

FNS and HEIV: Relating Two Vision Parameter Estimation Frameworks

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

School of Computer Science, University of Adelaide, SA 5005, Australia
 CRC for Sensor Signal and Information Processing, Mawson Lakes, SA 5095, Australia

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

Abstract

Problems requiring accurate determination of parameters from image-based quantities arise often in computer vision. Two recent, independently developed frameworks for estimating such parameters are the FNS and HEIV schemes. Here it is shown that FNS and a core version of HEIV are essentially equivalent, solving a common underlying equation via different means. The analysis is driven by the search for a non-degenerate form of a certain generalised eigenvalue problem, and effectively leads to a new derivation of the relevant case of the HEIV algorithm. This work may be seen as an extension of previous efforts to rationalise and inter-relate a spectrum of estimators, including the renormalisation method of Kanatani and the normalised eight-point method of Hartley.

1. Introduction

Estimation of the parameters that describe a relationship between image feature locations, possibly across multiple cameras, is a central problem in computer vision. Basic examples include the stereo and motion problems of estimating coefficients of the *epipolar equation* [7] and the *differential epipolar equation* [1], and conic fitting [8]. The *principal equation* applicable in a variety of situations, including those specified above, takes the form

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

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. In some cases, the parameters are subject to an *ancillary constraint* not involving feature locations. A

common form of the ancillary constraint is

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

where, for some real number κ , ϕ is a scalar-valued function *homogeneous* of degree κ —that is such that $\phi(t\boldsymbol{\theta}) = t^\kappa \phi(\boldsymbol{\theta})$ for every $\boldsymbol{\theta}$ and every non-zero scalar t .

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. The *Gaussian model of errors* in data combined with the *principle of maximum likelihood* leads to the cost function

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

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 (see [2, 4, 9]). 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* $\hat{\boldsymbol{\theta}}_{\text{AML}}$. The *unconstrained* minimiser of J_{AML} obtained by ignoring the ancillary constraint and searching over all of the parameter space defines the *unconstrained approximated maximum likelihood estimate*, $\hat{\boldsymbol{\theta}}_{\text{AML}}^u$. The function $\boldsymbol{\theta} \mapsto J_{\text{AML}}(\boldsymbol{\theta}; \mathbf{x}_1, \dots, \mathbf{x}_n)$ is homogeneous of degree zero and the zero set of ϕ is unaffected by multiplication by non-zero scalars, so both $\hat{\boldsymbol{\theta}}_{\text{AML}}$ and $\hat{\boldsymbol{\theta}}_{\text{AML}}^u$ are determined only up to scale.

Various methods are available for finding $\hat{\boldsymbol{\theta}}_{\text{AML}}^u$. One is the *fundamental numerical scheme* (FNS) introduced by Chojnacki et al. in [4]. Another is, as will be revealed shortly, a certain version of the *heteroscedastic errors-in-variables* (HEIV) *scheme* that was first proposed by Leedan

1. Set $\theta_0 = \hat{\theta}_{\text{ALS}}$.
2. Assuming that θ_{k-1} is known, compute the matrix $\mathbf{X}_{\theta_{k-1}}$.
3. Compute a normalised eigenvector of $\mathbf{X}_{\theta_{k-1}}$ corresponding to the eigenvalue closest to zero (in absolute value) and take this eigenvector for θ_k .
4. If θ_k is sufficiently close to θ_{k-1} , then terminate the procedure; otherwise increment k and return to Step 2.

FIGURE 1. Fundamental numerical scheme.

and Meer [10] and further developed by Matei and Meer [11, 12]. The FNS method operates over the entire parameter space, whereas the HEIV method operates essentially on a subspace of one dimension less and recuperates the missing dimension in a single final step. This paper aims to understand the previously unclear relationship between the two schemes. It is shown that the algorithms are two different, but intimately related, means for numerically solving one and the same equation characterising $\hat{\theta}_{\text{AML}}^u$. In the analysis that follows, FNS is taken as a starting point, and HEIV is evolved via reduction of a certain generalised eigenvalue problem to a non-degenerate form. This approach effectively results in a new derivation of the relevant case of the HEIV algorithm.

