



FNS, CFNS and HEIV: A Unifying Approach

WOJCIECH CHOJNACKI, MICHAEL J. BROOKS, ANTON VAN DEN HENGEL AND DARREN GAWLEY
School of Computer Science, University of Adelaide, Adelaide, SA 5005, Australia

wojtek@cs.adelaide.edu.au

mjb@cs.adelaide.edu.au

hengel@cs.adelaide.edu.au

dg@cs.adelaide.edu.au

First online version published in September, 2005

Abstract. Estimation of parameters from image tokens is a central problem in computer vision. FNS, CFNS and HEIV are three recently developed methods for solving special but important cases of this problem. The schemes are means for finding unconstrained (FNS, HEIV) and constrained (CFNS) minimisers of cost functions. In earlier work of the authors, FNS, CFNS and a core version of HEIV were applied to a specific cost function. Here we extend the approach to more general cost functions. This allows the FNS, CFNS and HEIV methods to be placed within a common framework.

Keywords: statistical methods, maximum likelihood, (un)constrained minimisation, fundamental matrix, epipolar equation, conic fitting

1. Introduction

A common task in computer vision is the estimation of the parameters that describe a relationship between image feature locations. The estimation problem can often be reduced to minimising a cost function. FNS, CFNS and HEIV are three recently developed techniques for finding minimisers of cost functions underpinning a special but important class of estimation problems. FNS and HEIV aim to determine unconstrained minimisers, while CFNS seeks to isolate constrained minimisers. In earlier work of the authors [2–4], FNS, CFNS and a core version of HEIV were applied to a specific cost function. Here we extend the approach to more general cost functions including the cost function that pertains to the original version of HEIV [8] which is different from the core version. This allows the FNS, CFNS and HEIV methods to be placed within a common framework.

The rest of the paper develops as follows. We first derive a certain form for the optimal, maximum like-

lihood cost function that is appropriate for a class of estimation problems. We then evolve two approximations to this function. One of these is the cost function to which the standard versions of FNS and CFNS and the core version of HEIV apply. The other is the cost function recognised here as the function underlying the original version of HEIV. Both of these functions have a similar form and can be viewed as specialisations of a single model function. The subsequent development, based largely on a critical review of our earlier work, is concerned with advancing variants of FNS, HEIV and CFNS for this model function. Finally, we discuss the effects of applying the derived methods to the two approximated maximum likelihood functions.

2. Background

2.1. Estimation Problem

Relationships between image tokens can often be arranged into parametric models. This paper focuses on

models expressed by means of a *principal constraint* of the form

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

Here $\boldsymbol{\theta} = [\theta_1, \dots, \theta_l]^T$ is a vector of parameters describing a particular model; $\mathbf{x} = [x_1, \dots, x_k]^T$ is a vector that represents an ideal data point conforming to the model; and $\mathbf{u}(\mathbf{x}) = [u_1(\mathbf{x}), \dots, u_l(\mathbf{x})]^T$ is a vector with the ideal data point transformed so that: (i) each component $u_i(\mathbf{x})$ is a quadratic form in the compound vector $[\mathbf{x}^T, 1]^T$, (ii) the last component $u_l(\mathbf{x})$ is equal to 1. The set of all points \mathbf{x} satisfying (1) for a given $\boldsymbol{\theta}$ is a geometric representation of the model parametrised by $\boldsymbol{\theta}$ and is termed the *geometric primitive* associated with $\boldsymbol{\theta}$. In some cases, the parameters are subject to an *ancillary constraint* not involving model points. 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 non-zero scalar t .

Associated with (1) and (2) is the following estimation problem: Given a collection $\mathbf{x}_1, \dots, \mathbf{x}_n$ of *observed* data points and a meaningful *cost function* that characterises the extent to which any particular $\boldsymbol{\theta}$ fails to satisfy the system of copies of Eq. (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.

Example estimation problems of the above form include the estimation of the coefficients of the *epipolar equation* [5] and the *differential epipolar equation* [1], and *conic fitting* [6]. Each of the first two problems involves a separate ancillary *cubic* constraint, while the last problem involves no constraint.

2.2. The ML Cost Function

To evolve a useful cost function for our estimation problem, a model for the data generation process is required. Here we assume that the data are produced by choosing a geometric primitive, selecting points on it, and then perturbing these points using Gaussian noise. The observed data points $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ will be viewed as a sample drawn from a vector-valued random variable $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ whose distribution is not exactly specified but is an element of a set of candidate distributions. Each of these is labelled by an *extended* parameter

vector $\boldsymbol{\psi} = (\boldsymbol{\theta}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$, where $\boldsymbol{\theta}$ is the *principal* part of $\boldsymbol{\psi}$ representing a particular geometric primitive and $(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$ is the *subsidiary* part representing a specific selection of points on the primitive. The two components of $\boldsymbol{\psi}$ are such that $\boldsymbol{\theta} \neq \mathbf{0}$ and

$$\boldsymbol{\theta}^T \mathbf{u}(\bar{\mathbf{x}}_1) = \dots = \boldsymbol{\theta}^T \mathbf{u}(\bar{\mathbf{x}}_n) = 0. \quad (3)$$

With $(\mathbf{y}_1, \dots, \mathbf{y}_n)$ a member of the sample space, the distribution associated with a particular $\boldsymbol{\psi}$ is described by the probability density function

$$\begin{aligned} p(\mathbf{y}_1, \dots, \mathbf{y}_n | \boldsymbol{\psi}) \\ = (2\pi)^{-kn/2} \prod_{i=1}^n (\det \Lambda_{\mathbf{x}_i})^{-1/2} \\ \times \exp \left\{ -\frac{1}{2} \sum_{i=1}^n (\mathbf{y}_i - \bar{\mathbf{x}}_i)^T \Lambda_{\bar{\mathbf{x}}_i}^{-1} (\mathbf{y}_i - \bar{\mathbf{x}}_i) \right\}. \end{aligned}$$

