

What Value Covariance Information in Estimating Vision Parameters?

Michael J. Brooks, Wojciech Chojnacki, Darren Gawley, Anton van den Hengel
Department of Computer Science, Adelaide University
Adelaide, SA 5005, Australia

Abstract

Many parameter estimation methods used in computer vision are able to utilise covariance information describing the uncertainty of data measurements. This paper considers the value of this information to the estimation process when applied to measured image point locations. Covariance matrices are first described and a procedure is then outlined whereby covariances may be associated with image features located via a measurement process. An empirical study is made of the conditions under which covariance information enables generation of improved parameter estimates. Also explored is the extent to which the noise should be anisotropic and inhomogeneous if improvements are to be obtained over covariance-free methods. 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. Finally, an experiment is carried out to assess the value of covariances in estimating the fundamental matrix from real images.

1. Introduction

Many problems in computer vision can be couched in terms of parameter estimation from image-based measurements. Such problems arise in stereo vision, with the estimation of the fundamental matrix [8, 12, 15, 18], in conic fitting, with the estimation of an ellipse's coefficients [3, 17], and in many other areas. Because such problems are typically very sensitive to noise, there has recently been considerable interest in assessing how parameter estimation might be improved if additional information is available characterising the uncertainty of the data [1, 4, 10, 11, 16]. This uncertainty information is usually expressed in terms of *covariance matrices*.

This paper investigates conditions under which the use of covariance matrices enables parameter estimates of im-

proved quality to be obtained. Several novel experiments are carried out under carefully controlled conditions. Prior to describing these, however, we first examine the mathematical form of covariance matrices.

2. Covariance matrices

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

$$p = \bar{p} + \Delta p, \quad \Delta p \sim N(0, \sigma^2),$$

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

$$E[\Delta p] = 0, \quad E[(\Delta p)^2] = \sigma^2,$$

where $E[y]$ denotes the expected value of the random variable y .

A multi-dimensional analogue of this model is readily developed using covariance matrices. Confining our attention to the two-dimensional case, suppose that a two-entry vector \mathbf{p} represents a contaminated measurement of a true value $\bar{\mathbf{p}}$, with measurement errors in each entry conforming to a zero-mean Gaussian distribution. We then have, for some positive definite matrix \mathbf{A} ,

$$\mathbf{p} = \bar{\mathbf{p}} + \Delta \mathbf{p}, \quad \Delta \mathbf{p} \sim N(\mathbf{0}, \mathbf{A}),$$

where $N(\mathbf{0}, \mathbf{A})$ denotes the bivariate Gaussian distribution with mean $\mathbf{0}$ and covariance matrix \mathbf{A} . With $\Delta \mathbf{p} = (\Delta x, \Delta y)^T$, the parameters of the distribution are defined explicitly by

$$(E[\Delta x], E[\Delta y])^T = \mathbf{0},$$

$$\mathbf{A} = E[(\Delta \mathbf{p})(\Delta \mathbf{p})^T] = \begin{bmatrix} E[(\Delta x)^2] & E[\Delta x \Delta y] \\ E[\Delta y \Delta x] & E[(\Delta y)^2] \end{bmatrix}.$$

A simple form of covariance matrix is given by

$$\mathbf{A} = \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix}.$$

When $\sigma_x = \sigma_y$, 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*. In the event that $\sigma_x \neq \sigma_y$, the distribution is termed *anisotropic*. Perturbations in the x - and y -directions governed by the corresponding pdf are independent of each other and have variances σ_x^2 and σ_y^2 , respectively. If $\sigma_x > \sigma_y$, the level sets of the graph of the pdf are ellipses with the major and minor axes aligned with the x - and y -axes, respectively. It is useful to note that \mathbf{A} may equivalently be represented as

$$\mathbf{A} = \alpha \begin{bmatrix} \beta & 0 \\ 0 & 1 - \beta \end{bmatrix},$$

where $\alpha = \sigma_x^2 + \sigma_y^2$ and $\beta = \sigma_x^2 / (\sigma_x^2 + \sigma_y^2)$ are the *scale* and *eccentricity* of \mathbf{A} .