Determination of $\hat{\theta}_{\text{AML}}$ is a much more complicated task than isolation of $\hat{\theta}_{\text{AML}}^u$. Recently, an integrated method for calculating $\hat{\theta}_{\text{AML}}$ was proposed that extends the FNS technique [5, 13]. The present contribution may provide a basis for designing a similar extension to the HEIV framework. From a broader perspective, this work may also be seen as an extension of previous efforts to rationalise and interrelate a spectrum of estimators, including the renormalisation method of Kanatani [3] and the normalised eight-point method of Hartley [6].

2. Fundamental Numerical Scheme

The unconstrained minimiser $\hat{\theta}_{\text{AML}}^u$ satisfies the *variational equation* for unconstrained minimisation

$$[\partial_{\theta} J_{\text{AML}}(\theta; \mathbf{x}_1, \dots, \mathbf{x}_n)]_{\theta=\hat{\theta}_{\text{AML}}^u} = \mathbf{0}^T \quad (3)$$

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

$$[\partial_{\theta} J_{\text{AML}}(\theta; \mathbf{x}_1, \dots, \mathbf{x}_n)]^T = 2\mathbf{X}_{\theta}\theta, \quad (4)$$

where \mathbf{X}_{θ} is an $l \times l$ symmetric matrix given by

$$\mathbf{X}_{\theta} = \sum_{i=1}^n \frac{\mathbf{A}_i}{\theta^T \mathbf{B}_i \theta} - \sum_{i=1}^n \frac{\theta^T \mathbf{A}_i \theta}{(\theta^T \mathbf{B}_i \theta)^2} \mathbf{B}_i,$$

$$\mathbf{A}_i = \mathbf{u}(\mathbf{x}_i)\mathbf{u}(\mathbf{x}_i)^T, \quad \mathbf{B}_i = \partial_{\mathbf{x}}\mathbf{u}(\mathbf{x}_i)\Lambda_{\mathbf{x}_i}\partial_{\mathbf{x}}\mathbf{u}(\mathbf{x}_i)^T.$$

Thus (3) can be written as

$$[\mathbf{X}_{\theta}\theta]_{\theta=\hat{\theta}_{\text{AML}}^u} = \mathbf{0}. \quad (5)$$

An algorithm for numerically solving this equation proposed in [4] exploits the fact that a vector θ satisfies (5) if and only if it is a solution of the *ordinary* eigenvalue problem

$$\mathbf{X}_{\theta}\xi = \lambda\xi \quad (6)$$

corresponding to the eigenvalue $\lambda = 0$. Thus if θ_{k-1} is an approximate solution, then an improved solution can be obtained by picking a vector θ_k from that eigenspace of $\mathbf{X}_{\theta_{k-1}}$ which most closely approximates the null space of \mathbf{X}_{θ} ; this eigenspace is, of course, the one corresponding to the eigenvalue closest to zero in absolute value. The fundamental numerical scheme [4] implementing this idea is presented in Figure 1. The scheme is seeded with the *algebraic least squares* (ALS) estimate, $\hat{\theta}_{\text{ALS}}$, defined as the unconstrained minimiser of the cost function $J_{\text{ALS}}(\theta; \mathbf{x}_1, \dots, \mathbf{x}_n) = \|\theta\|^{-2} \sum_{i=1}^n \theta^T \mathbf{A}_i \theta$, with $\|\theta\| = (\sum_{j=1}^l \theta_j^2)^{1/2}$. The estimate $\hat{\theta}_{\text{ALS}}$ 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* of the matrix $[\mathbf{u}(\mathbf{x}_1), \dots, \mathbf{u}(\mathbf{x}_n)]^T$.

