A new constrained parameter estimator for computer vision applications

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

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

CRC for Sensor Signal and Information Processing, Mawson Lakes, SA 5095, Australia

Abstract

A method of constrained parameter estimation is proposed for a class of computer vision problems. In a typical application, the parameters will describe a relationship between image feature locations, expressed as an equation linking the parameters and the image data, and will satisfy an ancillary constraint not involving the image data. A salient feature of the method is that it handles the ancillary constraint in an integrated fashion, not by means of a correction process operating upon results of unconstrained minimisation. The method is evaluated through experiments in fundamental matrix computation. Results are given for both synthetic and real images. It is demonstrated that the method produces results commensurate with, or superior to, previous approaches, with the advantage of being faster than comparable techniques.

Key words: Gaussian errors, maximum likelihood, constrained minimisation, fundamental matrix, epipolar equation, ancillary constraint, singularity constraint

1 Introduction

Many problems in computer vision involve estimating the parameters that constrain a set of image feature locations. In some cases, the parameters are

* Corresponding author

Email addresses: wojtek@cs.adelaide.edu.au (Wojciech Chojnacki), mjb@cs.adelaide.edu.au (Michael J. Brooks), hengel@cs.adelaide.edu.au (Anton van den Hengel), dg@cs.adelaide.edu.au (Darren Gawley).

subject to an *ancillary constraint* not involving feature locations. Two example problems of this form are the stereo and motion problems of estimating coefficients of the *epipolar equation* [1] and the *differential epipolar equation* [2], each involving an ancillary *cubic constraint*. The *principal equation* that applies to a wide class of problems, including those specified above, takes the form

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

Here $\boldsymbol{\theta} = [\theta_1, \dots, \theta_l]^T$ is a vector representing unknown parameters; $\mathbf{x} = [x_1, \dots, x_k]^T$ is a vector representing an element of the data (for example, the locations of a pair of corresponding points); and $\mathbf{u}(\mathbf{x}) = [u_1(\mathbf{x}), \dots, u_l(\mathbf{x})]^T$ is a vector with the data transformed in a problem-dependent manner such that: (i) each component $u_i(\mathbf{x})$ is a quadratic form in the compound vector $[\mathbf{x}^T, 1]^T$, (ii) one component is equal to 1. A common form of the ancillary constraint is

$$\phi(\boldsymbol{\theta}) = 0, \tag{2}$$

where, for some real number κ, ϕ is a scalar-valued function *homogeneous* of degree κ—that is such that

$$\phi(t\boldsymbol{\theta}) = t^{\kappa}\phi(\boldsymbol{\theta}) \tag{3}$$

for every $\boldsymbol{\theta}$ and every non-zero scalar t. In the above example problems the ancillary constraints are both given by homogeneous functions of degree 3. The estimation problem associated with (1) and (2) can be stated as follows: Given a collection $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ of *image data* and a meaningful *cost function* that characterises the extent to which any particular $\boldsymbol{\theta}$ fails to satisfy the system of the copies of equation (1) associated with $\mathbf{x} = \mathbf{x}_i$ $(i = 1, \dots, n)$, find $\boldsymbol{\theta} \neq \mathbf{0}$ satisfying (2) for which the cost function attains its minimum. Use of the *Gaussian model of errors* in data in conjunction with the *principle of maximum likelihood* leads to a complicated cost function with a tractable first-order approximation defined as

$$J_{\mathrm{AML}}(\boldsymbol{\theta}; \mathbf{x}_1, \dots, \mathbf{x}_n) = \sum_{i=1}^{n} \frac{\boldsymbol{\theta}^T \mathbf{u}(\mathbf{x}_i) \mathbf{u}(\mathbf{x}_i)^T \boldsymbol{\theta}}{\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i)^T \boldsymbol{\theta}},$$

where, for any length k vector $\mathbf{y}$, $\partial_{\mathbf{x}}\mathbf{u}(\mathbf{y})$ denotes the $l \times k$ matrix of the partial derivatives of the function $\mathbf{x} \mapsto \mathbf{u}(\mathbf{x})$ evaluated at $\mathbf{y}$, and, for each $i = 1, \dots, n$, $\boldsymbol{\Lambda}_{\mathbf{x}_i}$ is a $k \times k$ symmetric *covariance matrix* describing the uncertainty of the data point $\mathbf{x}_i$. A full derivation of J_{AML} may be found in [3]; see also [4] for the inception of J_{AML}, and [5] for an instance of its adoption. If J_{AML} is minimised over those non-zero parameter vectors for which (2) holds, then the vector at which the minimum of J_{AML} is attained, the *constrained* minimiser of J_{AML}, defines the *approximated maximum likelihood estimate* $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$. The function $\boldsymbol{\theta} \mapsto J_{\mathrm{AML}}(\boldsymbol{\theta}; \mathbf{x}_1, \ldots, \mathbf{x}_n)$ is homogeneous of degree zero and the zero set of ϕ is unaffected by multiplication by non-zero scalars, so $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$ is determined only up to scale.

