

A Consistency Result for the Normalized Eight-Point Algorithm

Wojciech Chojnacki, Michael J. Brooks

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

{wojciech.chojnacki, michael.brooks}@adelaide.edu.au

Abstract

A recently proposed argument to explain the improved performance of the eight-point algorithm that results from using normalized data [IEEE Trans. Pattern Anal. Mach. Intell., 25(9):1172–1177, 2003] relies upon adoption of a certain model for statistical data distribution. Under this model, the cost function that underlies the algorithm operating on the normalized data is statistically more advantageous than the cost function that underpins the algorithm using unnormalized data. Here we extend this explanation by introducing a more refined, structured model for data distribution. Under the extended model, the normalized eight-point algorithm turns out to be approximately consistent in a statistical sense. The proposed extension provides a link between the existing statistical rationalization of the normalized eight-point algorithm and the approach of Mühlich and Mester for enhancing total least squares estimation methods via equilibration. Our contribution forms part of a wider effort to rationalize and interrelate foundational methods in vision parameter estimation.

1. Introduction

In a version adapted for computation of the *fundamental matrix*, the *eight-point algorithm* of Longuet-Higgins [12] produces significantly more accurate estimates if normalized rather than unnormalized image data are used as input for the scheme. Hartley [7], who made this fundamental discovery, attributed the improvement in performance resulting from data normalization to the better numerical conditioning of a pivotal matrix used in solving an eigenvalue problem. Recently, an alternative explanation was proposed based on a certain model for statistical data distribution [3]. Under this model, the summands of the cost function underlying the normalized version of the eight-point algorithm are more balanced in terms of spread compared with the summands of the cost function underlying the unnormalized version. The improved statistical design of the associated cost function is seen as being conducive to the normal-

ized algorithm's gains in accuracy. Here we reinforce this line of argumentation by introducing a more refined, structured model for the data distribution. We show that under the extended model the normalized eight-point algorithm is approximately consistent in a statistical sense. This result is then directly linked to the work of Mühlich and Mester [15], where a similar argument is used to validate an algorithm akin to Hartley's. The study reported below forms part of a wider effort to rationalize and interrelate foundational methods in vision parameter estimation (e.g., see [1, 2, 4, 9, 11]).

2. Algebraic Least Squares

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

$$\mathbf{m}'^T \mathbf{F} \mathbf{m} = 0, \quad (1)$$

where $\mathbf{F} = [f_{ij}]$ is a 3×3 fundamental matrix that incorporates information about the relative orientation and internal geometry of the cameras [5, 8]. In addition to (1), \mathbf{F} is subject to the *singularity constraint* (or, equivalently, the *rank-2 constraint*)

$$\det \mathbf{F} = 0. \quad (2)$$

Given a set $\{(\mathbf{m}_n, \mathbf{m}'_n)\}_{n=1}^N$ of *measured* correspondences, a simple “linear” estimate of \mathbf{F} providing an approximate solution of the system

$$\mathbf{m}_n'^T \mathbf{F} \mathbf{m}_n = 0 \quad (n = 1, \dots, N)$$

is obtained by minimizing the cost function

$$J_{\text{ALS}}(\mathbf{F}) = \|\mathbf{F}\|_{\mathbf{F}}^{-2} \sum_{n=1}^N (\mathbf{m}_n'^T \mathbf{F} \mathbf{m}_n)^2$$

with $\|\mathbf{F}\|_F = (\sum_{i,j} f_{ij}^2)^{1/2}$ the Frobenius norm of \mathbf{F} . Here, $\mathbf{m}_n^T \mathbf{F} \mathbf{m}_n$ is the signed *algebraic distance* between the individual correspondence $(\mathbf{m}_n, \mathbf{m}'_n)$ and the candidate matrix \mathbf{F} . The minimizer of J_{ALS} constitutes the *algebraic least squares (ALS) estimate* and is denoted $\hat{\mathbf{F}}_{\text{ALS}}$. $\hat{\mathbf{F}}_{\text{ALS}}$ is determined only up to scale, as the value of $J_{\text{ALS}}(\mathbf{F})$ does not change when \mathbf{F} is multiplied by a nonzero scalar. As a rule, $\hat{\mathbf{F}}_{\text{ALS}}$ does not satisfy (2). However, this estimate, like many other estimates not obeying (2), can further be upgraded to a genuine rank-2 fundamental matrix by applying one of a variety of available methods [8, 9]. In this paper, only the ALS method and its variant for normalized data will be of concern and the singularity constraint will be set aside.

