A Statistical Rationalisation of Hartley's Normalised Eight-Point Algorithm

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
{wojtek,mjb,hengel,dg}@cs.adelaide.edu.au

Abstract

The eight-point algorithm of Hartley occupies an important place in computer vision, notably as a means of providing an initial value of the fundamental matrix for use in iterative estimation methods. In this paper, a novel explanation is given for the improvement in performance of the eight-point algorithm that results from using normalised data. A first step is singling out a cost function that the normalised algorithm acts to minimise. The cost function is then shown to be statistically better founded than the cost function associated with the non-normalised algorithm. This augments the original argument that improved performance is due to the better conditioning of a pivotal matrix. Experimental results are given that support the adopted approach. This work continues a wider effort to place a variety of estimation techniques within a coherent framework.

1. Introduction

In a landmark paper, Longuet-Higgins [12] proposed the *eight-point algorithm*—a simple direct method for computation of the *essential matrix*. The algorithm extends straightforwardly to computation of the *fundamental matrix*, the uncalibrated analogue of the essential matrix [5,9]. While simple and fast, the algorithm is very sensitive to noise in the specification of the image coordinates serving as input for computation, and as such is of limited use. Many alternative methods have been advanced since Longuet-Higgins' proposal, including more sophisticated and computationally intensive iterative algorithms [6, 8]. Hartley [7] discovered that the accuracy of the eight-point algorithm can be greatly improved if, prior to applying the method, a simple normalisation of image data is performed. This fundamental modification dramatically extended the applicability of the algorithm, and, in particular, rendered it an excellent tool for generation of initial estimates for iterative methods.

Hartley attributed the improved performance of the normalised eight-point algorithm to the better numerical conditioning of a pivotal matrix used in solving an eigenvalue problem. Here we offer a new insight into the working of the method. A crucial observation is that the estimate produced by the normalised eight-point algorithm can be identified with the minimiser of a cost function. The minimiser can be directly calculated by solving a generalised eigenvalue problem. We confirm experimentally that the estimate obtained as a solution of the generalised eigenproblem coincides with the estimate generated by Hartley's original method. Exploiting the cost function, we propose an alternative explanation of the improved performance of the normalised eight-point algorithm, based on a certain statistical model of data distribution. Under this model, the summands of the cost function underlying the normalised eight-point algorithm turn out to be more balanced in terms of spread than the summands of the cost function underlying the standard eight-point algorithm. Summation of more balanced terms leads to a statistically more appropriate expression for minimisation, and this in turn translates into a more accurate estimator. The proposed approach continues a line of research due to Torr [14], Mühlich and Mester [13], and Torr and Fitzgibbon [15], in which variants of the normalised eight-point algorithm are analysed statistically. The work presented here also forms part of a wider effort to place a variety of estimation techniques within a coherent framework (e.g. see [1, 2, 4, 10, 11]).

2. Estimation Problem

A 3D point in a scene perspectively projected onto the image plane of a camera gives rise to an *image point* represented by a pair (m_1, m_2) of coordinates, or equivalently, by the 'homogeneous' vector $\boldsymbol{m} = [m_1, m_2, 1]^T$. A 3D point projected onto the image planes of two cameras endowed with separate coordinate systems gives rise to a pair of *corresponding points*. When represented by $(\boldsymbol{m}, \boldsymbol{m}')$, this pair satisfies the *epipolar constraint*

$$\boldsymbol{m}'^T \boldsymbol{F} \boldsymbol{m} = 0, \tag{1}$$

where $\boldsymbol{F} = [f_{ij}]$ is a 3×3 fundamental matrix that incorporates information about the relative orientation and internal geometry of the cameras [6, 8]. In addition to (1), $\boldsymbol{F}$ is subject to the *singularity constraint* (or, equivalently, the *rank-2 constraint*)