The present paper presents a method for determining $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$. Unlike previous approaches, the proposed method handles the ancillary constraint in an integrated fashion, not by means of a correction process operating upon results of unconstrained minimisation. The results of experiments in fundamental matrix estimation given here indicate that the method attains a very high level of performance in terms of accuracy and speed.

2 Fundamental numerical scheme

Isolating the constrained minimiser of J_{AML} is a challenging problem. Much easier to find is the *unconstrained approximated maximum likelihood estimate*, $\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}$, defined as the *unconstrained* minimiser of J_{AML}; it is obtained by ignoring the ancillary constraint and searching over all of the parameter space. While $\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}$ cannot be expressed in closed form, a numerical approximation to it can be calculated by employing a suitable numerical scheme [3]. Underpinning numerical calculation is the fact that $\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}$ satisfies the *variational equation* for unconstrained minimisation

$$[\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}(\boldsymbol{\theta}; \mathbf{x}_1, \ldots, \mathbf{x}_n)]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}} = \mathbf{0}^T \tag{4}$$

with $\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}$ the row vector of the partial derivatives of J_{AML} with respect to $\boldsymbol{\theta}$. Direct computation shows that

$$[\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}(\boldsymbol{\theta}; \mathbf{x}_1, \ldots, \mathbf{x}_n)]^T = 2\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}, \tag{5}$$

where

$$\begin{aligned}
\mathbf{X}_{\boldsymbol{\theta}} &= \sum_{i=1}^{n} \frac{\mathbf{A}_i}{\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta}} - \sum_{i=1}^{n} \frac{\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}}{(\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta})^2} \mathbf{B}_i, \\
\mathbf{A}_i &= \mathbf{u}(\mathbf{x}_i)\mathbf{u}(\mathbf{x}_i)^T, \\
\mathbf{B}_i &= \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i)^T.
\end{aligned}$$

Thus (4) can be written as

$$[\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}} = \mathbf{0}, \tag{6}$$

providing a convenient basis for determining $\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}$. A straightforward algorithm for numerically solving the last equation can be derived by realising that a vector $\boldsymbol{\theta}$ satisfies (6) if and only if it falls into the null space of the matrix $\mathbf{X}_{\boldsymbol{\theta}}$. Thus if $\boldsymbol{\theta}_{k-1}$ is a tentative approximate solution, then an improved solution can be obtained by picking a vector $\boldsymbol{\theta}_k$ from that eigenspace of $\mathbf{X}_{\boldsymbol{\theta}_{k-1}}$ which most closely approximates the null space of $\mathbf{X}_{\boldsymbol{\theta}}$; this eigenspace is, of course, the one corresponding to the eigenvalue closest to zero in absolute value. The *fundamental numerical scheme* (FNS) [3] implementing this idea is presented in Algorithm 1. The scheme is seeded with the *algebraic least squares* (ALS) *estimate*, $\widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$, defined as the unconstrained minimiser of the cost function $J_{\mathrm{ALS}}(\boldsymbol{\theta}; \mathbf{x}_1, \ldots, \mathbf{x}_n) = \|\boldsymbol{\theta}\|^{-2} \sum_{i=1}^n \boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}$, with $\|\boldsymbol{\theta}\| = (\sum_{j=1}^l \theta_j^2)^{1/2}$. The estimate $\widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$ coincides, up to scale, with an eigenvector of $\sum_{i=1}^n \mathbf{A}_i$ associated with the smallest eigenvalue, and this can be found by performing singular-value decomposition (SVD) of the matrix $[\mathbf{u}(\mathbf{x}_1), \ldots, \mathbf{u}(\mathbf{x}_n)]^T$.

Algorithm 1. Fundamental numerical scheme.

1. Set $\boldsymbol{\theta}_0 = \widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$.
2. Assuming $\boldsymbol{\theta}_{k-1}$ is known, compute the matrix $\mathbf{X}_{\boldsymbol{\theta}_{k-1}}$.
3. Compute a normalised eigenvector of $\mathbf{X}_{\boldsymbol{\theta}_{k-1}}$ 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.

Different but related schemes for numerically solving equations similar to (6) were developed by Leedan and Meer [6], and Matei and Meer [7]. Yet another approach is Kanatani's [4, Chap. 9] *renormalisation* scheme, in which an estimate is sought at which $\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}$ is approximately zero (see [8] for details).