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

$$\mathbf{A} = \alpha \mathbf{R}_\gamma \begin{bmatrix} \beta & 0 \\ 0 & 1 - \beta \end{bmatrix} \mathbf{R}_\gamma^T, \quad (1)$$

where

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

is the matrix representing anti-clockwise rotation by the angle γ . The level sets of the graph of the corresponding pdf are ellipses as before, but rotated anti-clockwise by the angle γ . 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

$$\alpha = \text{Tr } \mathbf{A}, \quad \beta = \frac{1}{2} (1 - [1 - (4 \det \mathbf{A}) / (\text{Tr } \mathbf{A})^2]^{1/2}).$$

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

2.1. Collections of covariance matrices

It is useful to review the terminology associated with collections of covariance 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 the same isotropic distribution. This is depicted graphically in Figure 1(a) with covariance matrices represented




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

by circles of the same size. We shall term such matrices *identity* covariances, given that they are all a common (possibly unknown) multiple of the identity matrix. In the event that noise is isotropic but that the standard deviation of each 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 type of noise we shall deal with, namely *inhomogeneous anisotropic* noise.

3. Generating covariance matrices for real-image features

Underlying our analysis is the theoretical assumption that a *measurement process* locates a feature, and that this process generates a *measurement error* that conforms to a bivariate normal distribution about the true value. Our aim at this stage is to associate with each measured feature location a covariance matrix capturing the nature of this distribution. Importantly, we note that different measurement processes applied to the same feature (image) data will, in general, give rise to different covariance matrices.

With the parametrisation of covariance matrices given in (1), the following procedure represents one approach to associating a covariance matrix with a given measurement process:

- Determine the *maximal* direction along which the measurement process is susceptible to the largest perturbations. In view of our assumption that errors conform to a bivariate normal distribution, it immediately follows that the *minimal* direction is orthogonal to this. The (anti-clockwise) angle that the maximal direction makes with the horizontal corresponds to the angle γ in (1).




Figure 2. \mathbf{A}_1 covariances for image features.

- Compute the ratio r of the expected values of perturbations in the maximal and minimal directions. This ratio determines the eccentricity parameter β given in (1) via the formula $\beta = r(1 + r)^{-1}$.
- Compute an overall measure of the scale of the error distribution by adding the expected values of perturbations in the maximal and minimal directions. The sum determines the parameter α in (1). Note that subsequent parameter estimation permits use of covariance matrices with each α multiplied by a common scalar factor.

This procedure is in practice difficult to follow precisely and so heuristic methods are usually employed. Perhaps the most commonly adopted covariance matrix takes the form

$$\mathbf{A}_1 = \frac{1}{\widehat{E}_y^2 \widehat{E}_x^2 - \widehat{E}_y \widehat{E}_x \widehat{E}_x \widehat{E}_y} \begin{bmatrix} \widehat{E}_y^2 & -\widehat{E}_y \widehat{E}_x \\ -\widehat{E}_x \widehat{E}_y & \widehat{E}_x^2 \end{bmatrix}.$$

Here E_x and E_y denote the first partial derivatives of the image E with respect to x and y . \widehat{E}_x^2 stands for the square of the derivative (at a point (x, y)) smoothed over a neighbourhood (of (x, y)). One way of computing this is:

$$\begin{aligned} \widehat{E}_x^2 &= \widehat{E}_x^2(x, y) = \sum_{(i,j) \in P} G(i, j) \times (E_x(x + i, y + j))^2, \\ E_x(x, y) &= (E(x + 1, y) - E(x - 1, y)) / 2, \\ G(i, j) &= (2\pi\sigma^2)^{-1} \exp(-(i^2 + j^2)/2\sigma_P^2), \end{aligned}$$

where P is a window of pixel values centred at $(0, 0)$ (e.g. $P = [-k, k] \times [-k, k]$ for some positive integer k) and σ_P is a fixed standard deviation expressed in pixels. \widehat{E}_y^2 , $\widehat{E}_y \widehat{E}_x$ and $\widehat{E}_x \widehat{E}_y$ are defined analogously.

This covariance matrix captures the notion that features are best and worst located in the direction of high and low intensity changes, respectively. The factor $(\widehat{E}_y^2 \widehat{E}_x^2 - \widehat{E}_y \widehat{E}_x \widehat{E}_x \widehat{E}_y)^{-1}$ ensures that greater uncertainty is associated with lower overall intensity change. \mathbf{A}_1 has proved useful in such applications as optical flow computation, texture segmentation, and factorization [2, 13, 14].

