From FNS to HEIV: A Link between Two Vision Parameter Estimation Methods

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

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 nondegenerate form of a certain generalized 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 rationalize and interrelate a spectrum of estimators, including the renormalization method of Kanatani and the normalized eight-point method of Hartley.

Index Terms—Statistical methods, maximum likelihood, (un)constrained minimization, fundamental matrix, epipolar equation.

———————————— ♦ ————————————

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* [1] and the *differential epipolar equation* [2], and conic fitting [3]. The *principal equation* applicable in a variety of situations, including those specified above, takes 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 an element of the data (for example, the locations of a pair of corresponding points); and $\boldsymbol{u}(\boldsymbol{x}) = [u_1(\boldsymbol{x}), \ldots, u_l(\boldsymbol{x})]^T$ is a vector with the data transformed in a problem-dependent manner such that: 1) each component $u_i(\boldsymbol{x})$ is a quadratic form in the compound vector $[\boldsymbol{x}^T, 1]^T$ and 2) 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

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

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

The estimation problem associated with (1) and (2) can be stated as follows: Given a collection $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ of *image data* and a meaningful *cost function* that characterizes the extent to which any particular $\boldsymbol{\theta}$ fails to satisfy the system of the copies of (1) associated with $\boldsymbol{x} = \boldsymbol{x}_i$ ($i = 1, \ldots, n$), find $\boldsymbol{\theta} \neq \boldsymbol{0}$ 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

———————————————

• *The authors are with the School of Computer Science, University of Adelaide, SA 5005, Australia and the CRC for Sensor Signal and Information Processing, Mawson Lakes, SA 5095, Australia.*
E-mail: {wojtek, mjb, hengel, dg}@cs.adelaide.edu.au.

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

where, for any length k vector $\boldsymbol{y}$, $\partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{y})$ denotes the $l \times k$ matrix of the partial derivatives of the function $\boldsymbol{x} \mapsto \boldsymbol{u}(\boldsymbol{x})$ evaluated at $\boldsymbol{y}$, and, for each $i = 1, \ldots, n$, $\boldsymbol{\Lambda}_{\boldsymbol{x}_i}$ is a $k \times k$ symmetric *covariance matrix* describing the uncertainty of the data point $\boldsymbol{x}_i$ (see [4], [5], [6]). If J_{AML} is minimized over those nonzero parameter vectors for which (2) holds, then the vector at which the minimum of J_{AML} is attained, the *constrained* minimizer of J_{AML}, defines the *approximated maximum likelihood estimate* $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$. The *unconstrained* minimizer of J_{AML} obtained by ignoring the ancillary constraint and searching over all of the parameter space defines the *unconstrained approximated maximum likelihood estimate* $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$. The function $\boldsymbol{\theta} \mapsto J_{\mathrm{AML}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ is homogeneous of degree zero and the zero set of ϕ is unaffected by multiplication by nonzero scalars, so both $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$ and $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$ are determined only up to scale.

Various methods are available for finding $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$. One is the *fundamental numerical scheme* (FNS) introduced by Chojnacki et al. in [5]. Another is, as will be revealed shortly, a certain version of the *heteroscedastic errors-in-variables* (HEIV) *scheme* that was first proposed by Leedan and Meer [7] and further developed by Matei and Meer [8], [9]. 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 characterizing $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$. In the analysis that follows, FNS is taken as a starting point, and HEIV is evolved via reduction of a certain generalized eigenvalue problem to a nondegenerate form. This approach effectively results in a new derivation of the relevant case of the HEIV algorithm.

Determination of $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$ is a much more complicated task than isolation of $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$. Recently, an integrated method for calculating $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}$ was proposed that extends the FNS technique [10], [11]. 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 rationalize and interrelate a spectrum of estimators, including the renormalization method of Kanatani [12] and the normalized eight-point method of Hartley [13].

