Estimating Vision Parameters Given Data with Covariances

Wojciech Chojnacki, Michael J. Brooks
Anton van den Hengel, Darren Gawley
Department of Computer Science, University of Adelaide
Adelaide, SA 5005, Australia
{wojtek,mjb,hengel,dg}@cs.adelaide.edu.au

Abstract

A new parameter estimation method is presented, applicable to many computer vision problems. It operates under the assumption that the data (typically image point locations) are accompanied by covariance matrices characterising data uncertainty. An MLE-based cost function is first formulated and a new minimisation scheme is then developed. Unlike Sampson's method or the renormalisation technique of Kanatani, the new scheme has as its theoretical limit the true minimum of the cost function. It also has the advantages of being simply expressed, efficient, and unsurpassed in our comparative testing.

1 Introduction

A wide class of computer vision problems may be couched in terms of an equation of the form

$$\boldsymbol{\theta}^T \boldsymbol{u}(\boldsymbol{x}) = 0. \tag{1}$$

Here $\boldsymbol{\theta} = [\theta_1, \ldots, \theta_l]^T$ is a vector representing unknown parameters; $\boldsymbol{x} = [x_1, \ldots, x_k]^T$ is a vector representing the data; and $\boldsymbol{u}(\boldsymbol{x}) = [u_1(\boldsymbol{x}), \ldots, u_l(\boldsymbol{x})]^T$ is a vector with the data transformed in such a way that: (i) each component is a quadratic form in the compound vector $[\boldsymbol{x}^T, 1]^T$, (ii) one component of $\boldsymbol{u}(\boldsymbol{x})$ is equal to 1. An *ancillary constraint* may also apply that does not involve the data, and this can be expressed as $\psi(\boldsymbol{\theta}) = 0$ for some scalar-valued function ψ. We consider the following estimation problem: Given a collection $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ of image data, determine $\boldsymbol{\theta} \neq \mathbf{0}$ satisfying the ancillary equation (if applicable) together with the system of equations obtained from (1) by substituting $\boldsymbol{x}_i$ into $\boldsymbol{x}$ for each $1 \leq i \leq n$. When $n > l$ and noise is present, the corresponding system of equations is overdetermined and as such may fail to have a non-zero solution. In this situation, we are concerned with finding $\boldsymbol{\theta}$ that best fits the data. The form of this vision problem involving covariance information was first studied in detail by Kanatani [9], and later by various others (see, e.g., [3, 10, 13, 18, 19]).

Estimating the coefficients of the *epipolar equation* [5] is one problem of this kind. Consider a camera with an image plane that is equipped with a coordinate system. Any 3D point in the scene perspectively projected onto this plane gives rise to an *image point* represented by a pair (m_1, m_2) of coordinates, or equivalently, by the vector $\boldsymbol{m} =$

$[m_1, m_2, 1]^T$. A 3D point projected onto the image planes of two different cameras endowed with two separate coordinate systems gives rise to a pair of *corresponding points*. When represented by $(\boldsymbol{m}, \boldsymbol{m}')$, this pair satisfies the epipolar equation $\boldsymbol{m}'^T \boldsymbol{F} \boldsymbol{m} = 0$, where $\boldsymbol{F} = [f_{ij}]$ is a 3×3 *fundamental matrix* that incorporates information about the relative orientation and internal geometry of the cameras [5]. The matrix $\boldsymbol{F}$ is subject to the *rank-2 constraint* $\det \boldsymbol{F} = 0$. Let $\boldsymbol{\theta} = [f_{11}, f_{12}, f_{13}, f_{21}, f_{22}, f_{23}, f_{31}, f_{32}, f_{33}]^T$ be the vector of parameters, let $\boldsymbol{x} = [m_1, m_2, m'_1, m'_2]^T$ be the vector of variables, and let