FNS has the advantages that it aims to minimise a cost function which is statistically well founded, provides a genuine means for theoretically calculating the minimiser, and is simply expressed and efficient. Furthermore, FNS proves useful in determining theoretical relationships between various existing methods of minimising J_{AML} (see [8]).

3 Constrained fundamental numerical scheme

With all its virtues, FNS still lacks a proper accommodation of the ancillary constraint. A standard way of dealing with this problem is to adopt an adjustment procedure as a separate post-process (see e.g. [4, Chap. 9], [7]). In contrast, as shown next, it is possible to merge FNS and the ancillary constraint in a fully consistent, integrated fashion. The resulting scheme, the first of its kind, is a variant of FNS in which $\mathbf{X}_{\boldsymbol{\theta}}$ is replaced by a more complicated matrix.

The starting point for the development of the new algorithm is the *variational system* for constrained minimisation

$$\begin{aligned} [\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}(\boldsymbol{\theta}) + \lambda \partial_{\boldsymbol{\theta}} \phi(\boldsymbol{\theta})]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}} &= \mathbf{0}^T, \\ \phi(\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}) &= 0, \end{aligned} \tag{7}$$

where λ is scalar. This system is readily derived using the standard method of Lagrange multipliers. Another crucial ingredient is the identity $\partial_{\boldsymbol{\theta}} \phi(\boldsymbol{\theta})\boldsymbol{\theta} = \kappa\phi(\boldsymbol{\theta})$ obtained by differentiating (3) with respect to t and evaluating at $t = 1$. With this identity, the equation $\phi(\boldsymbol{\theta}) = 0$ is equivalent to $\mathbf{a}_{\boldsymbol{\theta}}^T \boldsymbol{\theta} = 0$, where $\mathbf{a}_{\boldsymbol{\theta}} = [\partial_{\boldsymbol{\theta}} \phi(\boldsymbol{\theta})]^T/2$. Combining this with (5), the system (7) becomes

$$\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta} + \lambda \mathbf{a}_{\boldsymbol{\theta}} = \mathbf{0}, \tag{8}$$

$$\mathbf{a}_{\boldsymbol{\theta}}^T \boldsymbol{\theta} = 0, \tag{9}$$

where all evaluations at $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$ are dropped for clarity. Premultiplying both sides of (8) by $\mathbf{a}_{\boldsymbol{\theta}}^T$ and bearing in mind that $\mathbf{a}_{\boldsymbol{\theta}}^T \mathbf{a}_{\boldsymbol{\theta}} = \|\mathbf{a}_{\boldsymbol{\theta}}\|^2$, we find that

$$\mathbf{a}_{\boldsymbol{\theta}}^T \mathbf{X}_{\boldsymbol{\theta}} \boldsymbol{\theta} + \|\mathbf{a}_{\boldsymbol{\theta}}\|^2 \lambda = \mathbf{0}$$

whence $\lambda = -\|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2} \mathbf{a}_{\boldsymbol{\theta}}^T \mathbf{X}_{\boldsymbol{\theta}} \boldsymbol{\theta}$. Consequently, (8) can be written as

$$(\mathbf{X}_{\boldsymbol{\theta}} - \|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2} \mathbf{a}_{\boldsymbol{\theta}} \mathbf{a}_{\boldsymbol{\theta}}^T \mathbf{X}_{\boldsymbol{\theta}})\boldsymbol{\theta} = \mathbf{0}. \tag{10}$$

Let $\mathbf{P}_{\boldsymbol{\theta}}$ be the $l \times l$ matrix given by

$$\mathbf{P}_{\boldsymbol{\theta}} = \mathbf{I}_l - \|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2} \mathbf{a}_{\boldsymbol{\theta}} \mathbf{a}_{\boldsymbol{\theta}}^T,$$

where $\mathbf{I}_l$ denotes the $l \times l$ identity matrix. Note that $\mathbf{P}_{\boldsymbol{\theta}}$ is symmetric and obeys $\mathbf{P}_{\boldsymbol{\theta}}^2 = \mathbf{P}_{\boldsymbol{\theta}}$. With use of $\mathbf{P}_{\boldsymbol{\theta}}$, (10) can be written as

$$\mathbf{P}_{\boldsymbol{\theta}} \mathbf{X}_{\boldsymbol{\theta}} \boldsymbol{\theta} = \mathbf{0}. \tag{11}$$

In view of (9),

$$(\mathbf{I}_l - \mathbf{P}_{\boldsymbol{\theta}})\boldsymbol{\theta} = \|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2}\mathbf{a}_{\boldsymbol{\theta}}\mathbf{a}_{\boldsymbol{\theta}}^T\boldsymbol{\theta} = \mathbf{0}. \tag{12}$$

Hence $\mathbf{P}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \boldsymbol{\theta}$, and so (11) immediately leads to

$$\mathbf{P}_{\boldsymbol{\theta}}\mathbf{X}_{\boldsymbol{\theta}}\mathbf{P}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}. \tag{13}$$

