are presented in Figures 3 and 4. As ρ increases, the error of the estimate generated using covariance matrices reduces. In our circumstances, this improvement is more marked for the scale parameter than the skew parameter. Interestingly, the average error for the two methods not using covariances remains constant regardless of the value of ρ.

5. RESULTS USING INEXACT COVARIANCES

In practice, it is not possible to measure exactly the covariances driving the noise in the data. Therefore, in this section, we investigate the robustness of estimators to inaccuracies in the covariance matrices. This we do by supplying an estimator not with the true underlying covariance matrix, but with a version to which noise has been added. This we may regard as *meta noise*. Two experiments were carried out using inexact covariances.

5.1. Partially Correct Covariances

The first experiment involved generating true covariance matrices by randomly choosing all three parameters (as described earlier) and producing points perturbed consistently with these covariances. However, the estimation method was supplied with a covariance matrix that used only one of the original parameter values, with default values being supplied for the others.

The results in Figure 5 show average errors obtanained under the following curcumstances:

True. The correct covariances generated are passed to the estimator.

Identity. Identity covariances are used.

Skew. A new covariance is formed using the original skew parameter, setting scale to $\sigma/2$, and rotation at 0.

Angle. The original rotation angle is used, scale is set to $\sigma/2$, and skew to $1/3$.

Scale. The original scale is used, with skew set to $1/2$, and rotation to 0.

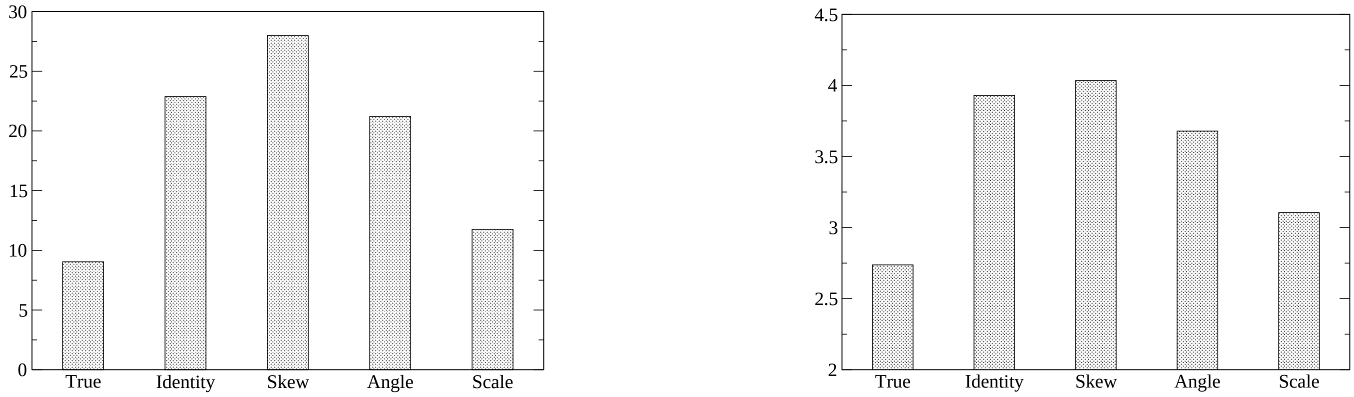

Figure 5. Errors obtained using partially correct covariances in ellipse fitting (left) and the fundamental matrix estimation (right).

We infer from Figure 5 that, if we can only have one correct component of covariance, then scale is the most valuable for the purposes of estimation. Using only the skew parameter results unsurprisingly in larger errors than those obtained using identity covariances.

5.2. Fully Perturbed Covariances

The second experiment involved contaminating the true covariance matrices with noise by multiplying each of the underlying parameters by a random factor, $(\alpha, \beta, \gamma) \mapsto (\kappa_1\alpha, \kappa_2\beta, \kappa_3\gamma)$, where $\kappa_1, \kappa_2, \kappa_3$ are chosen independently from the Gaussian distribution with unit mean and standard deviation τ. In this way, the level of noise added to the covariances is controlled by the deviation τ. It should be noted that whenever the multiplication of the parameter caused it to exit the range specified in Section 3.1, it was clipped at the maximum or minimum appropriately.

The algorithm can be summarised as follows:

1. Generate synthetic *true* data.
2. Generate random covariance matrices as per Section 3.1, with a fixed average noise level.
3. Perturb data in accordance with the covariance matrices to obtain noisy data.
4. Create a new set of covariance matrices by adding noise to the original matrices.
5. Compare estimates obtained by supplying identity, true, and noisy covariance matrices.

Again, these steps were repeated 2000 times to obtain an average error for the three estimation methods. However, in contrast with the previous tests where this process was repeated for different values of σ, these tests were repeated with varying τ.

Figure 6 shows the results. The error of the estimate computed by **FNS** using the noisy covariances is seen to rise linearly with the error. In our experiments, we see that noisy covariances offer advantage over identity defaults when $\tau < 0.5$ in the case of ellipse estimation, and when $\tau < 0.3$ in the case of fundamental matrix estimation.