Here, for each $i = 1, \dots, n$, $\Lambda_{\mathbf{x}_i}$ is a $k \times k$ symmetric *covariance matrix*, assumed to be known, that quantifies errors in the measurement of the data point \mathbf{x}_i . To identify the most desirable parameters given the data, we apply the strategy of *maximum likelihood inference*. This reflects the lack of any special preferences as to how a particular geometric primitive and points on it have been chosen. The preferred extended parameter vector will be the vector that makes the observed data as likely as possible. This *maximum likelihood estimate* of $\boldsymbol{\psi}$, $\hat{\boldsymbol{\psi}}_{\text{ML}}$, is formally defined as the maximiser of the *likelihood function* $\boldsymbol{\psi} \mapsto p(\mathbf{x}_1, \dots, \mathbf{x}_n | \boldsymbol{\psi})$. The likelihood function has such a form that $\hat{\boldsymbol{\psi}}_{\text{ML}}$ can alternatively be characterised as the extended parameter vector whose subsidiary part minimises the squared *Mahalanobis distance* from the data. The squared Mahalanobis distance between the data points $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ and model points $(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$ is given by

$$\begin{aligned} d_{\text{Mahal}}^2(\mathbf{x}_1, \dots, \mathbf{x}_n; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n) \\ = \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}}_i)^T \Lambda_{\bar{\mathbf{x}}_i}^{-1} (\mathbf{x}_i - \bar{\mathbf{x}}_i). \end{aligned}$$

For each $\boldsymbol{\theta} \neq \mathbf{0}$, when restricted to the set of those $(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$ for which (3) holds, the function

$$\begin{aligned} d_{\text{Mahal}}^2(\mathbf{x}_1, \dots, \mathbf{x}_n; \bullet): (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n) \\ \mapsto d_{\text{Mahal}}^2(\mathbf{x}_1, \dots, \mathbf{x}_n; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n) \end{aligned}$$

attains a constrained minimum at some point $(\bar{\mathbf{x}}_1^\theta, \dots, \bar{\mathbf{x}}_n^\theta)$. All these minima can be assembled into

a cost function by setting

$$J_{\text{ML}}(\boldsymbol{\theta}) = d_{\text{Mahal}}^2(\mathbf{x}_1, \dots, \mathbf{x}_n; \bar{\mathbf{x}}_1^\theta, \dots, \bar{\mathbf{x}}_n^\theta). \quad (4)$$

Let $\hat{\boldsymbol{\theta}}_{\text{ML}}$ be the minimiser of J_{ML} . Clearly, $J_{\text{ML}}(\hat{\boldsymbol{\theta}}_{\text{ML}})$ is the smallest of all the values that $d_{\text{Mahal}}^2(\mathbf{x}_1, \dots, \mathbf{x}_n; \bullet)$ attains at the subsidiary parts of extended parameter vectors. Consequently, $\hat{\boldsymbol{\theta}}_{\text{ML}}$ and $(\bar{\mathbf{x}}_1^{\hat{\boldsymbol{\theta}}_{\text{ML}}}, \dots, \bar{\mathbf{x}}_n^{\hat{\boldsymbol{\theta}}_{\text{ML}}})$ are the principal and subsidiary parts of $\hat{\boldsymbol{\psi}}_{\text{ML}}$. This justifies calling J_{ML} the maximum likelihood cost function for $\boldsymbol{\theta}$ -estimation, and $\hat{\boldsymbol{\theta}}_{\text{ML}}$ the maximum likelihood estimate of $\boldsymbol{\theta}$. It goes without saying that $\hat{\boldsymbol{\theta}}_{\text{ML}}$ is the preferred estimate of $\boldsymbol{\theta}$.

3. Approximating the ML Cost Function

3.1. The ML Cost Function: An Alternate Form

Finding $(\bar{\mathbf{x}}_1^\theta, \dots, \bar{\mathbf{x}}_n^\theta)$ for each $\boldsymbol{\theta}$ is a daunting task, and so direct minimisation of J_{ML} is rather impractical. A more feasible approach is to seek to minimise an appropriate approximation of J_{ML} that captures near-optimality. With a view to evolving a couple of such approximations, we now derive an alternative formula for J_{ML} .

Applied to the constrained minimiser $(\bar{\mathbf{x}}_1^\theta, \dots, \bar{\mathbf{x}}_n^\theta)$, the method of Lagrange Multipliers implies that, for each $i = 1, \dots, n$, the gradient (the column vector of the partial derivatives) of $(\mathbf{x}_i - \mathbf{y})^T \boldsymbol{\Lambda}_{\mathbf{x}_i}^{-1}(\mathbf{x}_i - \mathbf{y})$ with respect to \mathbf{y} is proportional to the gradient of $\boldsymbol{\theta}^T \mathbf{u}(\mathbf{y})$ with respect to \mathbf{y} , provided that both these gradients are evaluated at $\bar{\mathbf{x}}_i^\theta$. Since the first gradient is equal to $-2\boldsymbol{\Lambda}_{\mathbf{x}_i}^{-1}(\mathbf{x}_i - \mathbf{y})$ and the second is equal to $\partial_{\mathbf{x}} \mathbf{u}(\mathbf{y})^T \boldsymbol{\theta}$, where

$$\partial_{\mathbf{x}} \mathbf{u}(\mathbf{y}) = [(\partial u_i / \partial x_j)(\mathbf{y})]_{1 \leq i \leq l, 1 \leq j \leq k}$$

is the Jacobian matrix of \mathbf{u} at \mathbf{y} , it follows that

$$\boldsymbol{\Lambda}_{\mathbf{x}_i}^{-1}(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta) = \lambda_i \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta} \quad (5)$$

for some scalar λ_i . Multiplying both sides of this equation by $\boldsymbol{\Lambda}_{\mathbf{x}_i}$ and then multiplying both sides of the resulting equation by $\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)$, we find that

$$\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta) = \lambda_i \boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta}.$$