3. Basic HEIV Scheme

An alternative parameter estimation framework, derived in a quite different manner to FNS, has been proposed by Leedan and Meer [10] and further extended by Matei and Meer [11, 12]. As will become apparent shortly, a core method of this framework that we will term *HEIV with carrier bias correction eliminated* is effectively a different means for numerically solving (5). In one form this method relies upon re-expressing \mathbf{X}_{θ} as

$$\mathbf{X}_{\theta} = \mathbf{M}_{\theta} - \mathbf{N}_{\theta}$$

with

$$\mathbf{M}_{\theta} = \sum_{i=1}^n \frac{\mathbf{A}_i}{\theta^T \mathbf{B}_i \theta}, \quad \mathbf{N}_{\theta} = \sum_{i=1}^n \frac{\theta^T \mathbf{A}_i \theta}{(\theta^T \mathbf{B}_i \theta)^2} \mathbf{B}_i,$$

and restating the variational equation (5) as

$$\mathbf{M}_{\theta}\theta = \mathbf{N}_{\theta}\theta, \quad (7)$$

1. Set $\theta_0 = \hat{\theta}_{\text{ALS}}$.
2. Assuming that θ_{k-1} is known, compute the matrices $M_{\theta_{k-1}}$ and $N_{\theta_{k-1}}$.
3. Compute a normalised eigenvector of the eigenvalue problem
$$M_{\theta_{k-1}}\xi = \lambda N_{\theta_{k-1}}\xi$$
corresponding to the eigenvalue closest to 1 and take this eigenvector for θ_k .
4. If θ_k is sufficiently close to θ_{k-1} , then terminate the procedure; otherwise increment k and return to Step 2.

FIGURE 2. Basic HEIV scheme.

where the evaluation at $\hat{\theta}_{\text{AML}}^u$ is dropped for clarity. The matrices M_θ and N_θ are non-negative definite (with M_θ generically positive definite if $n \geq l$), so θ can be viewed as a solution of the *generalised* eigenvalue problem

$$M_\theta \xi = \lambda N_\theta \xi \quad (8)$$

corresponding to the eigenvalue $\lambda = 1$. The *basic* heteroscedastic errors-in-variables scheme is an algorithm for solving (7) that exploits the above eigenvalue problem in a manner analogous to that in which FNS utilises the eigenvalue problem (6). The scheme is a variation on the technique proposed in [10–12]. The details are given in Figure 2.

As is easily seen from (11) below, the null space of each matrix B_i contains the length l vector $[0, \dots, 0, 1]^T$. Consequently, N_θ is singular. It turns out that the eigenvalue problem (8) can be reduced to a similar problem involving a positive definite right-hand side matrix. A specific reduction will be described next, namely one resulting from a suitable reformulation of the variational equation (7).

4. Reduced Variational Equation

The vector $u(x)$ has one entry equal to 1 and can be written as

$$u(x) = [z(x)^T, 1]^T, \quad (9)$$

where $z(x)$ is a ‘pure measurement’ vector of length $l - 1$. The vector of parameters θ can be partitioned conformally as

$$\theta = [\eta^T, \alpha]^T$$

with η a length $l - 1$ vector and α a scalar. An immediate consequence of (9) is the representation

$$A_i = u(x_i)u(x_i)^T = \begin{bmatrix} z_i z_i^T & z_i \\ z_i^T & 1 \end{bmatrix} \quad (10)$$

for each $i = 1, \dots, n$, where z_i is short for $z(x_i)$. Another consequence of (9) is the identity

$$\partial_x u(x) = \begin{bmatrix} \partial_x z(x) \\ \mathbf{0}^T \end{bmatrix},$$

which implies that, for each $i = 1, \dots, n$,

$$B_i = \begin{bmatrix} B_i^0 & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix}, \quad (11)$$

with $B_i^0 = \partial_x z(x_i) A_{x_i} \partial_x z(x_i)^T$. For each $i = 1, \dots, n$, define a weight

$$\beta_i = \frac{1}{\eta^T B_i^0 \eta}$$

that depends on the i th element of data x_i , its covariance A_{x_i} , and the parameter vector η . Let \tilde{z} be the centroid of the z_i given by

$$\tilde{z} = \frac{\sum_{i=1}^n \beta_i z_i}{\sum_{i=1}^n \beta_i},$$

and, for each $i = 1, \dots, n$, let

$$z'_i = z_i - \tilde{z}$$

be the i th pure measurement vector relative to \tilde{z} . Define two $(l - 1) \times (l - 1)$ matrices

$$M'_\eta = \sum_{i=1}^n \beta_i z'_i z'^T_i, \quad N'_\eta = \sum_{i=1}^n (\beta_i z'^T_i \eta)^2 B_i^0.$$