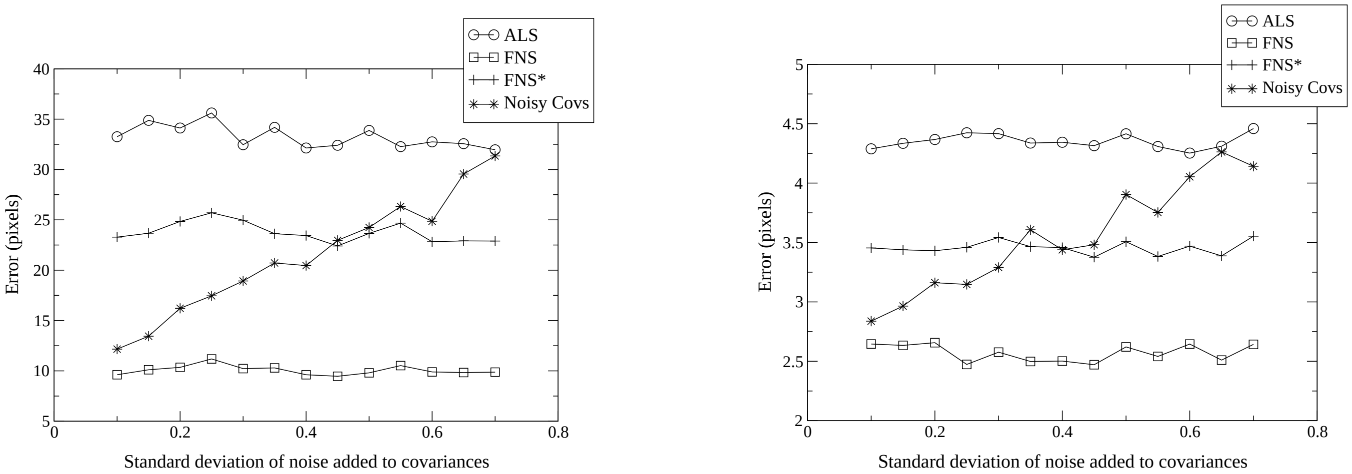

Figure 6. Average error for increasing average error in covariances.

6. CONCLUSION

A series of tests were presented that operate upon data in the form of image measurements. These were designed to determine the degree of benefit attained by using covariance information in the parameter estimation process. Our experiments indicated not only that covariance information can be useful, but also the extent to which this information may be inaccurate before the advantage is lost. It emerged that relatively inaccurate covariance information can sometimes improve the quality of parameter estimates, and that it is the relative scale of the covariances which is the most valuable factor.

ACKNOWLEDGMENTS

This work was in part funded by the Australian Research Council and the Cooperative Research Centre for Sensor Signal and Information Processing.