2 FUNDAMENTAL NUMERICAL SCHEME

The unconstrained minimizer $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$ satisfies the *variational equation* for unconstrained minimization

$$[\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)]_{\boldsymbol{\theta} = \widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u} = \boldsymbol{0}^T \tag{3}$$

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

$$[\partial_{\boldsymbol{\theta}} J_{\mathrm{AML}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)]^T = 2 \boldsymbol{X}_{\boldsymbol{\theta}} \boldsymbol{\theta}, \tag{4}$$

where $\boldsymbol{X}_{\boldsymbol{\theta}}$ is an $l \times l$ symmetric matrix given by

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

$$\boldsymbol{A}_i = \boldsymbol{u}(\boldsymbol{x}_i) \boldsymbol{u}(\boldsymbol{x}_i)^T, \quad \boldsymbol{B}_i = \partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{x}_i) \boldsymbol{\Lambda}_{\boldsymbol{x}_i} \partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{x}_i)^T.$$

Thus, (3) can be written as

$$[\boldsymbol{X}_{\boldsymbol{\theta}} \boldsymbol{\theta}]_{\boldsymbol{\theta} = \widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u} = \boldsymbol{0}. \tag{5}$$

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

1) Set $\boldsymbol{\theta}_0 = \widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$.

2) Assuming that $\boldsymbol{\theta}_{k-1}$ is known, compute the matrix $\boldsymbol{X}_{\boldsymbol{\theta}_{k-1}}$.

3) Compute a normalized 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.

Fig. 1. Fundamental numerical scheme.

$$\boldsymbol{X}_{\boldsymbol{\theta}}\boldsymbol{\xi} = \lambda\boldsymbol{\xi} \qquad (6)$$

corresponding to the eigenvalue $\lambda = 0$. Thus, if $\boldsymbol{\theta}_{k-1}$ is an approximate solution, then an improved solution 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 in absolute value. The fundamental numerical scheme [5] implementing this idea is presented in Fig. 1. The scheme is seeded with the *algebraic least squares* (ALS) *estimate*, $\widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$, defined as the unconstrained minimizer of the cost function $J_{\mathrm{ALS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \|\boldsymbol{\theta}\|^{-2} \sum_{i=1}^{n} \boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}$, 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 $\sum_{i=1}^{n} \boldsymbol{A}_i$ associated with the smallest eigenvalue, and this can be found by performing *singular-value decomposition* of the matrix $[\boldsymbol{u}(\boldsymbol{x}_1), \ldots, \boldsymbol{u}(\boldsymbol{x}_n)]^T$.

3 BASIC HEIV SCHEME

An alternative parameter estimation framework, derived in a quite different manner to FNS, has been proposed by Leedan and Meer [7] and further extended by Matei and Meer [8], [9]. 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 $\boldsymbol{X}_{\boldsymbol{\theta}}$ as

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

with

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

and restating the variational equation (5) as

$$\boldsymbol{M}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \boldsymbol{N}_{\boldsymbol{\theta}}\boldsymbol{\theta}, \qquad (7)$$

where the evaluation at $\widehat{\boldsymbol{\theta}}_{\mathrm{AML}}^u$ is dropped for clarity. The matrices $\boldsymbol{M}_{\boldsymbol{\theta}}$ and $\boldsymbol{N}_{\boldsymbol{\theta}}$ are nonnegative definite (with $\boldsymbol{M}_{\boldsymbol{\theta}}$ generically positive definite if $n \geq l$), so $\boldsymbol{\theta}$ can be viewed as a solution of the *generalized* eigenvalue problem

$$\boldsymbol{M}_{\boldsymbol{\theta}}\boldsymbol{\xi} = \lambda\boldsymbol{N}_{\boldsymbol{\theta}}\boldsymbol{\xi} \qquad (8)$$

corresponding to the eigenvalue $\lambda = 1$. The *basic* heteroscedastic errors-in-variables scheme (see Section 7 for a clue as to the name) is an algorithm for solving (7) that exploits the above eigenvalue problem in a manner analogous to that in which FNS utilizes the eigenvalue problem (6). The scheme is a variation on the technique proposed in [7], [8], [9]. The details are given in Fig. 2.