Let $\mathbf{Y}_{\boldsymbol{\theta}}$ be the $l \times l$ matrix defined by

$$\mathbf{Y}_{\boldsymbol{\theta}} = \|\boldsymbol{\theta}\|^2\mathbf{P}_{\boldsymbol{\theta}}\mathbf{X}_{\boldsymbol{\theta}}\mathbf{P}_{\boldsymbol{\theta}} + \mathbf{I}_l - \mathbf{P}_{\boldsymbol{\theta}}.$$

Clearly, $\mathbf{Y}_{\boldsymbol{\theta}}$ is symmetric. Since the function $\boldsymbol{\theta} \mapsto \mathbf{P}_{\boldsymbol{\theta}}$ is homogeneous of degree 0 and the function $\boldsymbol{\theta} \mapsto \mathbf{X}_{\boldsymbol{\theta}}$ is homogeneous of degree -2, it follows that the function $\boldsymbol{\theta} \mapsto \mathbf{Y}_{\boldsymbol{\theta}}$ is consistently homogeneous of degree 0. In view of (12) and (13),

$$\mathbf{Y}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}. \tag{14}$$

As is readily verified, this equation is in fact equivalent to the system comprising (8) and (9).

At this stage, one would hope that a modified FNS with $\mathbf{Y}_{\boldsymbol{\theta}}$ playing the role of $\mathbf{X}_{\boldsymbol{\theta}}$ would be a good tool for calculating $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$. However, experiments reveal that this algorithm fails to converge.

Surprising as it may be, the solution strategy along the lines of FNS is not always successful when applied to equations like (14). This reflects the rather complicated behaviour of the function that sends a symmetric matrix to the eigenspace corresponding to the eigenvalue closest to zero. Since this function is not available in closed form, the question of convergence of FNS-like schemes is subtle and difficult to tackle. If a particular equation, like (14), appears not to be amenable to an FNS solution procedure, it is natural to consider equivalent forms of the equation, hoping that one can find a form for which the FNS approach will work. In the case of (14), one can for example consider equivalent equations of the form $\mathbf{Y}'_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}$, where $\mathbf{Y}'_{\boldsymbol{\theta}} = \mathbf{Y}_{\boldsymbol{\theta}} + \mathbf{N}_{\boldsymbol{\theta}}$ and $\mathbf{N}_{\boldsymbol{\theta}}$ is a symmetric matrix such that $\mathbf{N}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}$ for *all* $\boldsymbol{\theta} \neq \mathbf{0}$ and such that the function $\boldsymbol{\theta} \mapsto \mathbf{N}_{\boldsymbol{\theta}}$ is homogeneous of degree zero. The search for an appropriate modification will often reveal a host of FNS-able forms.

Bearing the above in mind, we now advance a rather complex modification of $\mathbf{Y}_{\boldsymbol{\theta}}$ for which, as experiments indicate, the FNS approach works. The new matrix is a result of search through various options and is derived from a matrix $\mathbf{Z}_{\boldsymbol{\theta}}$ chosen so that (14) is equivalent to

$$\mathbf{Z}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}. \tag{15}$$

To define $\mathbf{Z}_{\boldsymbol{\theta}}$, we need a few preliminary expressions. Denote by $\mathbf{H}_{\boldsymbol{\theta}}$ the Hessian of J_{AML} at $\boldsymbol{\theta}$. Direct if tedious calculation shows that

$$\mathbf{H}_{\boldsymbol{\theta}} = 2(\mathbf{X}_{\boldsymbol{\theta}} - \mathbf{T}_{\boldsymbol{\theta}}),$$

where

$$\mathbf{T}_{\boldsymbol{\theta}} = \sum_{i=1}^{n} \frac{2}{(\boldsymbol{\theta}^T\mathbf{B}_i\boldsymbol{\theta})^2}\Big[\mathbf{A}_i\boldsymbol{\theta}\boldsymbol{\theta}^T\mathbf{B}_i + \mathbf{B}_i\boldsymbol{\theta}\boldsymbol{\theta}^T\mathbf{A}_i - 2\frac{\boldsymbol{\theta}^T\mathbf{A}_i\boldsymbol{\theta}}{\boldsymbol{\theta}^T\mathbf{B}_i\boldsymbol{\theta}}\mathbf{B}_i\boldsymbol{\theta}\boldsymbol{\theta}^T\mathbf{B}_i\Big].$$