$$\boldsymbol{u}(\boldsymbol{x}) = [m_1 m'_1, m_2 m'_1, m'_1, m_1 m'_2, m_2 m'_2, m'_2, m_1, m_2, 1]^T$$

be the vector of transformed variables. Then $\boldsymbol{m}'^T \boldsymbol{F} \boldsymbol{m} = \boldsymbol{\theta}^T \boldsymbol{u}(\boldsymbol{x})$, showing that the epipolar equation is subsumed by equation (1). The rank-2 constraint can in turn be expressed as the ancillary constraint if we let

$$\psi(\boldsymbol{\theta}) = \theta_1(\theta_5\theta_9 - \theta_6\theta_8) + \theta_2(\theta_6\theta_7 - \theta_4\theta_9) + \theta_3(\theta_4\theta_8 - \theta_5\theta_7).$$

Conic fitting [1,22] and estimating the coefficients of the *differential epipolar equation* [2, 20] are further examples of our general problem (see [3] for details).

2 Cost Functions and Estimators

A vast class of techniques for solving our problem rest upon the use of *cost functions* measuring the extent to which the data and candidate estimates fail to satisfy (1). If—for simplicity—the ancillary constraint is set aside, then, given a cost function $J = J(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$, a corresponding estimate $\widehat{\boldsymbol{\theta}}$ is defined by

$$J(\widehat{\boldsymbol{\theta}}) = \min_{\boldsymbol{\theta} \neq 0} J(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n).$$

Since (1) does not change if $\boldsymbol{\theta}$ is multiplied by a non-zero scalar, we consider only cost functions satisfying $J(t\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = J(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ for any non-zero scalar t. The assignment of $\widehat{\boldsymbol{\theta}}$ (uniquely defined up to a scalar factor) to $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ will be termed the J-based *estimator* of $\boldsymbol{\theta}$.

Once an estimate has been generated by minimising a specific cost function, the ancillary constraint (if it applies) can further be accommodated via an adjustment procedure. One possibility is to use a general scheme delivering an 'optimal correction' described in [9, Subsec. 9.5.2]. In what follows we shall confine our attention to the estimation phase that precedes adjustment.

2.1 Ordinary Least Squares Estimator

A straightforward estimator is derived from the cost function

$$J_{\mathrm{OLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \|\boldsymbol{\theta}\|^{-2} \sum_{i=1}^{n} \boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta},$$

where $\boldsymbol{A}_i = \boldsymbol{u}(\boldsymbol{x}_i)\boldsymbol{u}(\boldsymbol{x}_i)^T$ and $\|\boldsymbol{\theta}\| = (\theta_1^2 + \cdots + \theta_l^2)^{1/2}$. Here each summand $\boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}$ is the square of the *algebraic distance* $|\boldsymbol{\theta}^T \boldsymbol{u}(\boldsymbol{x}_i)|$. Accordingly, the J_{OLS}-based estimate of $\boldsymbol{\theta}$ is termed the *ordinary least squares* (OLS) *estimate* and is denoted $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$. It is uniquely determined, up to a scalar factor, by an eigenvector of $\sum_{i=1}^{n} \boldsymbol{A}_i$ associated with the smallest eigenvalue.

2.2 Weighted Least Squares Estimator

The OLS estimator treats all data as being equally valuable. When information characterising the measurement errors is available, it is desirable that better data be weighted more heavily than poorer data during the estimation process. We assume that the data come equipped with a collection $(\boldsymbol{\Lambda}_{\boldsymbol{x}_1}, \ldots, \boldsymbol{\Lambda}_{\boldsymbol{x}_n})$ of positive definite $k \times k$ *covariance matrices*. These matrices constitute repositories of prior information about the uncertainty of the data.