$$\det \boldsymbol{F} = 0. \tag{2}$$

Using $\boldsymbol{x} = [m_1, m_2, m'_1, m'_2]^T$ as a compact descriptor of the single *image datum* $(\boldsymbol{m}, \boldsymbol{m}')$, the estimation problem associated with (1) and (2) can be stated as follows: Given a collection $\{\boldsymbol{x}_1, \dots, \boldsymbol{x}_n\}$ of image data and a meaningful *cost function* that characterises the extent to which any particular $\boldsymbol{F}$ fails to satisfy the system of the copies of equation (1) associated with $\boldsymbol{x} = \boldsymbol{x}_i$ ($i = 1, \dots, n$), find an *estimate* $\widehat{\boldsymbol{F}} \neq \boldsymbol{0}$ satisfying (2) for which the cost function attains its minimum. Since (1) and (2) do not change when $\boldsymbol{F}$ is multiplied by a non-zero scalar, $\widehat{\boldsymbol{F}}$ is to be found only up to scale. If the singularity constraint is set aside, then the estimate associated with a particular cost function $J = J(\boldsymbol{F}; \boldsymbol{x}_1, \dots, \boldsymbol{x}_n)$ is defined as the *unconstrained* minimiser $\widehat{\boldsymbol{F}}$ of J,

$$\widehat{\boldsymbol{F}} = \arg\min_{\boldsymbol{F} \neq \boldsymbol{0}} J(\boldsymbol{F}; \boldsymbol{x}_1, \dots, \boldsymbol{x}_n).$$

Such an estimate can further be converted to a nearby rank-2 fundamental matrix by applying one of a variety of methods [8, 10]. In this paper, we shall confine our attention to the pivotal component of this overall process that determines exclusively the unconstrained minimiser, as this will prove critical to rationalising the Hartley method. For alternative *integrated* approaches to computing a constrained minimiser, see the CFNS method [3, 16] or the Gold Standard Method [8].

3. Algebraic Least Squares

A straightforward estimation method employs the cost function

$$J_{\mathrm{ALS}}(\boldsymbol{F}; \boldsymbol{x}_1, \dots, \boldsymbol{x}_n) = \|\boldsymbol{F}\|_F^{-2} \sum_{i=1}^{n} (\boldsymbol{m}'^T_i \boldsymbol{F} \boldsymbol{m}_i)^2$$

with $\|\boldsymbol{F}\|_F = (\sum_{i,j} f_{ij}^2)^{1/2}$ the Frobenius norm of $\boldsymbol{F}$. Here $\boldsymbol{m}'^T_i \boldsymbol{F} \boldsymbol{m}_i$ is the signed *algebraic distance* between the individual datum $\boldsymbol{x}_i$ and the candidate matrix $\boldsymbol{F}$. The *algebraic least squares* (ALS) *estimate*, $\widehat{\boldsymbol{F}}_{\mathrm{ALS}}$, is defined as the minimiser of J_{ALS}.

A practical means for finding $\widehat{\boldsymbol{F}}_{\mathrm{ALS}}$ is conveniently derived based on an alternative expression for $\boldsymbol{m}'^T \boldsymbol{F} \boldsymbol{m}$. Given a matrix $\boldsymbol{A}$, denote by $\mathrm{vec}(\boldsymbol{A})$ the *vectorisation* of $\boldsymbol{A}$, that is the vector obtained by stacking the columns of $\boldsymbol{A}$ on top of each other. Let $\boldsymbol{\theta} = \mathrm{vec}(\boldsymbol{F}^T)$ and $\boldsymbol{u}(\boldsymbol{x}) = \mathrm{vec}(\boldsymbol{m}\boldsymbol{m}'^T)$. Then, as is easily verified,

$$\boldsymbol{m}'^T \boldsymbol{F} \boldsymbol{m} = \boldsymbol{\theta}^T \boldsymbol{u}(\boldsymbol{x}).$$