Let $\boldsymbol{\Phi}_{\boldsymbol{\theta}}$ be the Hessian of ϕ at $\boldsymbol{\theta}$. For each $i \in \{1, \ldots, l\}$, let $\mathbf{e}_i$ be the length l vector whose ith entry is unital and all other entries are zero. With these preparations, set

$$\mathbf{Z}_{\boldsymbol{\theta}} = \mathbf{A}_{\boldsymbol{\theta}} + \mathbf{B}_{\boldsymbol{\theta}} + \mathbf{C}_{\boldsymbol{\theta}},$$

where

$$\begin{aligned}
\mathbf{A}_{\boldsymbol{\theta}} &= \mathbf{P}_{\boldsymbol{\theta}}\mathbf{H}_{\boldsymbol{\theta}}(2\boldsymbol{\theta}\boldsymbol{\theta}^T - \|\boldsymbol{\theta}\|^2\mathbf{I}_l),\\
\mathbf{B}_{\boldsymbol{\theta}} &= \|\boldsymbol{\theta}\|^2\|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2}\Big[\sum_{i=1}^{l}(\boldsymbol{\Phi}_{\boldsymbol{\theta}}\mathbf{e}_i\mathbf{a}_{\boldsymbol{\theta}}^T + \mathbf{a}_{\boldsymbol{\theta}}\mathbf{e}_i^T\boldsymbol{\Phi}_{\boldsymbol{\theta}})\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}\mathbf{e}_i^T - 2\|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2}\mathbf{a}_{\boldsymbol{\theta}}\mathbf{a}_{\boldsymbol{\theta}}^T\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}\mathbf{a}_{\boldsymbol{\theta}}^T\boldsymbol{\Phi}_{\boldsymbol{\theta}}\Big],\\
\mathbf{C}_{\boldsymbol{\theta}} &= \|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2}\kappa\Big[\frac{\phi(\boldsymbol{\theta})}{4}\boldsymbol{\Phi}_{\boldsymbol{\theta}} + \mathbf{a}_{\boldsymbol{\theta}}\mathbf{a}_{\boldsymbol{\theta}}^T - \frac{\phi(\boldsymbol{\theta})}{2}\|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2}\mathbf{a}_{\boldsymbol{\theta}}\mathbf{a}_{\boldsymbol{\theta}}^T\boldsymbol{\Phi}_{\boldsymbol{\theta}}\Big].
\end{aligned}$$

It should be emphasised that individually the matrices $\mathbf{A}_{\boldsymbol{\theta}}$, $\mathbf{B}_{\boldsymbol{\theta}}$ and $\mathbf{C}_{\boldsymbol{\theta}}$ do not have any special significance, and serve only to split the otherwise lengthy formula. It should also be stressed that $\mathbf{Z}_{\boldsymbol{\theta}}$ arises as a result of algebraic manipulations that have no obvious geometric interpretation. Direct calculation reveals that $\mathbf{A}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \|\boldsymbol{\theta}\|^2\mathbf{P}_{\boldsymbol{\theta}}\mathbf{H}_{\boldsymbol{\theta}}\boldsymbol{\theta}$, $\mathbf{B}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}$ and $\mathbf{C}_{\boldsymbol{\theta}}\boldsymbol{\theta} = (\mathbf{I}_l - \mathbf{P}_{\boldsymbol{\theta}})\boldsymbol{\theta}$ for each $\boldsymbol{\theta}$. Therefore (15) is equivalent to

$$(\|\boldsymbol{\theta}\|^2\mathbf{P}_{\boldsymbol{\theta}}\mathbf{H}_{\boldsymbol{\theta}} + \mathbf{I}_l - \mathbf{P}_{\boldsymbol{\theta}})\boldsymbol{\theta} = 0,$$

which is in turn equivalent to (14), as a simple argument shows.

Note that $\mathbf{Z}_{\boldsymbol{\theta}}$ is *not* symmetric. On the other hand, $\mathbf{Z}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}$ is equivalent to

$$\mathbf{Z}_{\boldsymbol{\theta}}^T\mathbf{Z}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0} \tag{16}$$

with $\mathbf{Z}_{\boldsymbol{\theta}}^T \mathbf{Z}_{\boldsymbol{\theta}}$ a symmetric matrix. This allows

$$\mathbf{Q}_{\boldsymbol{\theta}} = \mathbf{Z}_{\boldsymbol{\theta}}^T \mathbf{Z}_{\boldsymbol{\theta}}$$

to be adopted as the ultimate replacement for $\mathbf{X}_{\boldsymbol{\theta}}$. The steps of the resulting *constrained fundamental numerical scheme* (CFNS) are given in Algorithm 2.

Algorithm 2. Constrained fundamental numerical scheme.

1. Set $\boldsymbol{\theta}_0 = \widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$.
2. Assuming that $\boldsymbol{\theta}_{k-1}$ is known, compute the matrix $\mathbf{Q}_{\boldsymbol{\theta}_{k-1}}$.
3. Compute a normalised eigenvector of $\mathbf{Q}_{\boldsymbol{\theta}_{k-1}}$ 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.