Hence

$$\lambda_i = \frac{(\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta})^{-1}}{\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)}. \quad (6)$$

On the other hand, multiplying both sides of (5) by $(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)^T$, we obtain

$$(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\Lambda}_{\mathbf{x}_i}^{-1}(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta) = \lambda_i (\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta}.$$

Substituting the value of λ_i from (6) into this equation and taking into account that

$$\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta) = (\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta}$$

(both sides of the above identity represent one and the same scalar), we conclude that

$$\begin{aligned} & (\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\Lambda}_{\mathbf{x}_i}^{-1}(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta) \\ &= (\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta})^{-1} \\ & \times (\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta))^2. \end{aligned} \quad (7)$$

For each $\gamma = 1, \dots, l$, $u_\gamma(\mathbf{x})$ is a quadratic function in \mathbf{x} , and so

$$\partial_{\mathbf{xx}}^2 u_\gamma(\mathbf{y}) = [(\partial^2 u_\gamma / \partial x_i \partial x_j)(\mathbf{y})]_{1 \leq i, j \leq k},$$

the Hessian matrix of u_γ at \mathbf{y} , is independent of \mathbf{y} , and the partial derivatives of u_γ of order higher than two all vanish. Denote by \mathbf{H}_γ the unique value of $\partial_{\mathbf{xx}}^2 u_\gamma$. Taylor expanding u_γ about any particular \mathbf{y} gives

$$\begin{aligned} u_\gamma(\mathbf{x}) &= u_\gamma(\mathbf{y}) + \partial_{\mathbf{x}} u_\gamma(\mathbf{y})(\mathbf{x} - \mathbf{y}) \\ &+ \frac{1}{2}(\mathbf{x} - \mathbf{y})^T \mathbf{H}_\gamma(\mathbf{x} - \mathbf{y}). \end{aligned} \quad (8)$$

Let

$$\mu_\gamma(\mathbf{x}, \mathbf{y}) = \frac{1}{2}(\mathbf{x} - \mathbf{y})^T \mathbf{H}_\gamma(\mathbf{x} - \mathbf{y})$$

and

$$\boldsymbol{\mu}(\mathbf{x}, \mathbf{y}) = [\mu_1(\mathbf{x}, \mathbf{y}), \dots, \mu_l(\mathbf{x}, \mathbf{y})]^T.$$

Setting $\mathbf{y} = \bar{\mathbf{x}}_i^\theta$ in (8) and taking into account that $\boldsymbol{\theta}^T \mathbf{u}(\bar{\mathbf{x}}_i^\theta) = 0$ yields

$$\boldsymbol{\theta}^T (\mathbf{u}(\mathbf{x}_i) - \boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)) = \boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta).$$

With this equality, Eq. (7) can be written as

$$(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\Lambda}_{\mathbf{x}_i}^{-1}(\mathbf{x}_i - \bar{\mathbf{x}}_i^\theta) = \frac{(\boldsymbol{\theta}^T (\mathbf{u}(\mathbf{x}_i) - \boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)))^2}{\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta}}.$$

This immediately leads to the following alternative expression for J_{ML} :

$$J_{\text{ML}}(\boldsymbol{\theta}) = \sum_{i=1}^n \frac{(\boldsymbol{\theta}^T (\mathbf{u}(\mathbf{x}_i) - \boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)))^2}{\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)^T \boldsymbol{\theta}}. \quad (9)$$

3.2. Two AML Cost Functions

At first glance, the formula (9) looks much more complicated than the formula (4). Despite its apparent complexity, the former is ideally suited for generating approximations of J_{ML} . We now derive two such approximations. In both of them $\partial_{\mathbf{x}} \mathbf{u}(\bar{\mathbf{x}}_i^\theta)$ will be replaced by $\partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i)$. In addition, one approximation will set $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)$ to zero and the other will replace $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)$ by the expected value of $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i)$, where the underlying distribution of (x_1, \dots, x_n) is any distribution labelled by an extended parameter of the form $(\boldsymbol{\theta}; \bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$. Here, critically, the expected value will not depend on any particular choice of the subsidiary part $(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$ insofar as (3) is satisfied. Thus in essence the first approach is to take a first-order approximation and discard $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)$ as an irrelevant second-order term, whereas the second approach is to retain a remnant of $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)$ in the form of an average value of some kind.