With this formula, J_{ALS} can be written as

$$J_{\mathrm{ALS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \dots, \boldsymbol{x}_n) = \|\boldsymbol{\theta}\|^{-2} \boldsymbol{\theta}^T \boldsymbol{A} \boldsymbol{\theta}, \tag{3}$$

where $\boldsymbol{A} = \sum_{i=1}^{n} \boldsymbol{u}(\boldsymbol{x}_i)\boldsymbol{u}(\boldsymbol{x}_i)^T$ and $\|\boldsymbol{\theta}\| = (\theta_1^2 + \cdots + \theta_9^2)^{1/2}$. Now, using (3) to evolve a variational equation for the minimiser, $\widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$ can be characterised as an eigenvector of $\boldsymbol{A}$ associated with the smallest eigenvalue [2]. This eigenvector can be found in practice by performing *singular value decomposition* (SVD) of the matrix

$$\boldsymbol{M} = [\boldsymbol{u}(\boldsymbol{x}_1), \dots, \boldsymbol{u}(\boldsymbol{x}_n)]^T \tag{4}$$

and taking for the desired output the right singular vector of $\boldsymbol{M}$ associated with the smallest singular value (the minimum right singular vector). In this form, the ALS estimator is essentially identical to the eight-point algorithm of Longuet-Higgins [12].

4. Hartley's Approach

Let $\overline{\boldsymbol{m}}$ and $\overline{\boldsymbol{m}}'$ be the *centroids*, or 'centres of mass', of the $\boldsymbol{m}_i$ and the $\boldsymbol{m}'_i$, respectively, defined by

$$\overline{\boldsymbol{m}} = \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{m}_i \quad \text{and} \quad \overline{\boldsymbol{m}}' = \frac{1}{n} \sum_{i=1}^{n} \boldsymbol{m}'_i. \tag{5}$$

Let $\overline{\boldsymbol{m}} = [\overline{m}_1, \overline{m}_2, 1]^T$, $\overline{\boldsymbol{m}}' = [\overline{m}'_1, \overline{m}'_2, 1]^T$, $\boldsymbol{m}_i = [m_{1,i}, m_{2,i}, 1]^T$, $\boldsymbol{m}'_i = [m'_{1,i}, m'_{2,i}, 1]^T$ ($i = 1, \dots, n$). Following Hartley [7], let us shift the image coordinate systems to the respective centroids. In coordinates associated with the transformed systems, the points of the ith image datum can be written $[m_{1,i} - \overline{m}_1, m_{2,i} - \overline{m}_2, 1]^T$ and $[m'_{1,i} - \overline{m}'_1, m'_{2,i} - \overline{m}'_2, 1]^T$. Let

$$\begin{aligned} s &= \left(\frac{1}{2n} \sum_{i=1}^{n} (m_{1,i} - \overline{m}_1)^2 + (m_{2,i} - \overline{m}_2)^2 \right)^{1/2}, \\ s' &= \left(\frac{1}{2n} \sum_{i=1}^{n} (m'_{1,i} - \overline{m}'_1)^2 + (m'_{2,i} - \overline{m}'_2)^2 \right)^{1/2}. \end{aligned} \tag{6}$$

Define the *normalised* data as

$$\tilde{\boldsymbol{m}}_i = [(m_{1,i} - \overline{m}_1)/s, (m_{2,i} - \overline{m}_2)/s, 1]^T,$$
$$\tilde{\boldsymbol{m}}'_i = [(m'_{1,i} - \overline{m}'_1)/s', (m'_{2,i} - \overline{m}'_2)/s', 1]^T.$$

This amounts to setting $\tilde{\boldsymbol{m}}_i = \boldsymbol{T}\boldsymbol{m}_i$ and $\tilde{\boldsymbol{m}}'_i = \boldsymbol{T}'\boldsymbol{m}'_i$, where

$$\boldsymbol{T} = \begin{bmatrix} s^{-1} & 0 & -s^{-1}\overline{m}_1 \\ 0 & s^{-1} & -s^{-1}\overline{m}_2 \\ 0 & 0 & 1 \end{bmatrix},$$

$$\boldsymbol{T}' = \begin{bmatrix} s'^{-1} & 0 & -s'^{-1}\overline{m}'_1 \\ 0 & s'^{-1} & -s'^{-1}\overline{m}'_2 \\ 0 & 0 & 1 \end{bmatrix}.$$