As is easily seen from (12) below, the null space of each matrix $\boldsymbol{B}_i$ contains the length l vector $[0, \ldots, 0, 1]^T$. Consequently, $\boldsymbol{N}_{\boldsymbol{\theta}}$ is singular. Now, if $\boldsymbol{\theta}$ is a "true" parameter vector and $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ are noise-free data satisfying $\boldsymbol{\theta}^T \boldsymbol{u}(\boldsymbol{x}_i) = 0$ for each $i = 1, \ldots, n$, then $\boldsymbol{M}_{\boldsymbol{\theta}}$ is also singular, having $\boldsymbol{\theta}$ in its null space. As is well known, if both $\boldsymbol{M}_{\boldsymbol{\theta}}$ and $\boldsymbol{N}_{\boldsymbol{\theta}}$ are simultaneously rank-deficient (or almost rank-deficient, which happens, for example, when $\boldsymbol{\theta}$ is close to a "true"

1) Set $\boldsymbol{\theta}_0 = \widehat{\boldsymbol{\theta}}_{\mathrm{ALS}}$.

2) Assuming that $\boldsymbol{\theta}_{k-1}$ is known, compute the matrices $\boldsymbol{M}_{\boldsymbol{\theta}_{k-1}}$ and $\boldsymbol{N}_{\boldsymbol{\theta}_{k-1}}$.

3) Compute a normalized eigenvector of the eigenvalue problem

$$\boldsymbol{M}_{\boldsymbol{\theta}_{k-1}}\boldsymbol{\xi} = \lambda\boldsymbol{N}_{\boldsymbol{\theta}_{k-1}}\boldsymbol{\xi}$$

corresponding to the eigenvalue closest to 1 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.

Fig. 2. Basic HEIV scheme.

parameter vector and the data are just a tiny perturbation of noise-free data), then solving the eigenvalue problem (8) is prone to numerical instability [14]. One way to get around this difficulty is to reduce the eigenvector problem (8) to a similar problem involving a positive definite right-hand side matrix. Such a reduction is best achieved by reformulating the variational equation (7). This is described next.

4 REDUCED VARIATIONAL EQUATION

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

$$\boldsymbol{u}(\boldsymbol{x}) = [\boldsymbol{z}(\boldsymbol{x})^T, 1]^T, \qquad (9)$$

where $\boldsymbol{z}(\boldsymbol{x})$ is a "pure measurement" vector of length $l - 1$. The vector of parameters $\boldsymbol{\theta}$ can be partitioned conformally as

$$\boldsymbol{\theta} = [\boldsymbol{\eta}^T, \alpha]^T \qquad (10)$$

with $\boldsymbol{\eta}$ a length $l - 1$ vector and α a scalar. We are going to show that the variational equation (7) is equivalent to a system of two equations, one of which involves only $\boldsymbol{\eta}$ and can be solved in isolation, and the other expresses α in terms of $\boldsymbol{\eta}$. The first equation will lead to a desired nondegenerate eigenvalue problem.

We begin by noting that, in view of (9),

$$\boldsymbol{A}_i = \boldsymbol{u}(\boldsymbol{x}_i)\boldsymbol{u}(\boldsymbol{x}_i)^T = \begin{bmatrix} \boldsymbol{z}_i\boldsymbol{z}_i^T & \boldsymbol{z}_i \\ \boldsymbol{z}_i^T & 1 \end{bmatrix} \qquad (11)$$

for each $i = 1, \ldots, n$; here, of course, $\boldsymbol{z}_i$ is short for $\boldsymbol{z}(\boldsymbol{x}_i)$. Another consequence of (9) is the identity

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

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

$$\boldsymbol{B}_i = \begin{bmatrix} \boldsymbol{B}_i^0 & \boldsymbol{0} \\ \boldsymbol{0}^T & 0 \end{bmatrix}, \qquad (12)$$

with $\boldsymbol{B}_i^0 = \partial_{\boldsymbol{x}}\boldsymbol{z}(\boldsymbol{x}_i)\boldsymbol{\Lambda}_{\boldsymbol{x}_i}\partial_{\boldsymbol{x}}\boldsymbol{z}(\boldsymbol{x}_i)^T$. For each $i = 1, \ldots, n$, define a weight

$$\beta_i = \frac{1}{\boldsymbol{\eta}^T \boldsymbol{B}_i^0 \boldsymbol{\eta}} \qquad (13)$$

that depends on the ith element of data $\boldsymbol{x}_i$, its covariance $\boldsymbol{\Lambda}_{\boldsymbol{x}_i}$, and the parameter vector $\boldsymbol{\eta}$. Let $\tilde{\boldsymbol{z}}$ be the centroid of the $\boldsymbol{z}_i$ given by

$$\tilde{\boldsymbol{z}} = \frac{\sum_{i=1}^{n} \beta_i \boldsymbol{z}_i}{\sum_{i=1}^{n} \beta_i}, \qquad (14)$$