For the purpose of computing $\hat{\mathbf{F}}_{\text{ALS}}$, it is convenient to use the joint image point $\mathbf{x} = [m_x, m_y, m'_x, m'_y]^T$ as a compact descriptor of the single correspondence $(\mathbf{m}, \mathbf{m}')$. For an $m \times n$ matrix $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n]$ with \mathbf{a}_i the i th column vector of length m , denote by $\text{vec}(\mathbf{A})$ the *vectorization* of \mathbf{A} , i.e., the column vector of length mn defined by $\text{vec}(\mathbf{A}) = [\mathbf{a}_1^T, \dots, \mathbf{a}_n^T]^T$. Let $\boldsymbol{\theta} = \text{vec}(\mathbf{F}^T)$ and $\mathbf{u}(\mathbf{x}) = \text{vec}(\mathbf{m}\mathbf{m}'^T)$. Then, as one easily verifies,

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

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

$$J_{\text{ALS}}(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|^{-2} \boldsymbol{\theta}^T \mathbf{A} \boldsymbol{\theta},$$

where

$$\mathbf{A} = \sum_{n=1}^N \mathbf{u}(\mathbf{x}_n) \mathbf{u}(\mathbf{x}_n)^T$$

and $\|\boldsymbol{\theta}\| = (\theta_1^2 + \dots + \theta_3^2)^{1/2}$. Now, an optimality condition for the minimizer of J_{ALS} (expressed as the zeroing of the gradient of the cost function) can be used to characterize $\hat{\boldsymbol{\theta}}_{\text{ALS}}$ as the normalized eigenvector of \mathbf{A} associated with the smallest eigenvalue (the smallest eigenvector). The eigenvector is defined uniquely up to a sign [2]. In practice $\hat{\boldsymbol{\theta}}_{\text{ALS}}$ is found by performing *singular value decomposition (SVD)* of the matrix

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

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

3. Normalized Algebraic Least Squares

ALS estimates are highly sensitive to image noise. To curb noise-driven instabilities, Hartley proposed a simple

normalization of the input data to be performed prior to running the ALS method [7]. The normalization process involves two affine transformations

$$\tilde{\mathbf{m}}_n = \mathbf{T} \mathbf{m}_n, \quad \tilde{\mathbf{m}}'_n = \mathbf{T}' \mathbf{m}'_n,$$

where \mathbf{T} and \mathbf{T}' are 3×3 data-dependent matrices whose exact form will be presented shortly. The ALS method is applied to the normalized data $\{(\tilde{\mathbf{m}}_n, \tilde{\mathbf{m}}'_n)\}_{n=1}^N$ and then the result is suitably back-transformed. More specifically, if, for each $n = 1, \dots, N$, $\tilde{\mathbf{x}}_n = [\tilde{m}_{x,n}, \tilde{m}_{y,n}, \tilde{m}'_{x,n}, \tilde{m}'_{y,n}]^T$ is the joint image point corresponding to $(\tilde{\mathbf{m}}_n, \tilde{\mathbf{m}}'_n)$ and $\tilde{\mathbf{F}}_{\text{ALS}}$ is the ALS estimate based on $\{\tilde{\mathbf{x}}_n\}_{n=1}^N$, then the *normalized algebraic least squares (NALS)* estimate of \mathbf{F} , $\hat{\mathbf{F}}_{\text{NALS}}$, is defined by

$$\hat{\mathbf{F}}_{\text{NALS}} = \mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS}} \mathbf{T}.$$