Let $\tilde{x}_i = [\tilde{m}_{1,i}, \tilde{m}_{2,i}, \tilde{m}'_{1,i}, \tilde{m}'_{2,i}]^T$. Denote by $\widehat{\tilde{F}}_{\mathrm{ALS}}$ the minimiser of the ALS cost function seeded with the normalised data, that is, the minimiser of the function $F \mapsto J_{\mathrm{ALS}}(F; \tilde{x}_1, \ldots, \tilde{x}_n)$. Let $F \mapsto \tilde{F}$ be the mapping defined by

$$\tilde{F} = T'^{-T} F T^{-1}. \tag{7}$$

Clearly, if $\tilde{m} = Tm$ and $\tilde{m}' = T'm'$, then $m'^T Fm = \tilde{m}'^T \tilde{F} \tilde{m}$. Accordingly, the image of $\widehat{\tilde{F}}_{\mathrm{ALS}}$ by the inverse mapping $\tilde{F} \mapsto F$ can be viewed as a genuine estimate of F. We term this the *Hartley* (HRT) *estimate* of F and write $\widehat{F}_{\mathrm{HRT}}$; it is explicitly given by

$$\widehat{F}_{\mathrm{HRT}} = T'^T \widehat{\tilde{F}}_{\mathrm{ALS}} T. \tag{8}$$

The introduction of $\widehat{F}_{\mathrm{HRT}}$ is motivated by the fact that if the *modified condition number* of a non-negative definite matrix defined as the ratio of the greatest to the second smallest eigenvalues is large, then the two least eigenvalues are relatively close to one another, and consequently the eigenvectors associated with these nearby eigenvalues are "wobbly"—a small perturbation of the matrix entries can cause a significant change of these eigenvectors. The matrix $\tilde{A} = \sum_{i=1}^n u(\tilde{x}_i) u(\tilde{x}_i)^T$ serving to calculate $\widehat{\tilde{F}}_{\mathrm{ALS}}$ is in practice much better conditioned (in the above sense) than the matrix A with which $\widehat{F}_{\mathrm{ALS}}$ is calculated. As a result, Hartley's method is more stable, and in this sense more advantageous, than the ALS method.

5. Normalised Algebraic Least Squares

A useful modification of J_{ALS} is, as it turns out, the cost function defined by

$$J_{\mathrm{NALS}}(F; x_1, \ldots, x_n)$$
$$= \|T'^{-T} F T^{-1}\|_F^{-2} \sum_{i=1}^n (m_i'^T F m_i)^2. \tag{9}$$

The minimiser of J_{NALS} we call the *normalised algebraic least squares* (NALS) *estimate* of F and write $\widehat{F}_{\mathrm{NALS}}$. The precise sense in which the expressions entering J_{NALS} are normalised will be revealed later. A key property of J_{NALS} is that $J_{\mathrm{NALS}}(F; x_1, \ldots, x_n) = J_{\mathrm{ALS}}(\tilde{F}; \tilde{x}_1, \ldots, \tilde{x}_n)$ for any pair $(F, \tilde{F})$ satisfying (7). This property is instrumental in identifying the Hartley estimate as a minimiser of a cost function—namely it implies that

$$\widehat{F}_{\mathrm{HRT}} = \widehat{F}_{\mathrm{NALS}}. \tag{10}$$

The formula for J_{NALS} can be rewritten similarly to that for J_{ALS}. Denote by $C(s, t, m, n)$ the 9×9 matrix given by

$$C(s, t, m, n) = (t^2 I^* + nn^T) \otimes (s^2 I^* + mm^T),$$

where $I^* = \mathrm{diag}(1, 1, 0)$ and $\otimes$ denotes *Kronecker product*. A direct if tedious calculation shows that

$$\|T'^{-T} F T^{-1}\|_F^2 = \theta^T C \theta, \tag{11}$$

where C is given by

$$C = C(s, s', \overline{m}, \overline{m}'). \tag{12}$$

In view of (11), we can rewrite (9) as

$$J_{\mathrm{NALS}}(\theta; x_1, \ldots, x_n) = \frac{\theta^T A \theta}{\theta^T C \theta}.$$

One consequence of this formula is that $\widehat{\theta}_{\mathrm{NALS}}$ is a solution of the generalised eigenvalue problem

$$A\theta = \lambda C \theta \tag{13}$$

corresponding to the smallest eigenvalue. Since A may be ill-conditioned, solving the above eigenvector problem directly requires a numerically stable method. Leedan and Meer [11] proposed one such method which, when applied to the problem under consideration, employs *generalised singular value decomposition* (GSVD) of a pair of matrices (M, N) satisfying $A = M^T M$ and $C = N^T N$. Numerical experiments show (see later) that when this method is applied, the matrices A and C, of which the first is typically ill-conditioned, lead to a solution identical with the solution obtained using the well-conditioned matrix $\tilde{A}$ —in other words, equality (10) is experimentally confirmed.