Figure 2 shows the covariance ellipses generated when \mathbf{A}_1 is computed from various corner data. Each ellipse has

its minor axis coinciding with the direction of maximum gradient. This may or may not coincide with the error distribution associated with a given corner detector. Indeed, were an estimator to generate errors with distributions orthogonal to these, then the following matrix would be more appropriate:

$$\mathbf{A}_2 = \frac{1}{\widehat{E}_x^2 \widehat{E}_y^2 - \widehat{E}_x \widehat{E}_y \widehat{E}_y \widehat{E}_x} \begin{bmatrix} \widehat{E}_x^2 & \widehat{E}_x \widehat{E}_y \\ \widehat{E}_y \widehat{E}_x & \widehat{E}_y^2 \end{bmatrix}. \quad (2)$$

Covariances of similar form have been used for centroid estimation [6].

4. Experimental setup

In order to analyse the value of covariance information in the estimation process various tests were performed. The first set of tests aims to determine the accuracy required of covariances if they are to be of value. The second set of tests analyses the utility of the covariance estimation process given in Section 3.

4.1. Testing strategy

It is only in a purely synthetic environment that we can have complete control of the noise conditions. Our aim with this particular synthetic testing is to determine how accurate our covariance-based estimation process needs to be in order to improve the estimation of parameters.

Our experimental procedure was as follows. 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.

Methods for estimating parameters were:

- **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 can be found in [1]. **ALS** is employed as a simple non-iterative least squares scheme that uses singular-value decomposition and does not employ covariance information. **FNS** is used as a statistically well-founded 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. Clearly, the extent to which **FNS** improves upon **FNS*** is

a direct indication of the usefulness of the covariances employed.

4.2. Generating true and perturbed data

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 (\bar{p}, \bar{p}') .

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

Given a true point \bar{p} and an associated covariance matrix $A_{\bar{p}}$, a noisy point p consistent with $A_{\bar{p}}$ was obtained by adding a vector Δp to \bar{p} . The vector Δp was generated using the following algorithm:

1. Find a matrix U such that $A_{\bar{p}} = U U^T$; this can be done by performing, say, Cholesky decomposition of $A_{\bar{p}}$ [5].
2. Generate a random two-entry vector r , with each component drawn independently from the Gaussian distribution with zero mean and unit standard deviation.
3. Set $\Delta p = U r$.

Statistically, Δp has mean $\mathbf{0}$ and covariance $A_{\bar{p}}$, since

$$\begin{aligned} E[\Delta p] &= E[U r] = U E[r] = \mathbf{0}, \\ E[(\Delta p)(\Delta p)^T] &= E[(U r)(U r)^T] \\ &= E[U(r r^T) U^T] = U E[r r^T] U^T \\ &= U U^T = A_{\bar{p}}. \end{aligned}$$

Finally, each p was equipped with a covariance matrix A_p , which was taken to be $A_{\bar{p}}$.

5. Experiments with synthetic data

In order to provide a point of reference, we first conducted tests that estimated the fundamental matrix from noisy data, using perfect covariances. Initially, each test involved randomly choosing 60 *true points*, and generating an associated covariance matrix for each point as described in Section 4.2. These covariances were anisotropic and inhomogeneous, and parameterised by a (common) average level of noise. For each true point, a *noisy point* was

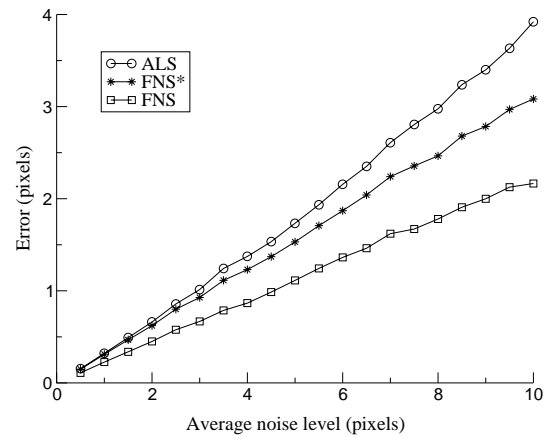


Figure 3. Error in fundamental matrix estimates.

randomly chosen consistent with the corresponding covariance matrix. The three methods **ALS**, **FNS**, and **FNS*** were then employed to estimate the fundamental matrix given the noisy points, and in the case of **FNS** the associated covariance matrices as well. An error measure was computed for each estimate. This measure was the sum of the Euclidean distances of the underlying *true points* (in both images) to the epipolar lines derived from the estimated fundamental matrix. The mean value of the error for each of the estimation methods was computed over 2000 repetitions. 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 3 show that employing covariance information in estimation can be advantageous. Errors obtained by **FNS** are around 70% of those obtained via **FNS***. **ALS** lags even further behind **FNS***.