Upon introducing

$$A_i^0 = z'_i z'^T_i,$$

the matrices M'_η and N'_η can also be written as

$$M'_\eta = \sum_{i=1}^n \frac{A_i^0}{\eta^T B_i^0 \eta}, \quad N'_\eta = \sum_{i=1}^n \frac{\eta^T A_i^0 \eta}{(\eta^T B_i^0 \eta)^2} B_i^0,$$

which reveals their resemblance to M_θ and N_θ . Obviously, M'_η and N'_η both depend not only on η but also on the data and their covariances.

A fundamental result that can now be established is that $\theta = [\eta^T, \alpha]^T$ satisfies (7) if and only if the following system of equations holds:

$$M'_\eta \eta = N'_\eta \eta, \quad (12)$$

$$\alpha = -\tilde{z}^T \eta. \quad (13)$$

The first of these equations involves only η and can be solved in isolation; the second expresses α in terms of η . Of the two constraints, the first plays a leading role and will be called the *reduced variational equation*. A key feature of this equation is that its right-hand side matrix N'_η , unlike N_θ , is generically *positive definite* if $n \geq l$.

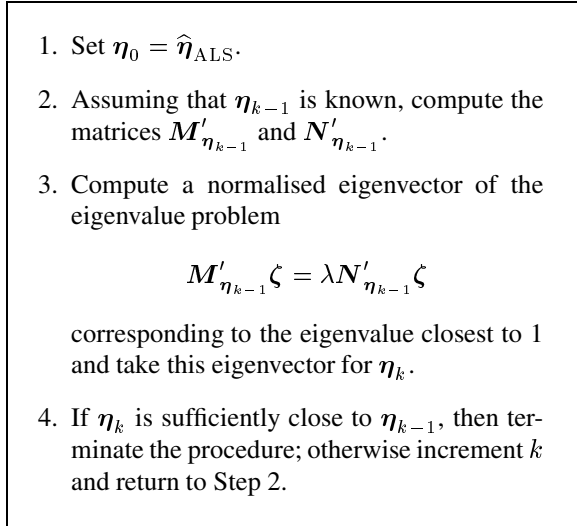


FIGURE 3. Reduced HEIV scheme.

5. Reduced HEIV Scheme

The algebraic least squares estimates $\hat{\eta}_{\text{ALS}}$ and $\hat{\alpha}_{\text{ALS}}$ are naturally defined as the respective components in the representation

$$\hat{\theta}_{\text{ALS}} = [(\hat{\eta}_{\text{ALS}})^T, \hat{\alpha}_{\text{ALS}}]^T.$$

Analogously, the unconstrained approximated maximum likelihood estimates $\hat{\eta}_{\text{AML}}^u$ and $\hat{\alpha}_{\text{AML}}^u$ are defined via the decomposition

$$\hat{\theta}_{\text{AML}}^u = [(\hat{\eta}_{\text{AML}}^u)^T, \hat{\alpha}_{\text{AML}}^u]^T.$$

In view of (13), $\hat{\eta}_{\text{AML}}^u$ is uniquely determined by $\hat{\eta}_{\text{AML}}^u$ —when the centroid \tilde{z} is taken with the weights

$$\beta_i = \frac{1}{(\hat{\eta}_{\text{AML}}^u)^T B_i^0 \hat{\eta}_{\text{AML}}^u},$$

then $\hat{\alpha}_{\text{AML}}^u = -\tilde{z}^T \hat{\eta}_{\text{AML}}^u$. Now, the generalised eigenvalue problem

$$M'_{\eta} \zeta = \lambda N'_{\eta} \zeta, \quad (14)$$

is non-degenerate: the matrix N'_{η} is positive-definite. Accordingly, $\hat{\eta}_{\text{AML}}^u$ can be determined with use of a simple modification of the HEIV algorithm. The steps of this *reduced* HEIV scheme are given in Figure 3. It is essentially in this form that the HEIV algorithm was first advanced [10, 12]. The original version employs a slightly different, bias-corrected form of the *vector of carriers* $u(x)$. The reduced scheme, based solely on $u(x)$, constitutes HEIV with carrier bias correction eliminated. Both versions are comparable in performance, but since the one with carrier bias correction eliminated is somewhat simpler, it is this version that was eventually recognised as the fundamental form of the HEIV algorithm [11].