To verify that the expected value of $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i)$ is independent of the subsidiary part $(\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_n)$ satisfying (3), note that

$$\begin{aligned} \mu_\gamma(\mathbf{x}_i, \bar{\mathbf{x}}_i) &= \frac{1}{2} \text{tr}((\mathbf{x}_i - \bar{\mathbf{x}}_i)^T \mathbf{H}_\gamma (\mathbf{x}_i - \bar{\mathbf{x}}_i)) \\ &= \frac{1}{2} \text{tr}(\mathbf{H}_\gamma (\mathbf{x}_i - \bar{\mathbf{x}}_i) (\mathbf{x}_i - \bar{\mathbf{x}}_i)^T). \end{aligned}$$

Here tr denotes trace. Hence, immediately,

$$\begin{aligned} \mathbb{E}[\mu_\gamma(\mathbf{x}_i, \bar{\mathbf{x}}_i)] &= \frac{1}{2} \text{tr}(\mathbf{H}_\gamma \mathbb{E}[(\mathbf{x}_i - \bar{\mathbf{x}}_i) (\mathbf{x}_i - \bar{\mathbf{x}}_i)^T]) \\ &= \frac{1}{2} \text{tr}(\mathbf{H}_\gamma \boldsymbol{\Lambda}_{\mathbf{x}_i}), \end{aligned}$$

where \mathbb{E} denotes expectation. At a stroke, this establishes the desired result and shows that what the second approximation takes instead of $\boldsymbol{\mu}(\mathbf{x}_i, \bar{\mathbf{x}}_i^\theta)$ is the i th *second-order correction* $\boldsymbol{\mu}(\mathbf{x}_i) = [\mu_1(\mathbf{x}_i), \dots, \mu_l(\mathbf{x}_i)]^T$ defined by

$$\mu_\gamma(\mathbf{x}) = \frac{1}{2} \text{tr}(\mathbf{H}_\gamma \boldsymbol{\Lambda}_{\mathbf{x}}) \quad (1 \leq \gamma \leq l).$$

Note that the last component of $\boldsymbol{\mu}(\mathbf{x})$, $\mu_l(\mathbf{x})$, is null, since $u_l(\mathbf{x}) = 1$ and, consequently, $\mathbf{H}_l = \mathbf{0}$.

It is now obvious that the first approximation of J_{ML} takes the form

$$J_{\text{AML}}^{(1)}(\boldsymbol{\theta}) = \sum_{i=1}^n \frac{(\boldsymbol{\theta}^T \mathbf{u}(\mathbf{x}_i))^2}{\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i)^T \boldsymbol{\theta}},$$

whereas the second approximation is given by

$$J_{\text{AML}}^{(2)}(\boldsymbol{\theta}) = \sum_{i=1}^n \frac{(\boldsymbol{\theta}^T (\mathbf{u}(\mathbf{x}_i) - \boldsymbol{\mu}_i))^2}{\boldsymbol{\theta}^T \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i)^T \boldsymbol{\theta}}.$$

For each $i = 1, \dots, n$, let $\mathbf{v}_1(\mathbf{x}_i) = \mathbf{u}(\mathbf{x}_i)$, $\mathbf{v}_2(\mathbf{x}_i) = \mathbf{u}(\mathbf{x}_i) - \boldsymbol{\mu}(\mathbf{x}_i)$, $\mathbf{A}_i^{(\alpha)} = \mathbf{v}_\alpha(\mathbf{x}_i) \mathbf{v}_\alpha(\mathbf{x}_i)^T$ ($\alpha = 1, 2$), and $\mathbf{B}_i = \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{u}(\mathbf{x}_i)^T$. With this notation, $J_{\text{AML}}^{(\alpha)}$ can be simply written as

$$J_{\text{AML}}^{(\alpha)}(\boldsymbol{\theta}) = \sum_{i=1}^n \frac{\boldsymbol{\theta}^T \mathbf{A}_i^{(\alpha)} \boldsymbol{\theta}}{\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta}} \quad (\alpha = 1, 2).$$

Note that, since $u_l(\mathbf{x}) = 1$ and $\mu_l(\mathbf{x}) = 0$, the last component of $\mathbf{v}_\alpha(\mathbf{x})$, $v_{\alpha,l}(\mathbf{x})$, equals 1.

4. Model Cost Function

The functions $J_{\text{AML}}^{(1)}$ and $J_{\text{AML}}^{(2)}$ have a similar structure and can be subsumed into a single model function

$$J_{\text{AML}}(\boldsymbol{\theta}) = \sum_{i=1}^n \frac{\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}}{\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta}},$$

where all the \mathbf{A}_i and \mathbf{B}_i are non-negative definite $l \times l$ matrices. By convention, J_{AML} will be referred to as the *approximated maximum likelihood* cost function. The unconstrained minimiser of J_{AML} will be denoted $\hat{\boldsymbol{\theta}}_{\text{AML}}^u$, and if an ancillary constraint (as per (2)) applies, the constrained minimiser of J_{AML} will be denoted $\hat{\boldsymbol{\theta}}_{\text{AML}}$. The function J_{AML} is homogeneous of degree zero and the zero set of the function ϕ defining the ancillary constraint 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. We shall mainly consider J_{AML} with the \mathbf{A}_i and \mathbf{B}_i such that

$$\mathbf{A}_i = \mathbf{v}(\mathbf{x}_i) \mathbf{v}(\mathbf{x}_i)^T, \quad (10)$$

$$\mathbf{B}_i = \partial_{\mathbf{x}} \mathbf{v}(\mathbf{x}_i) \boldsymbol{\Lambda}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{v}(\mathbf{x}_i)^T \quad (11)$$

for some $\mathbf{v}(\mathbf{x}) = [v_1(\mathbf{x}), \dots, v_l(\mathbf{x})]^T$ with $v_l(\mathbf{x}) = 1$. Note that with $\mathbf{v}(\mathbf{x}) = \mathbf{v}_\alpha(\mathbf{x})$, $J_{\text{AML}}^{(\alpha)}$ recovers J_{AML} .