The matrices \mathbf{T} , \mathbf{T}' come in two variants depending on whether the coordinates $m_{x,n}, m_{y,n}$ and $m'_{x,n}, m'_{y,n}$ are normalized separately or collectively. Let

$$\begin{aligned} \bar{\mathbf{m}} &= [\bar{m}_x, \bar{m}_y, 1] = \frac{1}{N} \sum_{n=1}^N \mathbf{m}_n, \\ \bar{\mathbf{m}}' &= [\bar{m}'_x, \bar{m}'_y, 1] = \frac{1}{N} \sum_{n=1}^N \mathbf{m}'_n \end{aligned}$$

be the centroids of the data points in each image. For the data normalized separately, *anisotropic* scaling factors are introduced by means of the formulae

$$\begin{aligned} s_x &= \left(\frac{1}{N} \sum_{n=1}^N (m_{x,n} - \bar{m}_x)^2 \right)^{1/2}, \\ s_y &= \left(\frac{1}{N} \sum_{n=1}^N (m_{y,n} - \bar{m}_y)^2 \right)^{1/2}, \\ s'_x &= \left(\frac{1}{N} \sum_{n=1}^N (m'_{x,n} - \bar{m}'_x)^2 \right)^{1/2}, \\ s'_y &= \left(\frac{1}{N} \sum_{n=1}^N (m'_{y,n} - \bar{m}'_y)^2 \right)^{1/2}, \end{aligned}$$

and the matrices \mathbf{T} and \mathbf{T}' are defined by

$$\begin{aligned} \mathbf{T} &= \begin{bmatrix} s_x^{-1} & 0 & -s_x^{-1} \bar{m}_x \\ 0 & s_y^{-1} & -s_y^{-1} \bar{m}_y \\ 0 & 0 & 1 \end{bmatrix}, \\ \mathbf{T}' &= \begin{bmatrix} s'_x{}^{-1} & 0 & -s'_x{}^{-1} \bar{m}'_x \\ 0 & s'_y{}^{-1} & -s'_y{}^{-1} \bar{m}'_y \\ 0 & 0 & 1 \end{bmatrix}. \end{aligned}$$

For the data normalized individually, *isotropic* scaling factors take the form

$$s = \left(\frac{1}{2N} \sum_{n=1}^N (m_{x,n} - \bar{m}_x)^2 + (m_{y,n} - \bar{m}_y)^2 \right)^{1/2},$$

$$s' = \left(\frac{1}{2N} \sum_{n=1}^N (m'_{x,n} - \bar{m}'_x)^2 + (m'_{y,n} - \bar{m}'_y)^2 \right)^{1/2},$$

and \mathbf{T} and \mathbf{T}' are given by

$$\mathbf{T} = \begin{bmatrix} s^{-1} & 0 & -s^{-1}\bar{m}_x \\ 0 & s^{-1} & -s^{-1}\bar{m}_y \\ 0 & 0 & 1 \end{bmatrix},$$

$$\mathbf{T}' = \begin{bmatrix} s'^{-1} & 0 & -s'^{-1}\bar{m}'_x \\ 0 & s'^{-1} & -s'^{-1}\bar{m}'_y \\ 0 & 0 & 1 \end{bmatrix}.$$

The numerical motivation for the NALS method comes from the fact that the matrix $\tilde{\mathbf{A}} = \sum_{n=1}^N \mathbf{u}(\tilde{\mathbf{x}}_n)\mathbf{u}(\tilde{\mathbf{x}}_n)^T$ serving to calculate $\hat{\mathbf{F}}_{\text{ALS}}$ is better conditioned—in a certain sense—than the matrix \mathbf{A} . The improved conditioning leads to the smallest eigenvector of $\tilde{\mathbf{A}}$ being less sensitive to small perturbations of the matrix entries than the smallest eigenvector of \mathbf{A} [7]. The NALS method can also be justified statistically [3]. The statistical argument exploits the fact that $\hat{\mathbf{F}}_{\text{NALS}}$ is the minimizer of the *normalized algebraic least squares* (NALS) cost function

$$J_{\text{NALS}}(\mathbf{F}) = \|\mathbf{T}'^{-T} \mathbf{F} \mathbf{T}^{-1}\|_F^{-2} \sum_{n=1}^N (\mathbf{m}_n^T \mathbf{F} \mathbf{m}_n)^2.$$