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

$$\mathbf{z}_i' = \mathbf{z}_i - \tilde{\mathbf{z}} \qquad (15)$$

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

$$\mathbf{M}_{\boldsymbol{\eta}}' = \sum_{i=1}^n \beta_i \mathbf{z}_i' \mathbf{z}_i'^T, \quad \mathbf{N}_{\boldsymbol{\eta}}' = \sum_{i=1}^n \left(\beta_i \mathbf{z}_i'^T \boldsymbol{\eta}\right)^2 \mathbf{B}_i^0.$$

Upon introducing $\mathbf{A}_i^0 = \mathbf{z}_i' \mathbf{z}_i'^T$, the matrices $\mathbf{M}_{\boldsymbol{\eta}}'$ and $\mathbf{N}_{\boldsymbol{\eta}}'$ can also be written as

$$\mathbf{M}_{\boldsymbol{\eta}}' = \sum_{i=1}^n \frac{\mathbf{A}_i^0}{\boldsymbol{\eta}^T \mathbf{B}_i^0 \boldsymbol{\eta}}, \quad \mathbf{N}_{\boldsymbol{\eta}}' = \sum_{i=1}^n \frac{\boldsymbol{\eta}^T \mathbf{A}_i^0 \boldsymbol{\eta}}{(\boldsymbol{\eta}^T \mathbf{B}_i^0 \boldsymbol{\eta})^2} \mathbf{B}_i^0,$$

which reveals their resemblance to $\mathbf{M}_{\boldsymbol{\theta}}$ and $\mathbf{N}_{\boldsymbol{\theta}}$. The choice of the weights β_i is largely motivated by the desire to achieve such a resemblance in the first place. Obviously, $\mathbf{M}_{\boldsymbol{\eta}}'$ and $\mathbf{N}_{\boldsymbol{\eta}}'$ both depend not only on $\boldsymbol{\eta}$, but also on the data and their covariances.

We now show that $\boldsymbol{\theta} = [\boldsymbol{\eta}^T, \alpha]^T$ satisfies (7) if and only if the following system of equations holds:

$$\mathbf{M}_{\boldsymbol{\eta}}' \boldsymbol{\eta} = \mathbf{N}_{\boldsymbol{\eta}}' \boldsymbol{\eta}, \qquad (16)$$
$$\alpha = -\tilde{\mathbf{z}}^T \boldsymbol{\eta}. \qquad (17)$$

Note that the first equation constrains solely $\boldsymbol{\eta}$ and, therefore, can be solved separately. Once $\boldsymbol{\eta}$ is determined, α is readily prescribed by the second equation. As will become apparent shortly, the equations decouple as a result of the specific choice of the centroid $\tilde{\mathbf{z}}$. 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 $\mathbf{N}_{\boldsymbol{\eta}}'$, unlike $\mathbf{N}_{\boldsymbol{\theta}}$, is generically *positive definite* if $n \geq l$.

To show the equivalence of (7) and the system comprising (16) and (17), first note that, by (10) and (12), $\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta} = \boldsymbol{\eta}^T \mathbf{B}_i^0 \boldsymbol{\eta}$ and, further, by (13),

$$\beta_i = \frac{1}{\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta}} \qquad (18)$$

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

$$\mathbf{M}_{\boldsymbol{\theta}} = \sum_{i=1}^n \beta_i \mathbf{A}_i. \qquad (19)$$

Observe next that, in view of (12),

$$\mathbf{N}_{\boldsymbol{\theta}} = \begin{bmatrix} \mathbf{N}_{\boldsymbol{\theta}}^0 & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix} \qquad (20)$$

with

$$\mathbf{N}_{\boldsymbol{\theta}}^0 = \sum_{i=1}^n \frac{\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}}{(\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta})^2} \mathbf{B}_i^0.$$