5. Unconstrained Minimisation

5.1. Variational Equation

The unconstrained minimiser $\widehat{\boldsymbol{\theta}}_{\text{AML}}^u$ satisfies the *variational equation*

$$[\partial_{\boldsymbol{\theta}} J_{\text{AML}}(\boldsymbol{\theta})]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{\text{AML}}^u} = \mathbf{0}^T \quad (12)$$

with $\partial_{\boldsymbol{\theta}} J_{\text{AML}}$ the row vector of the partial derivatives of J_{AML} with respect to $\boldsymbol{\theta}$. It is readily verified that

$$\partial_{\boldsymbol{\theta}} J_{\text{AML}}(\boldsymbol{\theta})^T = 2\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}, \quad (13)$$

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

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

Thus (12) can be rewritten as

$$[\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{\text{AML}}^u} = \mathbf{0}. \quad (14)$$

The latter equation provides the basis for isolating $\widehat{\boldsymbol{\theta}}_{\text{AML}}^u$.

There are two fundamental methods for solving (14). One is the *fundamental numerical scheme* (FNS) introduced by Chojnacki et al. [2]. Another is the *heteroscedastic errors-in-variables* (HEIV) *scheme* that was first proposed by Leedan and Meer [8] and further developed by Matei and Meer [9, 10].

5.2. Fundamental Numerical Scheme

A vector $\boldsymbol{\theta}$ satisfies (14) if and only if it is a solution of the *ordinary* eigenvalue problem

$$\mathbf{X}_{\boldsymbol{\theta}}\boldsymbol{\xi} = \lambda\boldsymbol{\xi} \quad (15)$$

corresponding to the eigenvalue $\lambda = 0$. This suggests an iterative method for solving (14) whereby if $\boldsymbol{\theta}_c$ is a current approximate solution, then an updated solution $\boldsymbol{\theta}_+$ is a vector chosen from that eigenspace of $\mathbf{X}_{\boldsymbol{\theta}_c}$ 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 process can be started by computing the *algebraic least squares* (ALS) *estimate*, $\widehat{\boldsymbol{\theta}}_{\text{ALS}}$, defined as the unconstrained minimiser of the cost function $J_{\text{ALS}}(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|^{-2} \sum_{i=1}^n \boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}$, with $\|\boldsymbol{\theta}\| = (\sum_{j=1}^l \theta_j^2)^{1/2}$. The

estimate $\widehat{\boldsymbol{\theta}}_{\text{ALS}}$ coincides, up to scale, with an eigenvector of $\sum_{i=1}^n \mathbf{A}_i$ associated with the smallest eigenvalue. When the \mathbf{A}_i satisfy (10), this eigenvector can be found by performing *singular-value decomposition* on the matrix $[\mathbf{v}(\mathbf{x}_1), \dots, \mathbf{v}(\mathbf{x}_l)]^T$. The overall procedure is summarised in Algorithm 1.

Algorithm 1. Fundamental numerical scheme

1. Set $\boldsymbol{\theta}$ to $\widehat{\boldsymbol{\theta}}_{\text{ALS}}$.
2. Repeat:
 - (a) Compute the matrix $\mathbf{X}_{\boldsymbol{\theta}}$;
 - (b) Compute a normalised eigenvector of $\mathbf{X}_{\boldsymbol{\theta}}$ corresponding to the eigenvalue closest to zero (in absolute value);
 - (c) Take the computed eigenvector for an update of $\boldsymbol{\theta}$;
 until convergence.

5.3. HEIV: A Basic Form

Given the representation $\mathbf{X}_{\boldsymbol{\theta}} = \mathbf{M}_{\boldsymbol{\theta}} - \mathbf{N}_{\boldsymbol{\theta}}$, where $\mathbf{M}_{\boldsymbol{\theta}} = \sum_{i=1}^n (\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta})^{-1} \mathbf{A}_i$ and $\mathbf{N}_{\boldsymbol{\theta}} = \sum_{i=1}^n (\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}) (\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta})^{-2} \mathbf{B}_i$, the variational equation (14) can be restated as

$$\mathbf{M}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{N}_{\boldsymbol{\theta}}\boldsymbol{\theta}, \quad (16)$$

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

$$\mathbf{M}_{\boldsymbol{\theta}}\boldsymbol{\xi} = \lambda\mathbf{N}_{\boldsymbol{\theta}}\boldsymbol{\xi} \quad (17)$$

corresponding to the eigenvalue $\lambda = 1$. The heteroscedastic errors-in-variables scheme in *basic* form, or *HEIV with intercept* [9], exploits the above eigenvalue problem in a manner analogous to that in which FNS utilises the eigenvalue problem (15). The details are given in Algorithm 2.

Algorithm 2. Basic HEIV scheme

1. Set $\boldsymbol{\theta}$ to $\widehat{\boldsymbol{\theta}}_{\text{ALS}}$.
2. Repeat:
 - (a) Compute the matrices $\mathbf{M}_{\boldsymbol{\theta}}$ and $\mathbf{N}_{\boldsymbol{\theta}}$;

- (b) Compute a normalised eigenvector of the eigenvalue problem

$$\mathbf{M}_\theta \xi = \lambda \mathbf{N}_\theta \xi$$

corresponding to the eigenvalue closest to 1;

- (c) Take the computed eigenvector for an update of θ ;

until convergence.