<i>Max. diff.</i>	ALS	FNS	HB	HR
HEIV	146	4.7×10^{-6}	7.5×10^{-5}	2.6×10^{-8}
HR	146	4.7×10^{-6}	7.5×10^{-5}	
HB	146	7.1×10^{-5}		
FNS	146			
<i>Avg. diff.</i>	ALS	FNS	HB	HR
HEIV	24.8	5.7×10^{-8}	2.0×10^{-6}	2.5×10^{-10}
HR	24.8	5.8×10^{-8}	2.0×10^{-6}	
HB	24.8	2.0×10^{-6}		
FNS	24.8			

TABLE 1. Maximum and average differences in J_{AML} values for different estimation methods.

6. Stable HEIV Scheme

The reduced HEIV scheme is locally convergent—to work it requires the initial iterate to be close to a solution of (12). A more stable version of the algorithm, able to cope with a less accurate initial iterate, results from selecting the eigenvector corresponding to the smallest eigenvalue instead of the eigenvector corresponding to the eigenvalue closest to 1. Leedan and Meer remark that this modified method converges successfully (in fact with high convergence rate) even when seeded with a random initial estimate. Typically, the minimal eigenvalues computed after a first iteration are also the closest to 1, and so from the second iteration onwards the modified algorithm acts effectively as the original version. Without the modification, the scheme may exhibit slow convergence or even divergence.

7. Experiments

Relative performance of the FNS and HEIV methods was experimentally assessed by running a series of simulations involving synthetic data. The particular problem considered was estimation of epipolar geometry. It turns out that in this case the vector of carriers is unbiased and the original version of HEIV involving bias corrected carriers coincides with the version with carrier bias correction eliminated. A single element of data took the form of matched corresponding points from left and right images of a stereo pair, and the goal was to estimate the associated *fundamental matrix*. Details of the various expressions involved are presented elsewhere [4].

In our experiments, five estimation methods were tested, denoted as ALS, FNS, HB, HR, and HEIV. ALS is the simple, direct algebraic least squares method described in Section 2. It is included as a method of a different category to

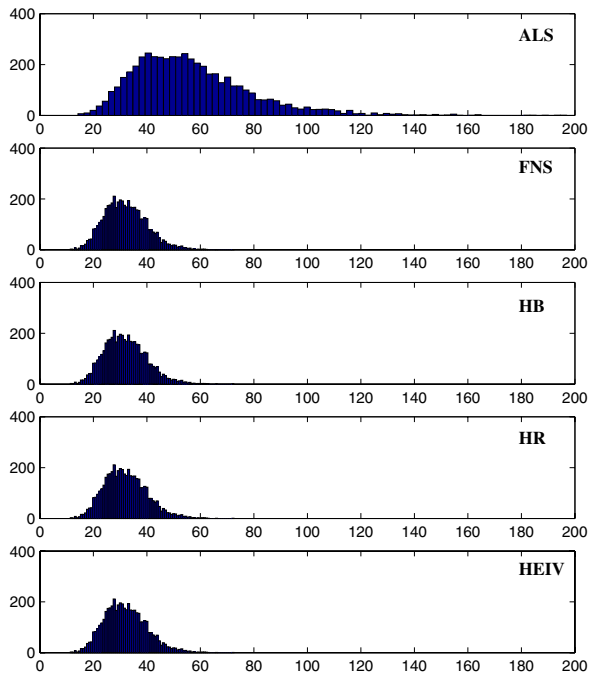


FIGURE 4. Performance histograms for each of the methods, with J_{AML} bins on the x -axes, and frequencies of occurrence on the y -axes.

give a sense of scale to the forthcoming numerical results. The FNS, HB, and HR methods were implemented as specified in Figures 1, 2, and 3, respectively. These iterative methods were terminated when the difference in norm between successive estimates was less than a common, very small threshold. Estimates of the final method, HEIV, were obtained using the MATLAB source code supplied by the authors of the original HEIV papers.¹

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. Only those scene points were considered that had both projections confined to the image size of $1,000 \times 1,000$ pixels.