Adopting a maximum likelihood approach and making some necessary concessions to tractability, a strong case may be mounted for adoption of the weighted least squares cost function given by

$$J_{\mathrm{WLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \sum_{i=1}^{n} \frac{\boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}}{\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta}}, \tag{2}$$

where $\boldsymbol{B}_i = \partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{x}_i) \boldsymbol{\Lambda}_{\boldsymbol{x}_i} \partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{x}_i)^T$ and $\partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{y}) = [(\partial u_i / \partial x_j)(\boldsymbol{y})]_{1 \le i \le l, 1 \le j \le k}$ [3, 9]. The J_{WLS}-based estimate of $\boldsymbol{\theta}$ will be called the *weighted least squares* (WLS) *estimate* and will be denoted $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$.

With the parameters to be estimated in matrix rather than in vector form, the function J_{WLS} underlying fundamental matrix estimation reduces to Zhang's favoured *gradient weighted least squares* cost function (see J_2 in the appendix of [24]) given by

$$J(\boldsymbol{F}) = \sum_{i=1}^{n} \frac{(\boldsymbol{m}_i'^T \boldsymbol{F} \boldsymbol{m}_i)^2}{\boldsymbol{m}_i{}^T \boldsymbol{F} \boldsymbol{\Lambda}_{\boldsymbol{m}_i'} \boldsymbol{F}^T \boldsymbol{m}_i + \boldsymbol{m}_i'^T \boldsymbol{F}^T \boldsymbol{\Lambda}_{\boldsymbol{m}_i} \boldsymbol{F} \boldsymbol{m}_i'}.$$

Here all of the covariance matrices have the form $\boldsymbol{\Lambda}_{\boldsymbol{m}} = \begin{bmatrix} \boldsymbol{\Lambda}_{\boldsymbol{p}} & \boldsymbol{0} \\ \boldsymbol{0}^T & 0 \end{bmatrix}$ with $\boldsymbol{\Lambda}_{\boldsymbol{p}}$ the 2×2 covariance matrix corresponding to the vector $\boldsymbol{p} = [m_1, m_2]^T$ that derives from the representation $\boldsymbol{m} = [m_1, m_2, 1]^T$.

3 Minimising J_{WLS}

Minimising J_{WLS} is a challenging problem. One strategy is to rely upon a general solver such as the Levenberg-Marquardt method [14, Sec. 15.5], the downhill simplex method of Nelder and Mead [14, Sec. 10.4], or one of the direction set methods of Powell [14, Sec. 10.5]. However, such a solver will not take into account the special form of the problem and may be unduly slow. A commonly adopted approach to minimising functions involving fractional expressions is that ascribed to Sampson [15]. When applied to J_{WLS}, Sampson's method (SMP) starts the search for a minimiser by substituting an initial estimate into the denominators $\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta}$ in (2), which transforms our problem into one of minimising a quadratic form in $\boldsymbol{\theta}$; the latter problem is straightforward and reduces to performing singular-value decomposition (SVD) of the matrix defining the form. This process is then repeated with the newly obtained estimate plugged into the denominators until a measure of convergence is obtained.

Kanatani showed that the SMP approach, involving a "freezing" of the denominators, is subject to systematic bias. Accordingly, he devised the technique of renormalisation (REN) in which an attempt is made at each iteration to undo the bias [9, Chap. 12]. We show in a companion paper [3] that the estimate $\widehat{\boldsymbol{\theta}}_{\mathrm{REN}}$ obtained via REN does not theoretically act to compute $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$. However, a case may be made that this is unimportant

given that $\widehat{\boldsymbol{\theta}}_{\mathrm{REN}}$ and $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ are both first-order approximations to $\widehat{\boldsymbol{\theta}}_{\mathrm{ML}}$ and as such are likely to be statistically equivalent.

In this section we derive a straightforward minimisation scheme that is a genuine means of theoretically determining the minimiser of J_{WLS}.