5.4. Reduced Variational Equation

If the \mathbf{A}_i and \mathbf{B}_i satisfy (10) and (11), respectively, for some $\mathbf{v}(\mathbf{x}) = [\mathbf{z}(\mathbf{x})^T, 1]^T$, where $\mathbf{z}(\mathbf{x})$ is a vector of length $l - 1$, then the variational equation can be re-expressed as a system of equations. To see how this can be done, first partition the parameter vector as $\theta = [\eta^T, \alpha]^T$ with η a length $l - 1$ vector and α a scalar. Further, let $\bar{\mathbf{z}} = (\sum_{i=1}^n \beta_i)^{-1} \sum_{i=1}^n \beta_i \mathbf{z}_i$ with $\beta_i = (\eta^T \mathbf{B}_i^0 \eta)^{-1}$ and $\mathbf{B}_i^0 = \partial_{\mathbf{x}} \mathbf{z}(\mathbf{x}_i) \mathbf{A}_{\mathbf{x}_i} \partial_{\mathbf{x}} \mathbf{z}(\mathbf{x}_i)^T$, and let $\mathbf{z}'_i = \mathbf{z}_i - \bar{\mathbf{z}}$ for each $i = 1, \dots, n$. Finally, define two $(l - 1) \times (l - 1)$ matrices $\mathbf{M}'_\eta = \sum_{i=1}^n \beta_i \mathbf{z}'_i \mathbf{z}'_i{}^T$ and $\mathbf{N}'_\eta = \sum_{i=1}^n (\beta_i \mathbf{z}'_i{}^T \eta)^2 \mathbf{B}_i^0$. A fundamental result that can now be established [3] is that $\theta = [\eta^T, \alpha]^T$ satisfies (16) if and only if the following system of equations holds:

$$\mathbf{M}'_\eta \eta = \mathbf{N}'_\eta \eta, \quad (18)$$

$$\alpha = -\bar{\mathbf{z}}^T \eta. \quad (19)$$

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 is called the *reduced variational equation*. A key feature of this equation is that its right-hand side matrix \mathbf{N}'_η is generically *positive definite* if $n \geq l$. In contrast, \mathbf{N}_θ is singular, since all the \mathbf{B}_i have the length l vector $[0, \dots, 0, 1]^T$ in their respective null spaces.

5.5. HEIV: A Reduced Form

Define the algebraic least squares estimates $\hat{\eta}_{\text{ALS}}$ and $\hat{\alpha}_{\text{ALS}}$ as the respective components in the representation $\hat{\theta}_{\text{ALS}} = [\hat{\eta}_{\text{ALS}}^T, \hat{\alpha}_{\text{ALS}}]^T$. Analogously, define the unconstrained approximated maximum likelihood estimates $\hat{\eta}_{\text{AML}}^u$ and $\hat{\alpha}_{\text{AML}}^u$ via the decomposition $\hat{\theta}_{\text{AML}}^u = [\hat{\eta}_{\text{AML}}^u{}^T, \hat{\alpha}_{\text{AML}}^u]^T$. In view of (19), $\hat{\alpha}_{\text{AML}}^u$ is

uniquely determined by $\hat{\eta}_{\text{AML}}^u$ —taking $\bar{\mathbf{z}}$ with the $\beta_i = (\hat{\eta}_{\text{AML}}^u{}^T \mathbf{B}_i^0 \hat{\eta}_{\text{AML}}^u)^{-1}$ results in $\hat{\alpha}_{\text{AML}}^u = -\bar{\mathbf{z}}^T \hat{\eta}_{\text{AML}}^u$. Now, the matrix \mathbf{N}'_η is generically positive-definite, and so the generalised eigenvalue problem $\mathbf{M}'_\eta \zeta = \lambda \mathbf{N}'_\eta \zeta$ is non-degenerate. Accordingly, $\hat{\eta}_{\text{AML}}^u$ can be determined via a simple modification of the HEIV algorithm. The steps of this HEIV scheme in *reduced* form, or *HEIV without intercept* [9], are given in Algorithm 3.

Algorithm 3. Reduced HEIV scheme

1. Set η to $\hat{\eta}_{\text{ALS}}$.
2. Repeat:
 - (a) Compute the matrices \mathbf{M}'_η and \mathbf{N}'_η ;
 - (b) Compute a normalised eigenvector of the eigenvalue problem

$$\mathbf{M}'_\eta \zeta = \lambda \mathbf{N}'_\eta \zeta$$

corresponding to the eigenvalue closest to 1 and take this eigenvector for η ;

- (c) Take the computed eigenvector for an update of η ;

until convergence.

5.6. Stable Versions

FNS and the HEIV schemes are locally convergent—to work they require an initial estimate sufficiently close to a solution of the equation underlying a particular method. More stable versions of the algorithms, able to cope with less accurate initial estimates, result from selecting at each iteration the eigenvector corresponding to the smallest eigenvalue instead of the eigenvector corresponding to the eigenvalue closest to 0 in the case of FNS, and closest to 1 in the case of the HEIV schemes. The HEIV method thus modified has been shown experimentally to converge successfully (in fact with high convergence rate) even when seeded with a random initial estimate [8]. Typically, the minimal eigenvalue computed after a few iterations is also the closest to 0 or 1 depending on the method, and once this stage is reached the modified algorithms act exactly as their original versions. Without these modifications, the schemes may exhibit slow convergence or even divergence.

6. Constrained Minimisation

6.1. Variational System

Applied to the constrained minimiser $\widehat{\boldsymbol{\theta}}_{\text{AML}}$, the method of Lagrange Multipliers yields