5.1. Varying the spread of covariances

A further experiment was carried out in which the level of anisotropy 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, variation of skew was governed by

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

and variation of scale by

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

The following tests were performed:

- Skew was selected as above, and scale was kept constant. This corresponds to homogeneous anisotropic noise.




Figure 4. Errors in fundamental matrix estimates against increasing spread of covariances.

- Scale was selected as above, and skew was set to $1/2$. This corresponds to inhomogeneous isotropic noise.
- Both parameters were selected as above. This corresponds in general to inhomogeneous anisotropic noise.

The results are presented in Figure 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. For a given average level of noise, we see that estimates obtained via FNS improve with increased diversity of covariance. We note that the average error for the two methods not using covariances remains constant, since the average level of noise remains constant.

5.2. Partially correct covariances

In practice, it is not possible to measure exactly the covariances driving the noise in the data. Therefore we undertook investigation of the robustness of estimators to inaccuracies in the covariance matrices. We conducted a test 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*.

Our 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 was 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 4.2, 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 4.2, with a fixed average noise level.

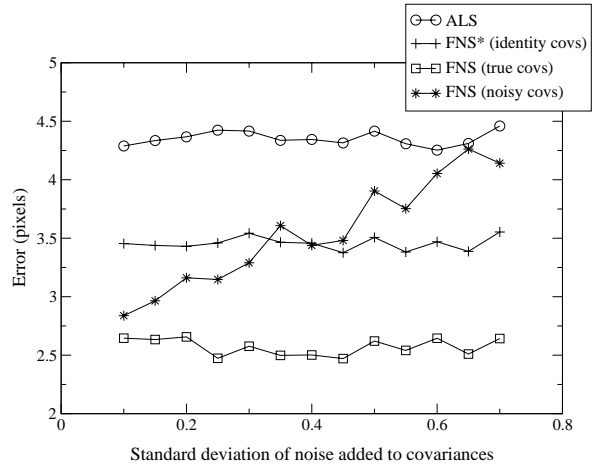


Figure 5. Average estimation error for increasing average error in covariances.

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 fixed σ and varying τ .

Figure 5 highlights the roughly linear increase in error of FNS estimates as the supplied covariance information becomes less accurate. In these experiments, we see that noisy covariances offer advantage over identity defaults when $\tau < 0.3$ pixels.




Figure 6. An enlarged section of a typical image, showing both the true corner and the corner detected by the Harris detector.

6. Experiments with real images

Our final test of the usefulness of covariance information was whether or not it can improve the accuracy of an estimated fundamental matrix relating two real images.

The Harris method [7] is commonly used as a means of extracting corner features from images for the purpose of stereo matching. This approach employs an ‘interest operator’ evaluated for each pixel, using image intensity gradients. Points which are local maxima of a function defined by the operator, and whose values are above a given threshold, are reported as features. The Harris method tends to find features “inside” of the true corner, as shown in Figure 6. The typical direction of this discrepancy is aligned with the intensity gradient, along the bisector of the corner. If this method is to be used to locate corners, it may therefore be seen as biased; certainly, it does not have the property of generating measurements whose errors are distributed in bivariate normal form.

Despite the theoretical inapplicability of the Harris detector, we nevertheless chose to use it in conjunction with the covariance matrix given in (2). For all but infrequently arising acute corners, this covariance matrix captures the notion that perturbations tend to be most significant along the bisector of the corner. Our experiment thus sought to determine whether estimation of the fundamental matrix using Harris corner data was improved by the incorporation of this covariance information.

Three example stereo image pairs were used in these tests, shown in Figure 7. The Harris detector was used to extract features from each image. For each of these points, an associated covariance matrix was computed using the form given in (2). Points in the left and right images were then matched. Between 70 and 90 points were matched for each image pair. Using these matching points, **FNS** and **FNS*** were used to compute respective fundamental matrices.




Figure 7. The grid, office, and library stereo image pairs.