Now, if $\boldsymbol{\theta} = [\boldsymbol{\eta}^T, \alpha]^T$ satisfies (7), then, in view of (11), (19), and (20), (7) can be rewritten as

$$\sum_{i=1}^n \beta_i \begin{bmatrix} \mathbf{z}_i \mathbf{z}_i^T & \mathbf{z}_i \\ \mathbf{z}_i^T & 1 \end{bmatrix} \begin{bmatrix} \boldsymbol{\eta} \\ \alpha \end{bmatrix} = \begin{bmatrix} \mathbf{N}_{\boldsymbol{\theta}}^0 & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\eta} \\ \alpha \end{bmatrix},$$

or equivalently as the system

$$\sum_{i=1}^n \beta_i (\alpha \mathbf{z}_i + \mathbf{z}_i \mathbf{z}_i^T \boldsymbol{\eta}) = \mathbf{N}_{\boldsymbol{\theta}}^0 \boldsymbol{\eta}, \qquad (21)$$

$$\sum_{i=1}^n \beta_i (\alpha + \mathbf{z}_i^T \boldsymbol{\eta}) = 0. \qquad (22)$$

On account of (14) and (22),

$$\alpha + \tilde{\mathbf{z}}^T \boldsymbol{\eta} = \frac{\sum_{i=1}^n \beta_i (\alpha + \mathbf{z}_i^T \boldsymbol{\eta})}{\sum_{i=1}^n \beta_i} = 0, \qquad (23)$$

and this immediately yields (17). To show that (16) also holds, note that, by (9) and (10), for each $i = 1, \ldots, n$,

$$\mathbf{u}(\mathbf{x}_i)^T \boldsymbol{\theta} = \alpha + \mathbf{z}_i^T \boldsymbol{\eta},$$

and, by (15) and (17),

$$\alpha + \mathbf{z}_i^T \boldsymbol{\eta} = (\mathbf{z}_i - \tilde{\mathbf{z}})^T \boldsymbol{\eta} = \mathbf{z}_i'^T \boldsymbol{\eta}.$$

Hence

$$\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta} = (\mathbf{u}(\mathbf{x}_i)^T \boldsymbol{\theta})^2 = (\mathbf{z}_i'^T \boldsymbol{\eta})^2.$$

This together with (18) implies that

$$\frac{\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}}{(\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta})^2} = (\beta_i \mathbf{z}_i'^T \boldsymbol{\eta})^2,$$

whence immediately $\mathbf{N}_{\boldsymbol{\theta}}^0 = \mathbf{N}_{\boldsymbol{\eta}}'$. With this identity, (21) can be rewritten as

$$\sum_{i=1}^n \beta_i (\alpha \mathbf{z}_i + \mathbf{z}_i \mathbf{z}_i^T \boldsymbol{\eta}) = \mathbf{N}_{\boldsymbol{\eta}}' \boldsymbol{\eta}. \qquad (24)$$

On the other hand, taking (15) into account, we see that

$$\sum_{i=1}^n \beta_i (\alpha \mathbf{z}_i + \mathbf{z}_i \mathbf{z}_i^T \boldsymbol{\eta}) = \sum_{i=1}^n \beta_i \mathbf{z}_i (\alpha + \mathbf{z}_i^T \boldsymbol{\eta}) = \sum_{i=1}^n \beta_i (\mathbf{z}_i' + \tilde{\mathbf{z}})(\alpha + \mathbf{z}_i^T \boldsymbol{\eta}). \qquad (25)$$

By (22),

$$\sum_{i=1}^n \beta_i \tilde{\mathbf{z}} (\alpha + \mathbf{z}_i^T \boldsymbol{\eta}) = \tilde{\mathbf{z}} \sum_{i=1}^n \beta_i (\alpha + \mathbf{z}_i^T \boldsymbol{\eta}) = \mathbf{0}, \qquad (26)$$

and by (15) and (23),

$$\sum_{i=1}^n \beta_i \mathbf{z}_i' (\alpha + \mathbf{z}_i^T \boldsymbol{\eta}) = \sum_{i=1}^n \beta_i \mathbf{z}_i' (\alpha + \tilde{\mathbf{z}}^T \boldsymbol{\eta} + \mathbf{z}_i'^T \boldsymbol{\eta}) = \sum_{i=1}^n \beta_i \mathbf{z}_i' \mathbf{z}_i'^T \boldsymbol{\eta} = \mathbf{M}_{\boldsymbol{\eta}}' \boldsymbol{\eta}. \qquad (27)$$

Combining (25), (26), and (27), we obtain

$$\sum_{i=1}^n \beta_i (\alpha \mathbf{z}_i + \mathbf{z}_i \mathbf{z}_i^T \boldsymbol{\eta}) = \mathbf{M}_{\boldsymbol{\eta}}' \boldsymbol{\eta}.$$

This jointly with (24) finally yields (16), as desired.