For each of $N = 5,000$ iterations, the true corresponding points were perturbed by homogeneous Gaussian jitter to produce noisy points. These noisy points were then used to generate a fundamental matrix estimate for each of the five estimation methods. For each estimate, the value of the J_{AML} cost function was computed. Comparison was undertaken in this realm as J_{AML} is the basis for our rationalising and linking of the various iterative methods considered. Note that the singularity constraint was not imposed as this would otherwise obfuscate comparison (the constraint is

¹<http://www.caip.rutgers.edu/riul/research/code.html>

usually implemented as a separate post-process). In these tests, the level of noise was fixed at $\sigma = 1.0$ pixels, although similar results were obtained using different noise levels.

Figure 4 shows the histograms of J_{AML} values associated with each of the estimators. In contrast with the ALS method, the iterative methods generate very similar response profiles. Table 1 compares estimators pairwise by showing both the maximum and average differences in associated J_{AML} values over the complete set of trials. The respective top left elements are computed via the expressions $\max_{1 \leq i \leq N} |J_{\text{AML}}(\hat{\theta}_{\text{HEIV}}^i) - J_{\text{AML}}(\hat{\theta}_{\text{ALS}}^i)|$ and $N^{-1} \sum_{i=1}^N |J_{\text{AML}}(\hat{\theta}_{\text{HEIV}}^i) - J_{\text{AML}}(\hat{\theta}_{\text{ALS}}^i)|$. The results demonstrate that the methods FNS, HB, HR, and HEIV deliver estimates whose associated J_{AML} values are extremely close. As would be expected from the earlier theory, the HR and HEIV methods prove to be almost numerically identical.

8. Conclusion

In this work, aspects of the FNS and HEIV frameworks for estimating parameters from image-based data were examined. It was shown that FNS and a core version of HEIV are essentially equivalent, both in terms of analytical formulation and numerical outcome. In this way, further understanding is gained about the inter-relationships between members of the spectrum of estimators available for computation of geometric parameters. Given that the FNS scheme has been recently upgraded to incorporate constraint in a fully integrated manner, the opportunity now exists to enhance the HEIV framework in a similar manner.

Acknowledgements

The support of the Australian Research Council is gratefully acknowledged.

References

- [1] M. J. Brooks, W. Chojnacki, and L. Baumela. Determining the egomotion of an uncalibrated camera from instantaneous optical flow. *J. Opt. Soc. Amer. A*, 14(10):2670–2677, 1997.
- [2] M. J. Brooks, W. Chojnacki, D. Gawley, and A. van den Hengel. What value covariance information in estimating vision parameters? In *Proc. Eighth Int. Conf. Computer Vision*, volume 1, pages 302–308, 2001.
- [3] 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.
- [4] 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.

- [5] 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.
- [6] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. A statistical rationalisation of Hartley’s normalised eight-point algorithm. In *Proc. 12th Int. Conf. Image Analysis and Processing*, 2003.
- [7] O. D. Faugeras. *Three-Dimensional Computer Vision: A Geometric Viewpoint*. MIT Press, Cambridge, Mass., 1993.
- [8] A. Fitzgibbon, M. Pilu, and R. B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 21(5):476–480, 1999.
- [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. *Int. J. Computer Vision*, 37(2):127–150, 2000.
- [11] B. Matei. *Heteroscedastic Errors-in-Variables Models in Computer Vision*. Ph.D. dissertation, Dept. of Electrical and Computer Engineering, Rutgers University, New Brunswick, NJ, May 2001. Available at <http://www.caip.rutgers.edu/riul/research/theses.html>.
- [12] B. Matei and P. Meer. A general method for errors-in-variables problems in computer vision. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, volume 2, pages 18–25, 2000.
- [13] 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 Conf.*, volume 2, pages 468–476, 2002.