6. Statistical Justification

Shifting the focus from matrices involved in computation of estimates (which may be well or ill conditioned) to cost functions, here we propose an alternative substantiation of the normalised eight-point algorithm. It is not *a priori* clear why J_{NALS} should be preferable to J_{ALS}. We shall present some explanation based on a statistical argument. Our reasoning will also provide the promised justification of the label 'normalised' for the terms forming J_{NALS}.

For each $i = 1, \ldots, n$, let r_i be the ith *residual* defined as

$$r_i = m_i'^T F m_i,$$

with F normalised for convenience. It is a fundamental observation that different residuals may carry different statistical weight. When m_i and m_i' are treated as sample values of independent multivariate random variables, the r_i are sample values of (typically) a *heteroscedastic* set of random variables, i.e. with member variables having different variances. The larger the variance of a particular r_i, the less reliable this residual is likely to be, and the more it should be devalued. Therefore, to account for heteroscedasticity,

it is natural to replace the simple cost function $\sum_{i=1}^{n} r_i^2$, effectively J_{ALS}, by the more complicated cost function $\sum_{i=1}^{n} r_i^2 / \mathrm{var}\,[r_i]$, where $\mathrm{var}\,[r]$ denotes the variance of r. The latter function is closer in form to a natural cost function derivable from the principle of maximum likelihood (cf. [2, 17]). We show that under a certain statistical model of data distribution, $\sum_{i=1}^{n} r_i^2 / \mathrm{var}\,[r_i]$ can be identified with J_{NALS}.

Assume that, for each $i = 1, \ldots, n$, the observed location $\mathbf{m}_i$ is a realisation of a random variable $\mathbf{m}_i = \overline{\mathbf{m}} + \Delta\mathbf{m}_i$, where $\overline{\mathbf{m}} = [\overline{m}_1, \overline{m}_2, 1]^T$ is a fixed, non-random location and $\Delta\mathbf{m}_i = [\Delta m_{1,i}, \Delta m_{1,2}, 0]^T$ is a random perturbation. Likewise, assume that m_i' is a realisation of a random variable $\mathbf{m}_i' = \overline{\mathbf{m}}' + \Delta\mathbf{m}_i'$ with non-random $\overline{\mathbf{m}}' = [\overline{m}_1', \overline{m}_2', 1]^T$ and random $\Delta\mathbf{m}_i' = [\Delta m_{1,i}', \Delta m_{2,i}', 0]^T$. Suppose that the following conditions hold:

- $\Delta\mathbf{m}_i, \Delta\mathbf{m}_j'$ $(i, j = 1, \ldots, n)$ are independent;
- $\mathsf{E}\,[\Delta\mathbf{m}_i] = \mathsf{E}\,[\Delta\mathbf{m}_i'] = \mathbf{0}$ for each $i = 1, \ldots, n$;
- there exist $\sigma > 0$ and $\sigma' > 0$ such that

$$\mathsf{E}\,\left[\Delta\mathbf{m}_i (\Delta\mathbf{m}_i)^T\right] = \sigma^2 \boldsymbol{I}^*,$$
$$\mathsf{E}\,\left[\Delta\mathbf{m}_i' (\Delta\mathbf{m}_i')^T\right] = \sigma'^2 \boldsymbol{I}^*$$

for each $i = 1, \ldots, n$.