Working backward, one can easily infer that if $\boldsymbol{\eta}$ satisfies (16) and α is given by (17), then $\boldsymbol{\theta} = [\boldsymbol{\eta}^T, \alpha]^T$ satisfies (7).

5   REDUCED HEIV SCHEME

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

$$\widehat{\boldsymbol{\theta}}_{\text{ALS}} = \left[(\widehat{\boldsymbol{\eta}}_{\text{ALS}})^T, \widehat{\alpha}_{\text{ALS}}\right]^T.$$

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

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

In view of (17), $\widehat{\alpha}_{\text{AML}}^u$ is uniquely determined by $\widehat{\boldsymbol{\eta}}_{\text{AML}}^u$—when the centroid $\tilde{\mathbf{z}}$ is taken with the weights
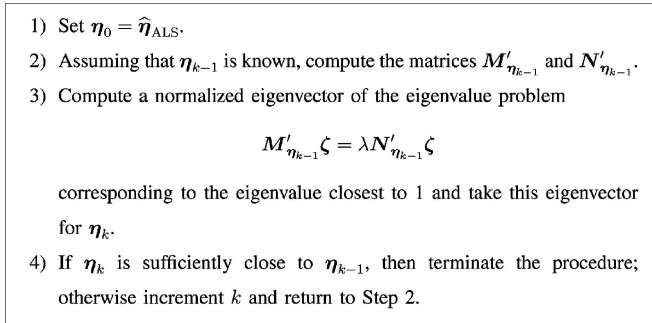
1) Set $\boldsymbol{\eta}_0 = \widehat{\boldsymbol{\eta}}_{\mathrm{ALS}}$.

2) Assuming that $\boldsymbol{\eta}_{k-1}$ is known, compute the matrices $\boldsymbol{M}'_{\boldsymbol{\eta}_{k-1}}$ and $\boldsymbol{N}'_{\boldsymbol{\eta}_{k-1}}$.

3) Compute a normalized eigenvector of the eigenvalue problem

$$\boldsymbol{M}'_{\boldsymbol{\eta}_{k-1}} \boldsymbol{\zeta} = \lambda \boldsymbol{N}'_{\boldsymbol{\eta}_{k-1}} \boldsymbol{\zeta}$$

corresponding to the eigenvalue closest to 1 and take this eigenvector for $\boldsymbol{\eta}_k$.

4) If $\boldsymbol{\eta}_k$ is sufficiently close to $\boldsymbol{\eta}_{k-1}$, then terminate the procedure; otherwise increment k and return to Step 2.

Fig. 3. Reduced HEIV scheme.

$$\beta_i = \frac{1}{(\widehat{\boldsymbol{\eta}}^u_{\mathrm{AML}})^T \boldsymbol{B}^0_i \widehat{\boldsymbol{\eta}}^u_{\mathrm{AML}}},$$

then $\widehat{\alpha}^u_{\mathrm{AML}} = -\tilde{\boldsymbol{z}}^T \widehat{\boldsymbol{\eta}}^u_{\mathrm{AML}}$. Now, the generalized eigenvalue problem

$$\boldsymbol{M}'_{\boldsymbol{\eta}} \boldsymbol{\zeta} = \lambda \boldsymbol{N}'_{\boldsymbol{\eta}} \boldsymbol{\zeta} \qquad (28)$$

is nondegenerate: the matrix $\boldsymbol{N}'_{\boldsymbol{\eta}}$ is positive definite. Accordingly, $\widehat{\boldsymbol{\eta}}^u_{\mathrm{AML}}$ can be determined with use of a simple modification of the HEIV algorithm. The steps of this *reduced* HEIV scheme are given in Fig. 3. It is essentially in this form that the HEIV algorithm was first advanced [7], [8]. The original version employs a slightly different, bias-corrected form of *the vector of carriers* $\boldsymbol{u}(\boldsymbol{x})$. The reduced scheme, based solely on $\boldsymbol{u}(\boldsymbol{x})$, 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 recognized as the fundamental form of the HEIV algorithm [9].