3.1 Variational Equation

As a minimiser of J_{WLS}, $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ satisfies $\partial_{\boldsymbol{\theta}} J_{\mathrm{WLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \mathbf{0}^T$, where $\partial_{\boldsymbol{\theta}} J_{\mathrm{WLS}}$ denotes the row vector of the partial derivatives of J_{WLS} with respect to $\boldsymbol{\theta}$. We term this the *variational equation*. Direct computation shows that

$$[\partial_{\boldsymbol{\theta}} J_{\mathrm{WLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)]^T = 2\boldsymbol{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}, \tag{3}$$

where $\boldsymbol{X}_{\boldsymbol{\theta}}$ is the $l \times l$ symmetric matrix

$$\boldsymbol{X}_{\boldsymbol{\theta}} = \sum_{i=1}^{n} \frac{\boldsymbol{A}_i}{\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta}} - \sum_{i=1}^{n} \frac{\boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}}{(\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta})^2} \boldsymbol{B}_i. \tag{4}$$

Thus the variational equation can be rephrased as

$$\boldsymbol{X}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}. \tag{5}$$

This is a non-linear equation and is unlikely to admit solutions in closed form.

3.2 Fundamental Numerical Scheme

Closed-form solutions of the variational equation may be infeasible, so in practice $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ has to be found numerically. We assume that $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ lies close to $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$ so as to guarantee that, when seeded with $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$, the numerical method that we are going to develop generates an estimate that coincides with $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$. Adopting $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$ as an initial guess $\boldsymbol{\theta}_0$, we construct a sequence $\{\boldsymbol{\theta}_k\}$ of successive updates. Under favourable conditions, the sequence will converge. Exploiting the assumption about the accuracy of the initial guess, we take the corresponding limit for the final estimate. In practice, the limit will be identified with the final term of $\{\boldsymbol{\theta}_k\}$ stopped after a finite number of steps. The stopping rule will be the choice of the first k such that the distance of some kind between $\boldsymbol{\theta}_{k-1}$ and $\boldsymbol{\theta}_k$, say $\|\boldsymbol{\theta}_{k-1} - \boldsymbol{\theta}_k\|$, is less than a pre-assigned quantity.

A straightforward algorithm for numerically solving the variational equation can be derived by realising that a vector $\boldsymbol{\theta}$ satisfies (5) if and only if it falls into the null space of the matrix $\boldsymbol{X}_{\boldsymbol{\theta}}$. Thus if $\boldsymbol{\theta}_{k-1}$ is a tentative guess, then an improved guess can be obtained by picking a vector $\boldsymbol{\theta}_k$ from that eigenspace of $\boldsymbol{X}_{\boldsymbol{\theta}_{k-1}}$ which most closely approximates the null space of $\boldsymbol{X}_{\boldsymbol{\theta}}$; this eigenspace is, of course, the one corresponding to the eigenvalue closest to zero. The *fundamental numerical scheme* implementing this idea is presented in Figure 1. It can be regarded as a variant of the Newton-Raphson method. As it turns out, the scheme is computationally efficient and accurate, and is more direct than alternative approaches (cf. [10]).

4 Experimental Results

The fundamental numerical scheme and other algorithms were tested on the problem of estimating the fundamental matrix [4,6–8,10,11,13,16,17,21–25]. Experiments reported

1. Set $\boldsymbol{\theta}_0 = \widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$.
2. Assuming that $\boldsymbol{\theta}_{k-1}$ is known, compute the matrix $\boldsymbol{X}_{\boldsymbol{\theta}_{k-1}}$.
3. Compute a normalised eigenvector of $\boldsymbol{X}_{\boldsymbol{\theta}_{k-1}}$ corresponding to the eigenvalue closest to zero (in absolute value) and take this eigenvector for $\boldsymbol{\theta}_k$.
4. If $\boldsymbol{\theta}_k$ is sufficiently close to $\boldsymbol{\theta}_{k-1}$, then terminate the procedure; otherwise increment k and return to Step 2.