A fundamental observation is that different *residuals*

$$r_n = \mathbf{m}_n^T \mathbf{F} \mathbf{m}_n \quad (n = 1, \dots, N)$$

carry, as a rule, different statistical weight. When the \mathbf{m}_n , \mathbf{m}'_n are treated as sample values of independent multivariate random variables, the r_n are realizations of (typically) *heteroscedastic* random variables—that is, ones having different variances. The larger the variance of a particular r_n , the less reliable this residual is likely to be, and the more it should be devalued. This suggests that, to account for heteroscedasticity, the simple cost function $\sum_{n=1}^N r_n^2$, which effectively is identical to J_{ALS} , should be replaced by the more complicated cost function $\sum_{n=1}^N r_n^2 / \text{var}[r_n]$, where $\text{var}[r]$ denotes the variance of r . As it turns out, one can conceive a model for data distribution such that all the $\text{var}[r_n]$ are equal to $\|\mathbf{T}'^{-T} \mathbf{F} \mathbf{T}^{-1}\|_F^2$ and consequently $\sum_{n=1}^N r_n^2 / \text{var}[r_n]$ reduces to J_{NALS} . Construed in this fashion, J_{NALS} has statistical advantage over J_{ALS} , leading to an improvement in performance of the normalized version of the eight-point algorithm over the unnormalized version.

The data distribution model that makes the above argument work is as follows. For each $n = 1, \dots, N$, one introduces random variables

$$\mathbf{m}_n = \bar{\mathbf{m}} + \Delta \mathbf{m}_n, \quad \mathbf{m}'_n = \bar{\mathbf{m}}' + \Delta \mathbf{m}'_n,$$

where $\bar{\mathbf{m}}$ and $\bar{\mathbf{m}}'$ are fixed, nonrandom points, and $\Delta \mathbf{m}_n = [\Delta m_{x,n}, \Delta m_{y,n}, 0]^T$ and $\Delta \mathbf{m}'_n = [\Delta m'_{x,n}, \Delta m'_{y,n}, 0]^T$ are random isotropic perturbations satisfying the following conditions:

- the collection $\{\Delta \mathbf{m}_n, \Delta \mathbf{m}'_n\}_{n=1}^N$ is jointly independent;
- $E[\Delta \mathbf{m}_n] = E[\Delta \mathbf{m}'_n] = \mathbf{0}$ for each $n = 1, \dots, N$;
- for some $s, s' > 0$ and each $n = 1, \dots, N$,

$$E[\Delta \mathbf{m}_n \Delta \mathbf{m}_n^T] = \text{diag}(s^2, s^2, 0),$$

$$E[\Delta \mathbf{m}'_n \Delta \mathbf{m}'_n^T] = \text{diag}(s'^2, s'^2, 0).$$

Here, E denotes the expectation value. The observed locations \mathbf{m}_n and \mathbf{m}'_n are viewed as realizations of \mathbf{m}_n and \mathbf{m}'_n , respectively.

4. Extended Statistical Model

A different approach to achieving improved ALS estimates was proposed by Mühlich and Mester [15, 16]. The idea is to search for 3×3 affine *weight matrices* \mathbf{W} and \mathbf{W}' with the property that if $\hat{\mathbf{F}}_{\text{ALS}}$ is the ALS estimate based on the data obtained by applying the transformations $\tilde{\mathbf{m}} = \mathbf{W} \mathbf{m}$ and $\tilde{\mathbf{m}}' = \mathbf{W}' \mathbf{m}'$ on the raw data, then the *equilibrated ALS estimate*, $\hat{\mathbf{F}}_{\text{EQL}}$, defined by

$$\hat{\mathbf{F}}_{\text{EQL}} = \mathbf{W}'^T \hat{\mathbf{F}}_{\text{ALS}} \mathbf{W}$$

is approximately unbiased in a statistical sense. The matrices \mathbf{W} and \mathbf{W}' are analogs of the normalization matrices \mathbf{T} and \mathbf{T}' used in the normalized eight-point algorithm. Various choices of weight matrices are possible. The resulting equilibrated ALS methods are akin to, but not identical with, Hartley's technique.