It is worth mentioning that Leedan and Meer [7] proposed a robust procedure for solving the eigenvalue problem (28) based upon *generalized singular value decomposition* of a pair of matrices $(\boldsymbol{S}_{\boldsymbol{\eta}}, \boldsymbol{T}_{\boldsymbol{\eta}})$ satisfying $\boldsymbol{M}'_{\boldsymbol{\eta}} = \boldsymbol{S}^T_{\boldsymbol{\eta}} \boldsymbol{S}_{\boldsymbol{\eta}}$ and $\boldsymbol{N}'_{\boldsymbol{\eta}} = \boldsymbol{T}^T_{\boldsymbol{\eta}} \boldsymbol{T}_{\boldsymbol{\eta}}$.

Finally, we remark that reduction to a nondegenerate form of eigenvalue problems similar to (28) is crucial for computing some other types of estimates, notably Kanatani-like renormalization estimates [12] [4, chapter 9] (these are approximates of $\widehat{\boldsymbol{\theta}}^u_{\mathrm{AML}}$ of some sort) and ellipse-specific estimates obtainable with use of an improved version of the direct least-squares fitting algorithm of Fitzgibbon et al. [3], due to Halír and Flusser [15].

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 (16). 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 ORIGINS OF HEIV

The original derivation of the HEIV algorithm, as given in [7], [8], [9], is different from the one presented here. In our exposition, the core of HEIV, namely its reduced form, results from reformulating the variational equation so that the associated generalized eigenvalue problem becomes nondegenerate. The original derivation is based on a direct application of the maximum likelihood principle to a statistical model operating with candidate prob-
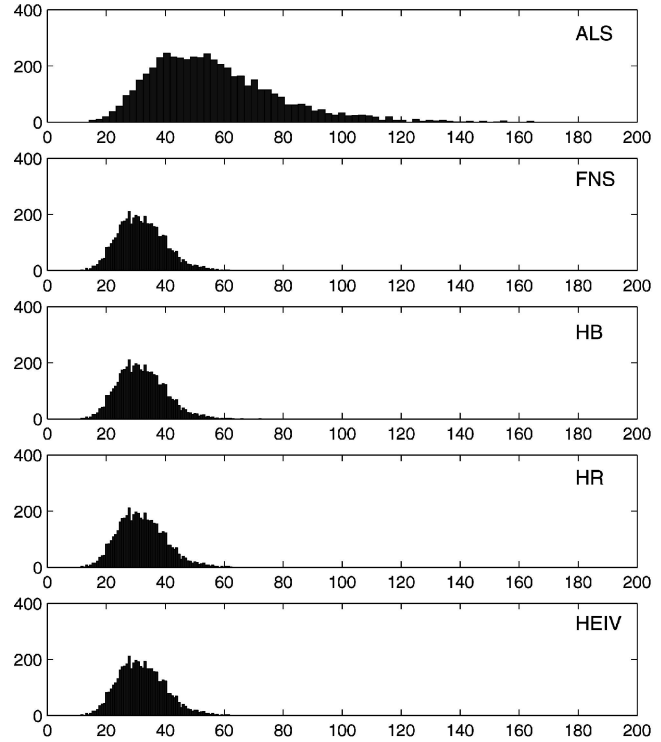
Fig. 4. Performance histograms for each of the methods, with J_{AML} bins on the x-axes, and frequencies of occurrence on the y-axes.

ability distributions for $\boldsymbol{z}(\mathbf{x}_1), \ldots, \boldsymbol{z}(\mathbf{x}_n)$, with the random variables $\mathbf{x}_1, \ldots \mathbf{x}_n$ modeling the image data $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$. When the $\mathbf{x}_i$ have equal variances (as is often assumed), the $\boldsymbol{z}(\mathbf{x}_i)$ form, as a rule, a *heteroscedastic* set of random variables, that is having different variances. This explains the term "heteroscedastic" in the name of the HEIV algorithm. The "errors-in-variables" part of the label alludes to the adopted statistical model being a so-called errors-in-variables model—the scalar components of each $\mathbf{x}_i$ are not segregated into two exclusive groups of explanatory (essentially nonrandom) and response (random) variables, and are all consistently treated as random variables.

8 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 [5].