Here E denotes expectation. Note that, effectively, *all* the $\mathbf{m}_i$ have common mean value $\overline{\mathbf{m}}$ and *all* the $\mathbf{m}_i'$ have common mean value $\overline{\mathbf{m}}'$. It is helpful to view $\overline{\mathbf{m}}$ and $\overline{\mathbf{m}}'$ as the centroids of some individual 'true' non-random locations $\overline{\mathbf{m}}_i$ and $\overline{\mathbf{m}}_i'$ that are not explicitly introduced, but are present in the background. An immediate upshot of this type of modelling is that the random perturbations $\Delta\mathbf{m}_i$ and $\Delta\mathbf{m}_i'$ cannot be regarded as small in typical situations—the magnitude of $\Delta\mathbf{m}_i$ and $\Delta\mathbf{m}_i'$ has to be big enough to account for the disparity between $\overline{\mathbf{m}}$ and the $\overline{\mathbf{m}}_i$, and $\overline{\mathbf{m}}'$ and the $\overline{\mathbf{m}}_i'$.

Denote by $\mathbf{r}_i = \mathbf{m}_i'^T \boldsymbol{F} \mathbf{m}_i$ the stochastic version of the ith residual. The above conditions on the $\mathbf{m}_i$ and $\mathbf{m}_i'$ together with the additional assumption that $\overline{\mathbf{m}}$ and $\overline{\mathbf{m}}'$ are 'true' locations bound by $\boldsymbol{F}$, in the sense that

$$\overline{\mathbf{m}}'^T \boldsymbol{F} \overline{\mathbf{m}} = \boldsymbol{\theta}^T \mathrm{vec}\,(\overline{\mathbf{m}}\,\overline{\mathbf{m}}'^T) = 0,$$

imply that all the residuals $\mathbf{r}_i$ have common variance $v = \boldsymbol{\theta}^T \boldsymbol{C}(\sigma, \sigma', \overline{\mathbf{m}}, \overline{\mathbf{m}}')\boldsymbol{\theta}$. Thus $\sum_{i=1}^{n} \mathbf{r}_i^2 / \mathrm{var}\,[\mathbf{r}_i]$, the random version of the cost function introduced earlier, can simply be written as $v^{-1} \sum_{i=1}^{n} \mathbf{r}_i^2$ with v^{-1} a common normalisation of all the residuals. Treating (5) and (6) as formulae for *estimates* of the parameters $\overline{\mathbf{m}}, \overline{\mathbf{m}}', \sigma, \sigma'$ used in our statistical model, replacing $\boldsymbol{C}(\sigma, \sigma', \overline{\mathbf{m}}, \overline{\mathbf{m}}')$ with $\boldsymbol{C}$ given by (12), and replacing the random residuals $\mathbf{r}_i$ with the non-random ones r_i, we arrive at the expression $(\boldsymbol{\theta}^T \boldsymbol{C} \boldsymbol{\theta})^{-1} \sum_{i=1}^{n} r_i^2$ which is identical with J_{NALS}. In this way, J_{NALS} becomes

statistically justified and its building blocks, the 'algebraic least squares' $(\boldsymbol{\theta}^T \boldsymbol{C} \boldsymbol{\theta})^{-1} r_i^2$, are found to be appropriately normalised.

7. Experimental Results

To assess whether the theoretical identity $\widehat{\boldsymbol{F}}_{\mathrm{HR\ T}} = \widehat{\boldsymbol{F}}_{\mathrm{NALS}}$ holds in practice, a series of simulations were run using synthetic data. The simulations were based on a set of 'true' pairs of corresponding points generated by selecting a realistic stereo camera configuration, randomly choosing many 3D points, and projecting the 3D points onto two image planes. Image resolution was chosen to be $1,000 \times 1,000$ pixels.

Two tests were conducted, each comprising $10,000$ trials. At each trial:

- the 'true' corresponding points were perturbed by homogeneous Gaussian jitter to produce noisy points;
- three fundamental matrices were generated from the noisy corresponding points using the non-normalised algebraic least-squares method (ALS), the normalised algebraic least-squares method (NALS), and Hartley's method (HRT);
- and the estimates were compared in the way described below.

The standard deviation of the noise was fixed at $\sigma = 1.0$ pixels (tests run with other levels of noise produced similar results).