A necessary condition for CFNS to converge to a solution $\boldsymbol{\theta}^*$ of (16) is that the zero eigenvalue of $\mathbf{Q}_{\boldsymbol{\theta}^*}$ should be simple, i.e., the null space of $\mathbf{Q}_{\boldsymbol{\theta}^*}$ should be one-dimensional, with all members being scalar multiples of $\boldsymbol{\theta}^*$. When this condition is satisfied, the algorithm seeded with an estimate close enough to $\boldsymbol{\theta}^*$ will produce updates quickly converging to $\boldsymbol{\theta}^*$. In practice it is required that, for each $k = 0, 1, \ldots$, the smallest (non-negative) eigenvalue of $\mathbf{Q}_{\boldsymbol{\theta}^*}$ should be sufficiently well separated from the remaining eigenvalues. Sometimes, to meet the condition, the data will have to be first suitably transformed and their covariances propagated; upon application of CNFS, the estimate will then have to be conformally readjusted (transformed back) to account for the data-cum-covariances transformation. Such is the case for fundamental matrix estimation, where an initial transformation of raw data and their covariances is necessary for a successful application of CFNS (this point will be elaborated upon in the next section).

4 Experimental evaluation

In this section, we present results of comparative tests carried out to evaluate the performance of CFNS. Several algorithms, including CFNS, were used to compute the fundamental matrix from both synthetic and real image data. A single item of data took the form of a quadruple obtained by concatenating the coordinates of a pair of corresponding points, the role of the principal constraint was played by the epipolar constraint, and the ancillary constraint was the condition that the determinant of the fundamental matrix should vanish. The covariances of the data were assumed to be the default 4×4 identity matrix corresponding to isotropic homogeneous noise in image point measurement.

The basic estimation methods considered were:

- **NALS** = Normalised Algebraic Least Squares Method,
- **FNS** = Fundamental Numerical Scheme,
- **CFNS** = Constrained FNS,
- **GS** = Gold Standard Method.

Here, NALS refers to the *normalised* ALS method of Hartley [9], which takes suitably normalised data as input to ALS and back-transforms the resulting estimate; GS refers to the (theoretically optimal) bundle-adjustment, maximum-likelihood method described by Hartley and Zisserman [10], seeded with the FNS estimate; FNS and CFNS are as described earlier. CFNS was applied in the Hartley-normalised data domain, with covariances assumed again to be the 4×4 identity matrix (covariance propagation for data normalisation amounts simply to multiplication of the original covariances by a single scalar factor, and this transformation does not affect the determination of the constrained minimiser). The data normalisation combined with back-transforming of estimates has no theoretical influence on the constrained minimiser, but in practice significantly improves separation of the smaller eigenvalues of the matrices $\mathbf{Q}_{\boldsymbol{\theta}}$ involved. The CFNS algorithm usually fails to converge when used with raw data, a phenomenon explained by the lack of sufficient eigenvalue separation.

When comparing the outputs of algorithms, it is critical that the ancillary constraint be perfectly satisfied. A convenient way to enforce this constraint is to correct an estimate of the fundamental matrix via a post-process. Any estimate $\widehat{\mathbf{F}}$ with $\|\widehat{\mathbf{F}}\|_F = 1$ can be modified to a rank-2 matrix $\widehat{\mathbf{F}}_c$ with $\|\widehat{\mathbf{F}}_c\|_F = 1$ by minimising the distance $\|\widehat{\mathbf{F}} - \widehat{\mathbf{F}}_c\|_F$ subject to the condition $\det \widehat{\mathbf{F}}_c = 0$; here $\|\cdot\|_F$ denotes the Frobenius norm. The minimiser can easily be found by performing a SVD of $\widehat{\mathbf{F}}$, setting the smallest singular value to zero and recomposing. For the estimate generated by FNS, a more sophisticated, Kanatani-like (cf. [4, Chap. 5], [7]) correction can be obtained by means of the iterative process

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - [\partial_{\boldsymbol{\theta}}\phi(\boldsymbol{\theta}_k)\mathbf{H}^-_{\boldsymbol{\theta}_k}[\partial_{\boldsymbol{\theta}}\phi(\boldsymbol{\theta}_k)]^T]^{-1}\phi(\boldsymbol{\theta}_k)\mathbf{H}^-_{\boldsymbol{\theta}_k}[\partial_{\boldsymbol{\theta}}\phi(\boldsymbol{\theta}_k)]^T, \tag{17}$$

where $\mathbf{H}^-_{\boldsymbol{\theta}_k}$ denotes the pseudo-inverse of $\mathbf{H}_{\boldsymbol{\theta}_k}$.