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 give a sense of scale to the forthcoming numerical results. The FNS, HB, and HR methods were implemented as specified in Figs. 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 [16].

TABLE 1
Maximum and Average Differences in J_{AML} Values for Different Estimation Methods

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

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

Fig. 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 $\max_{1\le i\le N}|J_{\mathrm{AML}}(\widehat{\boldsymbol{\theta}}^i_{\mathrm{HEIV}}) - J_{\mathrm{AML}}(\widehat{\boldsymbol{\theta}}^i_{\mathrm{ALS}})|$ and $N^{-1}\sum_{i=1}^{N}|J_{\mathrm{AML}}(\widehat{\boldsymbol{\theta}}^i_{\mathrm{HEIV}}) - J_{\mathrm{AML}}(\widehat{\boldsymbol{\theta}}^i_{\mathrm{ALS}})|$. 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.

Some MATLAB code implementing aspects of the above may be found at [17].

9 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 interrelationships 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.

ACKNOWLEDGMENTS

REFERENCES

[1] O.D. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint.* Cambridge, Mass.: MIT Press, 1993.
[2] M.J. Brooks, W. Chojnacki, and L. Baumela, "Determining the Egomotion of an Uncalibrated Camera from Instantaneous Optical Flow," *J. Optical Soc. Am. A,* vol. 14, no. 10, pp. 2670-2677, 1997.
[3] A. Fitzgibbon, M. Pilu, and R.B. Fisher, "Direct Least Square Fitting of Ellipses," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 21, no. 5, pp. 476-480, May 1999.
[4] K. Kanatani, *Statistical Optimization for Geometric Computation: Theory and Practice.* Amsterdam: Elsevier, 1996.
[5] 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,* vol. 22, no. 11, pp. 1294-1303, Nov. 2000.
[6] M.J. Brooks, W. Chojnacki, D. Gawley, and A. van den Hengel, "What Value Covariance Information in Estimating Vision Parameters?" *Proc. Eighth Int'l Conf. Computer Vision,* vol. 1, pp. 302-308, 2001.
[7] Y. Leedan and P. Meer, "Heteroscedastic Regression in Computer Vision: Problems with Bilinear Constraint," *Int'l J. Computer Vision,* vol. 37, no. 2, pp. 127-150, 2000.
[8] B. Matei and P. Meer, "A General Method for Errors-in-Variables Problems in Computer Vision," *Proc. IEEE Conf. Computer Vision and Pattern Recognition,* vol. 2, pp. 18-25, 2000.
[9] B. Matei, "Heteroscedastic Errors-in-Variables Models in Computer Vision," PhD dissertation, Dept. Electrical and Computer Eng., Rutgers Univ., New Brunswick, N.J., May 2001, http://www.caip.rutgers.edu/riul/research/theses.html.
[10] 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," *Proc. Statistical Methods in Video Processing Workshop,* held in conjunction with ECCV2002, pp. 43-48, 2002.
[11] A. van den Hengel, W. Chojnacki, M.J. Brooks, and D. Gawley, "A New Constrained Parameter Estimator: Experiments in Fundamental Matrix Computation," *Proc. 13th British Machine Vision Conf.,* vol. 2, pp. 468-476, 2002.
[12] W. Chojnacki, M.J. Brooks, A. van den Hengel, "Rationalising the Renormalisation Method of Kanatani," *J. Math. Imaging and Vision,* vol. 14, no. 1, pp. 21-38, 2001.
[13] W. Chojnacki, M.J. Brooks, A. van den Hengel, and D. Gawley, "Revisiting Hartley's Normalized Eight-Point Algorithm," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 25, no. 9, pp. 1172-1177, Sept. 2003.
[14] B.N. Parlett, *The Symmetric Eigenvalue Problem.* Englewood Cliffs, N.J.: Prentice-Hall, 1980.
[15] R. Halír and J. Flusser, "Numerically Stable Direct Least Squares Fitting of Ellipses," *Proc. Sixth Int'l Conf. Computer Graphics and Visualization,* vol. 1, pp. 125-132, 1998, http://wscg.zcu.cz/wscg98/wscg98.htm.
[16] http://www.caip.rutgers.edu/riul/research/code.html. 2003.
[17] http://www.cs.adelaide.edu.au/~hengel/Vision/Code/index.html. 2003.