REFERENCES

1. F. Bookstein, “Fitting conic sections to scattered data,” *Computer Graphics and Image Processing* **9**(1), pp. 56–71, 1979.
2. T. Ellis, “Ellipse detection and matching with uncertainty,” *Image and Vision Computing* **10**(2), pp. 271–276, 1992.
3. A. Fitzgibbon, M. Pilu, and R. B. Fisher, “Direct least square fitting of ellipses,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **21**(5), pp. 476–480, 1999.
4. A. W. Fitzgibbon and R. B. Fisher, “A buyer’s guide to conic fitting,” in *Proceedings of the Sixth British Machine Vision Conference*, D. Pycock, ed., vol. 2, pp. 513–522, BMVA, (Birmingham, England, 11-14 September, 1995), 1995.
5. W. Gander, G. H. Golub, and R. Strebel, “Least-square fitting of circles and ellipses,” *BIT* **43**, pp. 558–578, 1994.
6. K. Kanatani, “Statistical bias of conic fitting and renormalisation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **16**(3), pp. 320–326, 1994.
7. Y. Kanazawa and K. Kanatani, “Optimal conic fitting and reliability evaluation,” *IEICE Transactions on Information and Systems* **E79-D**, pp. 1323–1328, Sept. 1996.
8. B. Matei and P. Meer, “Reduction of bias in maximum likelihood ellipse fitting,” in *Proceedings of the 15th International Conference on Pattern Recognition*, vol. 3, pp. 802–806, IEEE Computer Society Press, Los Alamitos, CA, (Barcelona, Spain, September 3–8, 2000), 2000.
9. J. Porill, “Fitting ellipses and predicting confidence envelopes using a bias corrected Kalman filter,” *Image and Vision Computing* **8**(1), pp. 37–41, 1990.
10. P. L. Rosin, “A note on the least square fitting of ellipses,” *Pattern Recognition Letters* **14**(10), pp. 799–808, 1993.
11. P. L. Rosin and G. A. West, “Nonparametric segmentation of curves into various representations,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17**(12), pp. 1140–1153, 1995.
12. P. D. Sampson, “Fitting conic sections to ‘very scattered’ data: An iterative refinement of the Bookstein algorithm,” *Computer Graphics and Image Processing* **18**(1), pp. 97–108, 1982.
13. G. Taubin, “Estimation of planar curves, surfaces and nonplanar space curves defined by implicit equations, with applications to edge and range image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(11), pp. 1115–1138, 1991.
14. Z. Zhang, “Parameter estimation techniques: a tutorial with application to conic fitting,” *Image and Vision Computing* **15**(1), pp. 57–76, 1997.
15. O. Faugeras, “What can be seen in three dimensions with an uncalibrated stereo rig?,” in *Computer Vision—ECCV ’92, Second European Conference on Computer Vision*, G. Sandini, ed., vol. 588 of *Lecture Notes in Computer Science*, pp. 563–578, Springer, Berlin, (Santa Margherita Ligure, Italy, May 19–22, 1992), 1992.
16. R. Hartley, “In defense of the eight-point algorithm,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(6), pp. 580–593, 1997.
17. R. Hartley, R. Gupta, and T. Chang, “Stereo from uncalibrated cameras,” in *Proceedings, CVPR ’92, IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 761–764, IEEE Computer Society Press, Los Alamitos, CA, (Champaign, IL, June 15–18, 1992), 1992.
18. K. Kanatani, “Unbiased estimation and statistical analysis of 3-D rigid motion from two views,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **15**(1), pp. 37–50, 1993.
19. Q.-T. Luong and O. D. Faugeras, “The fundamental matrix: theory, algorithms, and stability analysis,” *International Journal of Computer Vision* **17**(1), pp. 43–75, 1996.
20. M. Mühlich and R. Mester, “The role of total least squares in motion analysis,” in *Computer Vision—ECCV’98, Fifth European Conference on Computer Vision*, H. Burkhardt and B. Neumann, eds., vol. 1407 of *Lecture Notes in Computer Science*, pp. 305–321, Springer, Berlin, (Freiburg, Germany, June 2–6, 1998), 1998.
21. M. Mühlich and R. Mester, “Subspace methods and equilibration in computer vision,” Tech. Rep. XP-TR-C-21, Institute for Applied Physics, Johann Wolfgang Goethe-Universität, Frankfurt am Main, Germany, Nov. 1999.
22. C. V. Stewart, “Robust parameter estimation in computer vision,” *SIAM Review* **41**(3), pp. 513–537, 1999.
23. P. H. S. Torr and D. W. Murray, “The development and comparison of robust methods for estimating the fundamental matrix,” *International Journal of Computer Vision* **24**(3), pp. 271–300, 1997.

24. J. Weng, T. S. Huang, and N. Ahuja, “Motion and structure from two perspective views: algorithms, error analysis, and error estimation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(5), pp. 451–476, 1989.
25. Z. Zhang, “Determining the epipolar geometry and its uncertainty: a review,” *International Journal of Computer Vision* **27**(2), pp. 161–195, 1998.
26. Z. Zhang, “On the optimization criteria used in two-view motion analysis,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **20**(7), pp. 717–729, 1998.
27. Z. Zhang, R. Deriche, O. Faugeras, and Q.-T. Luong, “A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry,” *Artificial Intelligence* **78**, pp. 87–119, 1995.
28. W. Förstner, “On weighting and choosing constraints for optimally reconstructing the geometry of image triplets,” in *Computer Vision—ECCV 2000, Sixth European Conference on Computer Vision*, D. Vernon, ed., vol. 2 of *Lecture Notes in Computer Science*, pp. 669–684, Springer, Berlin, (Trinity College, Dublin, Ireland, June 26–July 1, 2000), 2000.
29. Y. Leedan and P. Meer, “Estimation with bilinear constraints in computer vision,” in *Proceedings of the Sixth International Conference on Computer Vision*, pp. 733–738, IEEE Computer Society, Narosa Publishing House, New Delhi, (Bombay, India, January 4–7, 1998), 1998.
30. Y. Leedan and P. Meer, “Heteroscedastic regression in computer vision: problems with bilinear constraint,” *International Journal of Computer Vision* **37**(2), pp. 127–150, 2000.
31. B. Matei and P. Meer, “A general method for errors-in-variables problems in computer vision,” in *Proceedings, CVPR 2000, IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 18–25, IEEE Computer Society Press, Los Alamitos, CA, (Hilton Head Island, South Carolina, June 13-15, 2000), 2000.
32. B. Triggs, “A new approach to geometric fitting.” Available from `http://www.inrialpes.fr/movi/people/Triggs`.
33. B. Triggs, “Optimal estimation of matching constraints,” in *3D Structure from Multiple Images of Large-Scale Environments (SMILE), workshop organised by the CUMULI, PANORAMA and VANGUARD projects in conjunction with ECCV’98*, R. Koch and L. V. Gool, eds., pp. 30–32, Freiburg University, (Freiburg, Germany, June 6–7, 1998), 1998. Full version available from `http://www.inrialpes.fr/movi/people/Triggs`.
34. W. Chojnacki, M. J. Brooks, and A. van den Hengel, “On the fitting of surfaces to data with covariances,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22**(11), 2000.
35. G. H. Golub and C. F. van Loan, *Matrix Computations*, The Johns Hopkins University Press, 2nd ed., 1989.