All estimates were post-hoc rank-2 corrected except in the case of the NALS estimates, where SVD correction preceded the final back-transformation of estimates. In the following, we use the notation "+" to denote a post-process SVD correction, and "++" to denote an iterative correction (see (17)) followed by SVD correction. Thus, the composition of FNS and SVD correction is denoted by FNS+. Of the various methods listed here, only SVD correction is guaranteed to generate a perfectly rank-2 estimate, although CFNS, GS and

Table 1
J_{AML} and $|\phi|$ values for FNS and CFNS before and after SVD rank-2 correction.

	J_{AML}	$\lvert\phi\rvert$
FNS	50.18	1.56×10^{-13}
FNS+	57.48	0
CFNS	52.62	3.07×10^{-25}
CFNS+	52.62	0

the iterative correction usually get extremely close.

For reasons of clarity, testing of the Leedan–Meer and Matei–Meer methods and Kanatani's renormalisation technique was not included in the results presented here. As reported previously [3,11], these unconstrained methods deliver results almost identical to those of FNS, and here we are primarily concerned with the improvement over these methods stemming from the use of our new integrated, constrained approach.

4.1 Synthetic image tests

While perhaps under-appreciated in computer vision, synthetic tests are valuable, as we have ground truth available and may employ repeated trials yielding results of statistical significance.

The regime adopted was to generate true corresponding points for some stereo configuration and collect performance statistics over many trials in which random Gaussian perturbations were made to the image points. Many configurations were investigated and the results below are typical. Specifically, we conducted experiments by first choosing a realistic geometric configuration for the cameras. Next, 30 3D points were randomly selected in the field of view of both cameras, and these were then projected onto 500×500 pixel images to provide "true" matches. For each of 200 iterations, homogeneous Gaussian noise with standard deviation of 1.5 pixels was added to each image point, and the contaminated pairs were used as input to the various algorithms.

Table 1 examines the FNS and CFNS methods in terms of the cost function, J_{AML}, and the ancillary constraint residual, $|\phi|$. The values displayed are the averages of individual values obtained in all 200 trials. As is to be expected, and consistent with its design, FNS generates the smallest value of J_{AML}, but leaves a non-zero ancillary constraint value, $|\phi|$. CFNS reduces the value of $|\phi|$ almost to zero and (necessarily) incurs a small increase in J_{AML}. Note that an SVD correction (which ensures $\phi = 0$) of the FNS estimate results in an associated J_{AML} value that is substantially increased. In contrast, SVD correction of the CFNS estimate leaves the J_{AML} value virtually unaffected,

Table 2
J_{AML} residuals, reprojection errors and execution times for rank-2 estimates.

	J_{AML}	Reproj. error	Time
NALS+	57.50	1.278	0.02
FNS+	57.47	1.278	0.29
FNS++	53.42	1.265	0.61
CFNS+	52.62	1.263	0.23
GS+	52.62	1.263	3.50

and much smaller than the corrected FNS estimate. This test, which is typical, confirms that CFNS is operating as designed.

Table 2 compares the J_{AML} values generated by the methods NALS+, FNS+, FNS++, CFNS+, and GS+. Note that all of the methods undergo a final SVD rank-2 correction ensuring that the ancillary constraint is perfectly satisfied. Were we to avoid this step (in, say, the CFNS and GS approaches) it might be unclear whether a low J_{AML} value was due to the constraint not having been exactly satisfied.

The results show that, with respect to J_{AML}, GS+ and CFNS+ perform best and equally well, with FNS++ only a little behind; FNS+ and NALS+ are set further back. The same ordering occurs when using a measure in which the estimated fundamental matrix is employed to reproject the data and compute the distance of the data from the truth. This reprojection error from truth may be regarded as an optimal measure in the synthetic realm.

Finally, a timing test is also presented in Table 2. Here we give the average time over 100 trials to compute NALS, FNS, CFNS, and GS. Unsurprisingly, GS turns out to be by far the slowest of the methods. While it may be speeded up via the incorporation of sparse-matrix techniques, it is destined to be relatively slow given the high-dimensionality of its search strategy.

CFNS thus emerges as an excellent means of estimating the fundamental matrix. Its performance is commensurate with GS while being much faster. FNS++ is only a little short of CFNS in speed and accuracy. However, it does not have the advantage of being an *integrated* method of constrained minimisation.

4.2 Real image tests

The image pairs from which we estimate fundamental matrices are presented in Figure 1. They exhibit variation in subject matter, and in the camera set-up used for acquisition. Features were detected in each image using the

Fig. 1. The building and soccer ball stereo image pairs.

Table 3
J_{AML} residuals and reprojection errors for rank-2 estimates - soccer ball images.