Next we compared the quality of each of these estimates to see if incorporating measured covariances is advantageous. Note that we could not employ the error measure used in previous testing because the underlying true points were not known. Instead, the reprojection error was utilised in which the estimated fundamental matrix was used to back-project each matching pair, yielding a world point. In the process, the optimal triangulation method of Hartley [9] was used. This point was then reprojected onto the images to give camera consistent points. The error is the sum of the Euclidean distances between the original measured points and these reprojected points, in both the left and right images. The results of applying this to each of the estimated fundamental matrices are shown in Figure 8.

The results show that, for each of the three stereo pairs, the fundamental matrix estimated using Λ_2 covariances gives errors which are less than those obtained with identity matrices. Note that use instead of Λ_1 covariances yields greater errors.

7. Conclusion

A series of tests was undertaken to determine the degree of benefit obtained when using covariance information in




Figure 8. Errors in FNS estimation of fundamental matrix using identity, Λ_1 , and Λ_2 covariance forms.

estimating the fundamental matrix. Our experiments indicated not only that covariance information can be valuable, but also the extent to which this information may be inaccurate before the advantage is lost. We presented a heuristic method to measure covariances from real images, and showed that it may be used to improve the quality of estimates.

Acknowledgements

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

References

- [1] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. On the fitting of surfaces to data with covariances. *IEEE Trans. Pattern Analysis & Machine Intelligence*, 22(11):1294–1303, 2000.
- [2] A. Faber and W. Förstner. Scale characteristics of local auto-covariances for texture segmentation. *Int. Arch. Photogrammetry & Remote Sensing*, 32(7-4-3W6), 1999.
- [3] A. Fitzgibbon, M. Pilu, and R. B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. Pattern Analysis & Machine Intelligence*, 21(5):476–480, 1999.
- [4] W. Förstner. On weighting and choosing constraints for optimally reconstructing the geometry of image triplets. In *Proc. 6th European Conf. Computer Vision, Dublin, Ireland*, volume 1843 of *Lecture Notes in Computer Science*, pages 669–684, 2000.
- [5] G. H. Golub and C. F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, Maryland, 1989.
- [6] R. M. Haralick and L. G. Shapiro. *Computer and Robot Vision*, volume 2. Addison–Wesley, Reading, Massachusetts, 1993.
- [7] C. G. Harris. Determination of ego-motion from matched points. In *Proc. 3rd Alvey Vision Conf., Cambridge, UK*, pages 189–192, 1987.
- [8] R. Hartley. In defense of the eight-point algorithm. *IEEE Trans. Pattern Analysis & Machine Intelligence*, 19(6):580–593, 1997.
- [9] R. I. Hartley and P. Sturm. Triangulation. In *Proc. Conf. Computer Analysis of Images and Patterns, Prague, Czech Republic*, volume 970 of *Lecture Notes in Computer Science*, pages 190–197, 1995.
- [10] K. Kanatani. *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, Amsterdam, 1996.
- [11] Y. Leedan and P. Meer. Heteroscedastic regression in computer vision: problems with bilinear constraint. *Int. J. Computer Vision*, 37(2):127–150, 2000.
- [12] Q.-T. Luong and O. D. Faugeras. The fundamental matrix: theory, algorithms, and stability analysis. *Int. J. Computer Vision*, 17(1):43–75, 1996.
- [13] D. D. Morris and T. Kanade. A unified factorization algorithm for points, line segments and planes with uncertainty models. In *Proc. 6th Int. Conf. Computer Vision, Bombay, India*, pages 696–702, 1998.
- [14] N. Ohta. Image movement detection with reliability indices. *IEICE Trans.*, E74(10):3379–3388, 1991.
- [15] P. H. S. Torr and D. W. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *Int. J. Computer Vision*, 24(3):271–300, 1997.
- [16] B. Triggs. Optimal estimation of matching constraints. In *3D Structure from Multiple Images of Large-Scale Environments, European Workshop, Freiburg, Germany*, volume 1506 of *Lecture Notes in Computer Science*, pages 63–77, 1998.
- [17] Z. Zhang. Parameter estimation techniques: a tutorial with application to conic fitting. *Image & Vision Computing*, 15(1):57–76, 1997.
- [18] Z. Zhang. Determining the epipolar geometry and its uncertainty: a review. *Int. J. Computer Vision*, 27(2):161–195, 1998.