In our experiments, the ALS estimate was computed by performing SVD of $\boldsymbol{M}$ given in (4) and taking the minimum right singular vector. Similarly, the Hartley estimate was computed by performing SVD of the matrix $\tilde{\boldsymbol{M}} = [\boldsymbol{u}(\tilde{\boldsymbol{x}}_1), \ldots, \boldsymbol{u}(\tilde{\boldsymbol{x}}_n)]^T$ and applying the back transformation prescribed by (8) to the minimum right singular vector (a standard SVD-correction step to produce a usable, rank-2 fundamental matrix before back-transforming was ignored). The NALS estimate was computed by employing Leedan–Meer's method based on the GSVD of $(\boldsymbol{M}, \boldsymbol{N})$, with $\boldsymbol{M}$ given by (4) and $\boldsymbol{N} = (s'\boldsymbol{I}^* + \boldsymbol{e}\overline{\mathbf{m}}'^T) \otimes (s\boldsymbol{I}^* + \boldsymbol{e}\overline{\mathbf{m}}^T)$, where $\boldsymbol{e} = [0, 0, 1]^T$.

In the first test, comparison of the estimates involved calculating two distances $d_1 = \min \|\widehat{\boldsymbol{F}}_{\mathrm{HRT}} \pm \widehat{\boldsymbol{F}}_{\mathrm{NALS}}\|_F$ and $d_2 = \min \|\widehat{\boldsymbol{F}}_{\mathrm{HR\ T}} \pm \widehat{\boldsymbol{F}}_{\mathrm{ALS}}\|_F$, with $\widehat{\boldsymbol{F}}_{\mathrm{HRT}}, \widehat{\boldsymbol{F}}_{\mathrm{NALS}}$, and $\widehat{\boldsymbol{F}}_{\mathrm{ALS}}$ having unit Frobenius norm. The first of these measures quantifies the discrepancy between the HRT and NALS estimates, the second informally gauges the significance of the values of the first. All results are plotted in Figure 1. The histogram of d_1 values shows that $\widehat{\boldsymbol{F}}_{\mathrm{HR\ T}}$ and $\widehat{\boldsymbol{F}}_{\mathrm{NALS}}$ are almost identical, with all values of d_1 less than
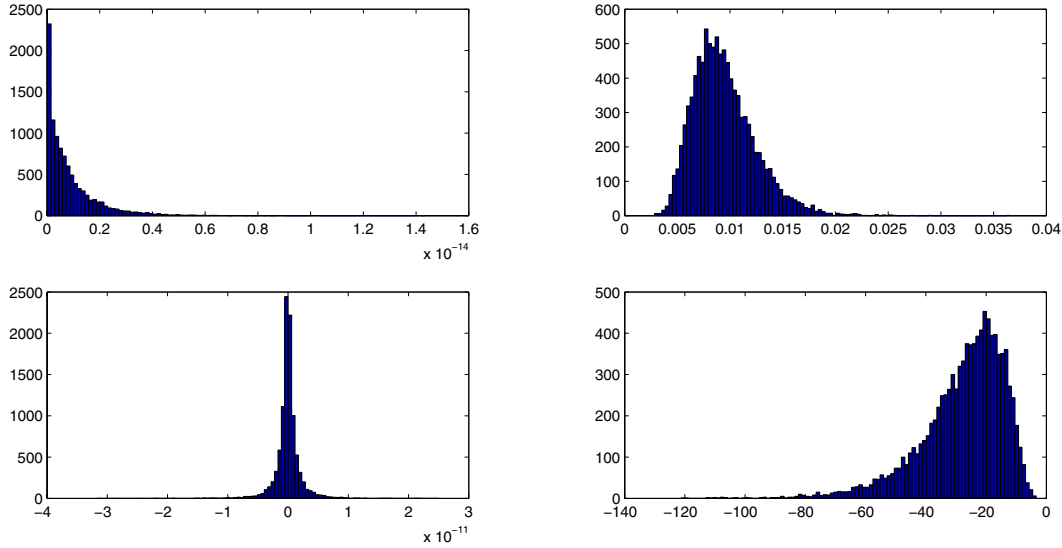
FIGURE 1. Histograms of values of d_1 (upper left), d_2 (upper right), d_3 (lower left), and d_4 (lower right) for $10,000$ tests.

1.5×10^{-14}. The significance of this may be gauged by noting that the d_2 histogram, capturing differences between the HRT and ALS estimates, exhibits values that are all greater than 1.5×10^{-3}.