Figure 1: Fundamental numerical scheme.

here are synthetic as these permit precise control of the conditions under which performance can be evaluated. In our tests, no special knowledge of the domain was utilised. Furthermore, no attempt was made to eliminate outliers, our experiments at this stage comparing general techniques for solving problems of the form specified in (1). The rank-2 constraint was not enforced as a post-process (it turns out that the fundamental numerical scheme always generated fundamental matrices which, when scaled to unit Frobenius norm, had determinants smaller than 10^{-11}).

Our experiments proceeded as follows. 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. A range of tests was then conducted. Each test was carried out with respect to an *average level of noise*, σ, defined as the expected value of the trace of the covariance matrices employed in the test.

For a given test, image points were perturbed by adding inhomogeneous anisotropic noise at some average level σ, consistent with covariance matrices generated via the method described below. Figure 2 depicts a typical image, showing a subset of (unperturbed) image points together with associated ellipses. Each of the ellipses represents a level set of the probability density function describing the noise spread around the ellipse centre, and as such captures graphically the covariances employed.

The following procedure was adopted for generating covariance matrices associated with image points, prescribing anisotropic and inhomogeneous noise at a given average level σ. Given an image point $\boldsymbol{p}$, the scale α of a prospective covariance matrix $\boldsymbol{\Lambda}_{\boldsymbol{p}}$ was first selected from a uniform distribution in the range $[0, 2\sigma]$. (Similar results were obtained using other distributions.) Next, an eccentricity parameter β was generated from a uniform distribution between 0 and 0.5. An intermediate covariance matrix was then formed by setting $\boldsymbol{\Lambda}'_{\boldsymbol{p}} = \alpha \left[\begin{smallmatrix} \beta & 0 \\ 0 & 1-\beta \end{smallmatrix}\right]$. This matrix was then 'rotated' by an angle γ selected from a uniform distribution between 0 and 2π to generate the final covariance $\boldsymbol{\Lambda}_{\boldsymbol{p}} = \boldsymbol{O}_\gamma \boldsymbol{\Lambda}'_{\boldsymbol{p}} {\boldsymbol{O}_\gamma}^T$ with $\boldsymbol{O}_\gamma = \left[\begin{smallmatrix} \cos\gamma & -\sin\gamma \\ \sin\gamma & \cos\gamma \end{smallmatrix}\right]$. Since $\mathsf{E}\left[\mathrm{Tr}\,\boldsymbol{\Lambda}\right] = \mathsf{E}\left[\mathrm{Tr}\,\boldsymbol{\Lambda}'\right] = \mathsf{E}\left[\alpha\right] = \sigma$, where $\mathrm{Tr}\,\boldsymbol{A}$ denotes the trace of the matrix $\boldsymbol{A}$, the noise generated by employing $\boldsymbol{\Lambda}_{\boldsymbol{p}}$ has the average level σ.

The following procedure was adopted for generating covariance matrices associated with points of the form $\boldsymbol{x} = [\boldsymbol{p}^T, \boldsymbol{p}'^T]^T$, where $\boldsymbol{p}$ and $\boldsymbol{p}'$ are corresponding points. First, covariance matrices $\boldsymbol{\Lambda}_{\boldsymbol{p}}$ and $\boldsymbol{\Lambda}_{\boldsymbol{p}'}$ were generated using the recipe given above. Next,

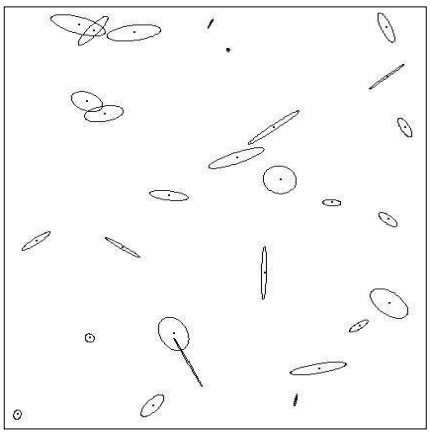

Figure 2: True image points and associated covariance ellipses

covariance matrices $\boldsymbol{\Lambda}_{\boldsymbol{x}}$ were constructed by setting $\boldsymbol{\Lambda}_{\boldsymbol{x}} = \begin{bmatrix} \boldsymbol{\Lambda}_{\boldsymbol{p}} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Lambda}_{\boldsymbol{p}'} \end{bmatrix}$.