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

where λ is scalar. When properly combined with the identity $\partial_{\boldsymbol{\theta}} \phi(\boldsymbol{\theta}) \boldsymbol{\theta} = \kappa \phi(\boldsymbol{\theta})$ obtained by differentiating (2) with respect to t and evaluating at $t = 1$, this *variational system* can be converted into a single equation similar to (14). Of many equivalent forms, the one useful to us reads

$$[\mathbf{Q}_{\boldsymbol{\theta}} \boldsymbol{\theta}]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{\text{AML}}} = \mathbf{0}, \quad (20)$$

where $\mathbf{Q}_{\boldsymbol{\theta}} = \mathbf{Z}_{\boldsymbol{\theta}}^T \mathbf{Z}_{\boldsymbol{\theta}}$ and $\mathbf{Z}_{\boldsymbol{\theta}}$ is an $l \times l$ matrix defined as follows. Let $\mathbf{P}_{\boldsymbol{\theta}} = \mathbf{I}_l - \|\mathbf{a}_{\boldsymbol{\theta}}\|^{-2} \mathbf{a}_{\boldsymbol{\theta}} \mathbf{a}_{\boldsymbol{\theta}}^T$, where \mathbf{I}_l denotes the $l \times l$ identity matrix and $\mathbf{a}_{\boldsymbol{\theta}} = \partial_{\boldsymbol{\theta}} \phi(\boldsymbol{\theta})^T / 2$. Denote by $\mathbf{H}_{\boldsymbol{\theta}}$ the Hessian of J_{AML} at $\boldsymbol{\theta}$; more explicitly, $\mathbf{H}_{\boldsymbol{\theta}} = 2(\mathbf{X}_{\boldsymbol{\theta}} - \mathbf{T}_{\boldsymbol{\theta}})$, where

$$\begin{aligned} \mathbf{T}_{\boldsymbol{\theta}} &= \sum_{i=1}^n \frac{2}{(\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta})^2} \\ &\quad \times \left[\mathbf{A}_i \boldsymbol{\theta} \boldsymbol{\theta}^T \mathbf{B}_i + \mathbf{B}_i \boldsymbol{\theta} \boldsymbol{\theta}^T \mathbf{A}_i - 2 \frac{\boldsymbol{\theta}^T \mathbf{A}_i \boldsymbol{\theta}}{\boldsymbol{\theta}^T \mathbf{B}_i \boldsymbol{\theta}} \mathbf{B}_i \boldsymbol{\theta} \boldsymbol{\theta}^T \mathbf{B}_i \right]. \end{aligned}$$

Let $\Phi_{\boldsymbol{\theta}}$ be the Hessian of ϕ at $\boldsymbol{\theta}$. For each $i \in \{1, \dots, l\}$, let \mathbf{e}_i be the length l vector whose i th entry is unital and all other entries are zero. With all the preparations now completed, we let $\mathbf{Z}_{\boldsymbol{\theta}} = \mathbf{A}_{\boldsymbol{\theta}} + \mathbf{B}_{\boldsymbol{\theta}} + \mathbf{C}_{\boldsymbol{\theta}}$, where

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

Here, let us recall, κ is the degree of homogeneity of ϕ . 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.

6.2. Constrained Fundamental Numerical Scheme

Letting $\mathbf{Q}_{\boldsymbol{\theta}}$ play the role of $\mathbf{X}_{\boldsymbol{\theta}}$, one can advance an algorithm fully analogous to FNS [4]. The steps of the resulting *constrained fundamental numerical scheme* (CFNS) are given in Algorithm 4.

Algorithm 4. Constrained fundamental numerical scheme

1. Set $\boldsymbol{\theta}$ to $\widehat{\boldsymbol{\theta}}_{\text{ALS}}$.
 2. Repeat:
 - (a) Compute the matrix $\mathbf{Q}_{\boldsymbol{\theta}}$;
 - (b) Compute a normalised eigenvector of $\mathbf{Q}_{\boldsymbol{\theta}}$ corresponding to the eigenvalue closest to zero (in absolute value);
 - (c) Take the computed eigenvector for an update of $\boldsymbol{\theta}$;
- until convergence.

For CFNS to converge to a vector $\boldsymbol{\theta}^*$ solving (20), the zero eigenvalue of $\mathbf{Q}_{\boldsymbol{\theta}^*}$ must be simple, i.e., the null space of $\mathbf{Q}_{\boldsymbol{\theta}^*}$ must 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 $\boldsymbol{\theta}^*$ produces updates quickly converging to $\boldsymbol{\theta}^*$. In practice it is required that, for each iterate $\boldsymbol{\theta}_c$, the smallest (non-negative) eigenvalue of $\mathbf{Q}_{\boldsymbol{\theta}_c}$ 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 CFNS, 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 [12].

Interestingly, many other, often simpler, equivalent forms of (20) like

$$[\mathbf{Y}_{\boldsymbol{\theta}} \boldsymbol{\theta}]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{\text{AML}}} = \mathbf{0} \quad \text{with} \quad \mathbf{Y}_{\boldsymbol{\theta}} = \|\boldsymbol{\theta}\|^2 \mathbf{P}_{\boldsymbol{\theta}} \mathbf{X}_{\boldsymbol{\theta}} \mathbf{P}_{\boldsymbol{\theta}} + \mathbf{I}_l - \mathbf{P}_{\boldsymbol{\theta}}$$

lead to non-converging algorithms, with divergence occurring irrespective of the distance of the initial estimate from the desired limit. This reflects the rather

complicated behaviour of the function that sends a symmetric matrix to the eigenspace corresponding to the eigenvalue closest to zero.

7. Discussion and Conclusion

With the general formulation of various algorithms finally accomplished, we proceed to discuss implications for the AML cost functions.

The function $J_{\text{AML}}^{(1)}$ was first proposed by Kanatani [7] as a cost function capturing *geometric fitting*. An important precursor of $J_{\text{AML}}^{(1)}$ was Sampson's [11] cost function for some form of *orthogonal regression*. FNS and CFNS were introduced by the authors [2, 4] to perform unconstrained and constrained minimisation of $J_{\text{AML}}^{(1)}$. Later it was recognised that the core version of HEIV that was first adopted in [9] computes—in both basic and reduced forms—the unconstrained minimiser of $J_{\text{AML}}^{(1)}$ [3].