The second test involved calculating the signed distances $d_3 = J_{\mathrm{AML}}(\widehat{\boldsymbol{F}}_{\mathrm{HRT}}) - J_{\mathrm{AML}}(\widehat{\boldsymbol{F}}_{\mathrm{NALS}})$ and $d_4 = J_{\mathrm{AML}}(\widehat{\boldsymbol{F}}_{\mathrm{HRT}}) - J_{\mathrm{AML}}(\widehat{\boldsymbol{F}}_{\mathrm{ALS}})$, where

$$
J_{\mathrm{AML}}(\boldsymbol{F}) = \sum_{i=1}^{n} \frac{(\boldsymbol{m}_i'^T \boldsymbol{F} \boldsymbol{m}_i)^2}{\boldsymbol{m}_i^T \boldsymbol{F} \boldsymbol{I}^* \boldsymbol{F}^T \boldsymbol{m}_i + \boldsymbol{m}_i'^T \boldsymbol{F}^T \boldsymbol{I}^* \boldsymbol{F} \boldsymbol{m}_i'}
$$

is the approximated maximum likelihood cost function commonly underlying more sophisticated iterative methods, associated with the default covariance $\boldsymbol{I}^*$ (e.g. see [2, 10, 17]). The d_3 histogram exhibits extremely small values centred on zero, confirming once again the practical equivalence of estimates $\widehat{\boldsymbol{F}}_{\mathrm{HRT}}$ and $\widehat{\boldsymbol{F}}_{\mathrm{NALS}}$. In contrast, the d_4 histogram shows differences in $\widehat{\boldsymbol{F}}_{\mathrm{HRT}}$ and $\widehat{\boldsymbol{F}}_{\mathrm{ALS}}$ that are very much larger.

8. Conclusion

A novel explanation has been presented for the improvement in performance of the normalised eight-point algorithm that results from using normalised data. It relies upon identifying a cost function that the algorithm effectively seeks to minimise. The advantageous character of the cost function is justified within a certain statistical model. The explanation avoids making any direct appeal to problem conditioning. Experimental results are presented that support the proposed approach.

Acknowledgements

References

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

[2] 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 and Machine Intelligence*, 22(11):1294–1303, 2000.

[3] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. A new approach to constrained parameter estimation applicable to some computer vision problems. In *Proc. Statistical Methods in Video Processing Workshop—held in conjunction with ECCV2002*, pages 43–48, 2002.

[4] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. FNS and HEIV: relating two vision parameter estimation frameworks. In *Proc. 12th Int. Conf. Image Analysis and Processing*, 2003.

[5] O. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In *Proc. Second European Conf. Computer Vision*, volume 588 of *Lecture Notes in Computer Science*, pages 563–578, 1992.

[6] O. Faugeras, Q.-T. Luong, and T. Papadopoulo. *The Geometry of Multiple Images*. MIT Press, Cambridge, Mass., 2001.

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

IEEE
COMPUTER
SOCIETY

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

[9] R. I. Hartley. Estimation of relative camera positions for uncalibrated cameras. In *Proc. Second European Conf. Computer Vision*, volume 588 of *Lecture Notes in Computer Science*, pages 579–587, 1992.

[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] H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(10):133–135, Sept. 1981.

[13] M. Mühlich and R. Mester. The role of total least squares in motion analysis. In *Proc. Fifth European Conf. Computer Vision*, volume 1407 of *Lecture Notes in Computer Science*, pages 305–321, 1998.

[14] P. H. S. Torr. *Outlier Detection and Motion Segmentation*. Ph.D. dissertation, Dept. of Engineering Science, University of Oxford, 1995.

[15] P. H. S. Torr and A. W. Fitzgibbon. Invariant fitting of two view geometry or "In defiance of the 8 point algorithm". Tech. Rep. MSR-TR-2002-50, Microsoft Research, May 2002.

[16] A. van den Hengel, W. Chojnacki, M. J. Brooks, and D. Gawley. A new constrained parameter estimator: experiments in fundamental matrix computation. In *Proc. 13th British Machine Vision Conference*, pages 468–476, 2002.

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