Each method under test was then challenged to compute the fundamental matrix. For each σ, the fundamental matrix was computed 250 times from a given set of 60 corresponding (true) points, but with fresh covariance matrices and perturbations generated each time. (Similar results were observed in trials using between 40 and 100 pairs.) For each fundamental matrix obtained, 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, in both the left and right images. A composite error measure was then obtained by averaging this error over all 250 trials. This entire process was then repeated for different average levels of noise (with σ varying from 1 to 10 pixels in steps of 1).

Note that the error measure used in the experiment took advantage of the fact that the underlying true points were known. Were these unknown, an alternative measure might be the sum of the Mahalanobis distances from the data points to the estimated epipolar lines.

The methods tested were as follows:

- **OLS** = ordinary least squares scheme,
- **SMP** = Sampson-like scheme,
- **REN** = renormalisation scheme,
- **LM** = Levenberg-Marquardt scheme,
- **FNS** = fundamental numerical scheme.

Here, SMP is effectively Sampson's method with covariances (see Section 3). The REN scheme chosen was a second-order method; in tests carried out in a companion paper, it performed as well as any of the renormalisation variants (see the SORIII scheme in [3]).

A Levenberg-Marquardt (LM) scheme was included so as to provide a baseline in both accuracy and timing trials; specifically, the MINPACK routine LMDER was used to directly minimise J_{WLS}, with the analytical derivatives of J_{WLS}, as in (3), supplied so as to improve the execution time.

The OLS method uses the LINPACK routine DSVDC to perform SVD, and the EISPACK routine RS is used in those methods requiring computation of eigenvalues and associated eigenvectors, since the matrices involved are symmetric. It should also be noted that the various iterative schemes were supplied with similar stopping conditions so as to enable fair comparison.

Table 1 shows the average epipolar-distance pixel errors obtained for each method. Figure 3 depicts the tabular data in graphical form. We note that errors mount approxi-

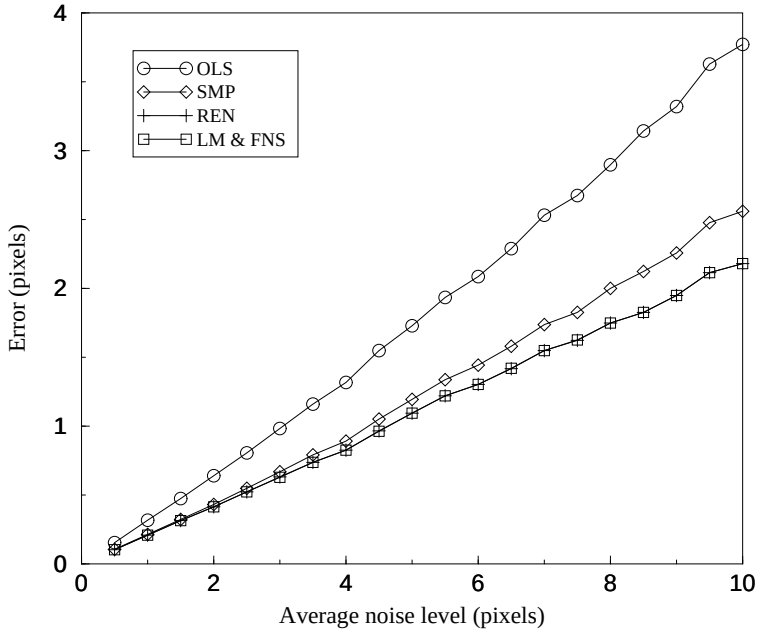

Figure 3: Fundamental matrix estimation errors (in pixels) vs. average noise level

mately linearly with increasing noise. The tests reveal that LM and FNS perform almost identically. In turn, REN generates answers in a very small neighbourhood of those delivered by FNS. Finally, SMP lags behind systematically, while OLS is by far the least successful method.