The function $J_{\text{AML}}^{(2)}$ introduced here is a second-order variation of $J_{\text{AML}}^{(1)}$. In some cases, like that of fundamental matrix estimation, the two functions coincide. Generally, they are different, as exemplified by the problem of estimating the coefficients of the differential epipolar equation. The significance of $J_{\text{AML}}^{(2)}$ is that it allows the original version of HEIV [8] to be placed within the operational framework of FNS, CFNS and the core version of HEIV. An inspection of Eqs. (22), (23), (24) in [8] reveals that the system of equations describing the estimate produced by the HEIV scheme is equivalent to the system comprising (18) and (19) for computing the unconstrained minimiser of $J_{\text{AML}}^{(2)}$. Thus the original form of HEIV turns out to be identical to the reduced HEIV scheme for computing the unconstrained minimiser of $J_{\text{AML}}^{(2)}$. Of course, this minimiser can also be recovered using FNS or the basic HEIV scheme, both based on $J_{\text{AML}}^{(2)}$. Note that the original derivation of HEIV did not utilise $J_{\text{AML}}^{(2)}$ —see [3] for a more detailed explanation.

On a final point, CFNS offers a simultaneous extension of the FNS and HEIV methods to the case of constrained minimisation. In particular, the $J_{\text{AML}}^{(1)}$ and $J_{\text{AML}}^{(2)}$ -based versions of CFNS are the constrained-minimisation counterparts of the core and original versions of HEIV, respectively. The first of these is the original version of CFNS as introduced in [3]. The second emerges as a new method still to be tested and compared with other techniques.

Summarising, this paper has presented a unifying approach to three recent estimation techniques: FNS, CFNS and HEIV. The proposed formulation allows for consistent analysis of various existing algorithms and advancement of new variants.

References

1. M.J. Brooks, W. Chojnacki, and L. Baumela, "Determining the egomotion of an uncalibrated camera from instantaneous optical flow," *Journal of the Optical Society of America A*, Vol. 14, No. 10, pp. 2670–2677, 1997.
2. W. Chojnacki, M.J. Brooks, A. van den Hengel, and D. Gawley, "On the fitting of surfaces to data with covariances," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 22, No. 11, pp. 1294–1303, 2000.
3. W. Chojnacki, M.J. Brooks, A. van den Hengel, and D. Gawley, "From FNS to HEIV: A link between two vision parameter estimation methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 26, No. 2, pp. 264–268, 2004.
4. W. Chojnacki, M.J. Brooks, A. van den Hengel, and D. Gawley, "A new constrained parameter estimator for computer vision applications," *Image and Vision Computing*, Vol. 22, No. 2, pp. 85–91, 2004.
5. O.D. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press: Cambridge, MA, 1993.
6. A. Fitzgibbon, M. Pilu, and R.B. Fisher, "Direct least square fitting of ellipses," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 21, No. 5, pp. 476–480, 1999.
7. K. Kanatani, *Statistical Optimization for Geometric Computation: Theory and Practice*, Elsevier: Amsterdam, 1996.
8. Y. Leedan and P. Meer, "Heteroscedastic regression in computer vision: Problems with bilinear constraint," *International Journal of Computer Vision*, Vol. 37, No. 2, pp. 127–150, 2000.
9. B. Matei, "Heteroscedastic errors-in-variables models in computer vision," PhD thesis, Department of Electrical and Computer Engineering, Rutgers University, New Brunswick, NJ, May 2001. Available at <http://www.caip.rutgers.edu/riul/research/theses.html>.
10. B. Matei and 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, 2000, IEEE Computer Society Press: Los Alamitos, CA, 2000, Vol. 2, pp. 18–25.
11. P.D. Sampson, "Fitting conic sections to 'very scattered' data: An iterative refinement of the Bookstein algorithm," *Computer Graphics and Image Processing*, Vol. 18, No. 1, pp. 97–108, 1982.
12. A. van den Hengel, W. Chojnacki, M.J. Brooks, and D. Gawley, "A new constrained parameter estimator: experiments in fundamental matrix computation," in *Proceedings of the 13th British Machine Vision Conference*, P.L. Rosin and D. Marshall (eds.), Cardiff, England, 2–5 September, 2002, BMVA Press, 2002. Vol. 2, pp. 468–476.



Wojciech Chojnacki is a professor of mathematics in the Department of Mathematics and Natural Sciences at Cardinal Stefan Wyszyński University in Warsaw. He is concurrently a senior research fellow in the School of Computer Science at the University of Adelaide working on a range of problems in computer vision. His research interests include differential equations, mathematical foundations of computer vision, functional analysis, and harmonic analysis. He is author of over 70 articles on pure mathematics and machine vision, and a member of the Polish Mathematical Society.



Michael J. Brooks holds the Chair in Artificial Intelligence within the University of Adelaide's School of Computer Science, which he heads. He is also leader of the Image Analysis Program within the Cooperative Research Centre for Sensor Signal and Information Processing, based in South Australia. His research interests include structure from motion, self-calibration, metrology, statistical vision-

parameter estimation, and video surveillance and analysis. He is author of over 100 articles on vision, actively involved in a variety of commercial applications, an Associate Editor of the International Journal of Computer Vision, and a Fellow of the Australian Computer Society.



Anton van den Hengel is a senior lecturer in the School of Computer Science within the University of Adelaide. He is also leader of the Video Surveillance and Analysis Project within the Cooperative Research Centre for Sensor Signal and Information Processing. His research interests include structure from motion, parameter estimation theory, and commercial applications of computer vision.



Darren Gawley graduated with first class honours from the School of Computer Science at the University of Adelaide. He holds a temporary lectureship at the same University, and is currently finalising his PhD in the field of computer vision.