	J_{AML}	Reproj. error (to data)
NALS+	0.799	0.0926
FNS+	0.813	0.0933
FNS++	0.442	0.0681
CFNS+	0.442	0.0681
GS+	0.442	0.0681

Table 4
J_{AML} residuals and reprojection errors for rank-2 estimates - building images.

	J_{AML}	Reproj. error (to data)
NALS+	5.35	0.285
FNS+	4.88	0.275
FNS++	2.05	0.173
CFNS+	1.88	0.163
GS+	1.88	0.163

Harris corner detector [12]. A set of corresponding points was generated for each image pair by manually matching the detected features. The number of matched points was 44 for the building, and 55 for the soccer ball. For each estimation method, the entire set of matched points was used to compute a fundamental matrix.

Each estimator was used to generate a fundamental matrix. Tables 3 and 4 show results obtained for various methods when dealing with the soccer-ball and building images, respectively. Measures used for comparison are J_{AML} and the reprojection error to data (the distance between the reprojected data and the original data). Note that the ancillary constraint is in all cases perfectly satisfied. CFNS+ and GS+ give the best results and are essentially inseparable, while FNS++ is only slightly behind. FNS+ and NALS+ lag much further behind.

5 Conclusion

The J_{AML} cost function has been shown in many previous studies to be valuable in obtaining a first-order approximation to a maximum likelihood estimate. Earlier work of the authors showed that FNS is a theoretically well-founded and efficient method for finding the minimiser of J_{AML}, given an initial estimate in the neighbourhood of the global minimum. However, our experiments with fundamental matrix estimation show that post-process SVD rank-2 correction of J_{AML}'s minimiser can act to considerably worsen the J_{AML} residual.

CFNS is a variant of the FNS method which incorporates constraint in an integrated manner, aiming to solve an equation characterising a constrained minimiser of J_{AML}. Our experiments suggest that CFNS achieves a smaller value of J_{AML} than is attained by either SVD rank-2 corrected FNS or SVD rank-2 corrected ALS. It produces a slightly smaller value of J_{AML} than the slower, iteratively rank-2 corrected FNS, and has the advantage of being a conceptually cogent, integrated method for constrained minimisation. Furthermore, when compared with the much slower FNS-seeded GS, it gives almost identical results, both in terms of J_{AML} residual and GS's MLE cost function residual.

Acknowledgements

This work was supported by the Australian Research Council. Thanks go to Ian Reid, David Murray, and the University of Oxford for hosting several of the authors. Anton van den Hengel extends his thanks to Luc van Gool and Marc Pollefeys for hosting him at the Katholieke Universiteit Leuven. The comments of Naoya Ohta and Kenichi Kanatani were greatly appreciated.

References

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

[2] M. J. Brooks, W. Chojnacki, L. Baumela, Determining the egomotion of an uncalibrated camera from instantaneous optical flow, J. Opt. Soc. Amer. A 14 (10) (1997) 2670–2677.

[3] W. Chojnacki, M. J. Brooks, A. van den Hengel, D. Gawley, On the fitting of surfaces to data with covariances, IEEE Trans. Pattern Anal. Mach. Intell. 22 (11) (2000) 1294–1303.

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

[5] M. J. Brooks, W. Chojnacki, D. Gawley, A. van den Hengel, What value covariance information in estimating vision parameters?, in: Proceedings, Eighth IEEE International Conference of Computer Vision, Vancouver, British Columbia, Canada, July 7–14, 2001, Vol. 1, IEEE Computer Society, Los Alamitos, CA, 2001, pp. 302–308.

[6] Y. Leedan, P. Meer, Heteroscedastic regression in computer vision: problems with bilinear constraint, Int. J. Comput. Vision 37 (2) (2000) 127–150.

[7] B. Matei, 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, Hilton Head Island, South Carolina, June 13-15, 2000, Vol. 2, IEEE Computer Society Press, Los Alamitos, CA, 2000, pp. 18–25.

[8] W. Chojnacki, M. J. Brooks, A. van den Hengel, Rationalising the renormalisation method of Kanatani, J. Math. Imaging Vision 14 (1) (2001) 21–38.

[9] R. Hartley, In defense of the eight-point algorithm, IEEE Trans. Pattern Anal. Mach. Intell. 19 (6) (1997) 580–593.

[10] R. Hartley, A. Zisserman, Multiple View Geometry in Computer Vision, Cambridge University Press, 2000.

[11] W. Chojnacki, M. J. Brooks, A. van den Hengel, D. Gawley, From FNS to HEIV: A link between two vision parameter estimation methods, IEEE Trans. Pattern Anal. Mach. Intell., to appear.

[12] C. G. Harris, M. Stephens, A combined corner and edge detector, in: C. J. Taylor (Ed.), Proceedings of the Fourth Alvey Vision Conference, Manchester, UK, August 31–September 2, 1988, 1988, pp. 147–151.