4.1 Timing Tests

Timing tests were conducted on the various schemes. Stopping conditions were devised so as to place similar demands upon each of the iterative methods. None of the schemes' timings were affected significantly by change in noise level. Figure 4 shows histogram timing data for the OLS, SMP, LM, and FNS methods. In each case, the bar denotes the time taken to complete a single test, averaged over the complete suite of experiments. REN was not included as SORIII is a particularly slow (but accurate) form of renormalisation. FNS typically converged within 4 or 5 iterations. Oscillation was not found to be a practical problem. Interestingly, FNS emerges as being significantly faster than LM, while generating essentially the same results.

AVERAGE	SCHEME				
NOISE LEVEL	OLS	SMP	REN	LM	FNS
1	0.317	0.214	0.210	0.210	0.210
2	0.642	0.432	0.414	0.414	0.414
3	0.982	0.668	0.631	0.630	0.630
4	1.317	0.893	0.826	0.825	0.825
5	1.729	1.195	1.094	1.094	1.094
6	2.085	1.443	1.302	1.303	1.303
7	2.532	1.736	1.548	1.549	1.549
8	2.898	2.001	1.748	1.749	1.749
9	3.319	2.258	1.950	1.949	1.949
10	3.771	2.559	2.181	2.181	2.181

Table 1: Fundamental matrix estimation errors (in pixels) vs. average noise level

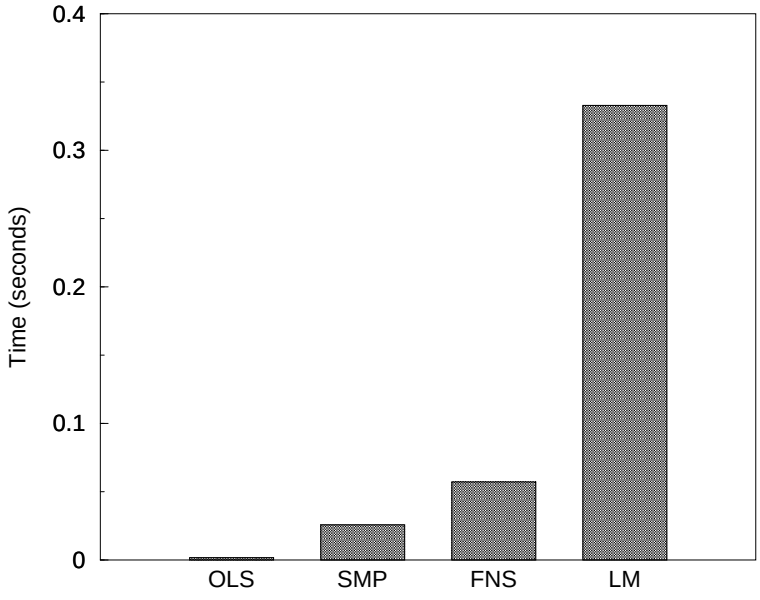

Figure 4: Timing results for various methods

Acknowledgements

The authors are grateful for the insightful comments of Marino Ivancic, Kenichi Kanatani, Garry Newsam, Naoya Ohta and Robyn Owens. 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):56–71, 1979.

[2] M. J. Brooks, W. Chojnacki, and L. Baumela. Determining the egomotion of an uncalibrated camera from instantaneous optical flow. *Journal of the Optical Society of America A*, 14(10):2670–2677, 1997.

[3] W. Chojnacki, M. J. Brooks, and A. van den Hengel. Fitting surfaces to data with covariance information: Fundamental methods applicable to computer vision. Technical Report TR 99-03, Department of Computer Science, University of Adelaide, South Australia, August 1999. Available from `http://www.cs.adelaide.edu.au/~mjb`.

[4] O. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In G. Sandini, editor, *Computer Vision—ECCV '92, Second European Conference on Computer Vision, Santa Margherita Ligure, Italy, May 19–22, 1992*, volume 588 of *Lecture Notes in Computer Science*, pages 563–578, Springer, Berlin, 1992.

[5] O. D. Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint*. The MIT Press, Cambridge, Mass., 1993.

[6] R. Hartley. In defense of the eight-point algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(6):580–593, 1997.

[7] 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, Champaign, IL, June 15–18, 1992*, pages 761–764, IEEE Computer Society Press, Los Alamitos, CA, 1992.

[8] 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):37–50, 1993.

[9] K. Kanatani. *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, Amsterdam, 1996.

[10] Y. Leedan and P. Meer. Heteroscedastic regression in computer vision: problems with bilinear constraint. *International Journal of Computer Vision*, 37(2):127–150, 2000.

[11] Q.-T. Luong and O. D. Faugeras. The fundamental matrix: theory, algorithms, and stability analysis. *International Journal of Computer Vision*, 17(1):43–75, 1996.

[12] B. Matei and P. Meer. Bootstrapping a heteroscedastic regression model with application to 3D rigid motion evaluation. In B. Triggs, A. Zisserman, and R. Szeliski, editors, *Vision Algorithms: Theory and Practice · International Workshop on Vision Algorithms Corfu, Greece, September 1999 Proceedings*, volume 1883 of *Lecture Notes in Computer Science*, pages 236–252, Berlin, 2000. Springer. Available from `http://www.caip.rutgers.edu/riul/`.

[13] M. Mühlich and R. Mester. The role of total least squares in motion analysis. In H. Burkhardt and B. Neumann, editors, *Computer Vision—ECCV'98, Fifth European Conference on Computer Vision, Freiburg, Germany, June 2–6, 1998*, volume 1407 of *Lecture Notes in Computer Science*, pages 305–321, Springer, Berlin, 1998.

[14] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, 1995.

[15] P. D. Sampson. Fitting conic sections to 'very scattered' data: An iterative refinement of the Bookstein algorithm. *Computer Graphics and Image Processing*, 18(1):97–108, 1982.

[16] C. V. Stewart. Robust parameter estimation in computer vision. *SIAM Review*, 41(3):513–537, 1999.

[17] 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):271–300, 1997.

[18] B. Triggs. A new approach to geometric fitting. Available from http://www.inrialpes.fr/movi/people/Triggs.

[19] B. Triggs. Optimal estimation of matching constraints. In R. Koch and L. Van Gool, editors, *3D Structure from Multiple Images of Large-Scale Environments (SMILE), workshop organised by the CUMULI, PANORAMA and VANGUARD projects in conjunction with ECCV'98, Freiburg, Germany, June 6–7, 1998*, pages 30–32, Freiburg University, 1998. Full version available from http://www.inrialpes.fr/movi/people/Triggs.

[20] T. Viéville and O. D. Faugeras. Motion analysis with a camera with unknown, and possibly varying intrinsic parameters. In *Proceedings of the Fifth International Conference on Computer Vision, Cambridge, MA, June 20–23, 1995*, pages 750–756, IEEE Computer Society Press, Los Alamitos, CA, 1995.

[21] 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):451–476, 1989.

[22] Z. Zhang. Parameter estimation techniques: a tutorial with application to conic fitting. *Image and Vision Computing*, 15(1):57–76, 1997.

[23] Z. Zhang. Determining the epipolar geometry and its uncertainty: a review. *International Journal of Computer Vision*, 27(2):161–195, 1998.

[24] Z. Zhang. On the optimization criteria used in two-view motion analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(7):717–729, 1998.

[25] 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*, 78:87–119, 1995.