Standard weight matrices are obtained by exploiting a measurement error model whereby noise is present in *one* image only [15]. In the next section we shall establish an unbiasedness property for the normalized eight-point algorithm using a statistical model that incorporates noise in *both* images. This model is an extension of the model described in the previous section. We present the extended model and operational entities related to it in this section. The argument leading to the unbiasedness property draws upon Mühlich–Mester's approach. While tailored to elucidate the utility of Hartley normalization for the eight-point algorithm, it can also serve, with due modifications,

to buttress Hartley normalization for other ALS methods like those appropriate for conic fitting or homography estimation. A distinctive feature of our result is its asymptotic character—only if the number of data points increases will the bias in estimates decrease to zero.

We proceed to describe the extended model. A set of N image correspondences is modeled statistically using a set of pairs of random variables $\{(\mathbf{m}_n, \mathbf{m}'_n)\}_{n=1}^N$. The number of corresponding pairs is not held constant but is allowed to vary from 1 to infinity. Thus, effectively, the model involves the infinite set of pairs of random variables $\{(\mathbf{m}_n, \mathbf{m}'_n)\}_{n=1}^\infty$. The given data points $\mathbf{m}_n, \mathbf{m}'_n$ are viewed as samples drawn from $\mathbf{m}_n, \mathbf{m}'_n$, respectively. The $\mathbf{m}_n, \mathbf{m}'_n$ are represented as

$$\mathbf{m}_n = \mathbf{n}_n + \Delta\mathbf{m}_n, \quad \mathbf{m}'_n = \mathbf{n}'_n + \Delta\mathbf{m}'_n,$$

where $\mathbf{n}_n, \mathbf{n}'_n$ are random variables describing “true” locations and $\Delta\mathbf{m}_n = [\Delta m_{x,n}, \Delta m_{y,n}, 0]^T$, $\Delta\mathbf{m}'_n = [\Delta m'_{x,n}, \Delta m'_{y,n}, 0]^T$ are random variables representing noise that corrupts the “true” locations. All the $\mathbf{n}_n, \mathbf{n}'_n, \Delta\mathbf{m}_n, \Delta\mathbf{m}'_n$ are assumed mutually independent. The $\Delta\mathbf{m}_n, \Delta\mathbf{m}'_n$ encapsulate homogeneous anisotropic noise with principal components aligned with the x and y axes—they all have zero mean and satisfy

$$\begin{aligned} E[\Delta\mathbf{m}_n \Delta\mathbf{m}_n^T] &= \text{diag}(\sigma_x^2, \sigma_y^2, 0), \\ E[\Delta\mathbf{m}'_n \Delta\mathbf{m}'_n^T] &= \text{diag}(\sigma_x'^2, \sigma_y'^2, 0) \end{aligned}$$

for some $\sigma_x, \sigma_y, \sigma_x', \sigma_y' > 0$. The $\mathbf{n}_n, \mathbf{n}'_n$ capture various possible arrangements of ideal data points and are assumed to undulate anisotropically around some fixed, nonrandom “centroids” \mathbf{c}, \mathbf{c}' :

$$\mathbf{n}_n = \mathbf{c} + \Delta\mathbf{n}_n, \quad \mathbf{n}'_n = \mathbf{c}' + \Delta\mathbf{n}'_n.$$

Here $\Delta\mathbf{n}_n = [\Delta n_{x,n}, \Delta n_{y,n}, 0]^T$, $\Delta\mathbf{n}'_n = [\Delta n'_{x,n}, \Delta n'_{y,n}, 0]^T$ are zero-mean random variables satisfying

$$\begin{aligned} E[\Delta\mathbf{n}_n \Delta\mathbf{n}_n^T] &= \text{diag}(\tau_x^2, \tau_y^2, 0), \\ E[\Delta\mathbf{n}'_n \Delta\mathbf{n}'_n^T] &= \text{diag}(\tau_x'^2, \tau_y'^2, 0) \end{aligned}$$

for some fixed $\tau_x, \tau_y, \tau_x', \tau_y' > 0$. Note that under the above model, all the $\Delta\mathbf{m}_n, \Delta\mathbf{m}'_n, \Delta\mathbf{n}_n, \Delta\mathbf{n}'_n$ are mutually independent. Note also that the $\mathbf{n}_n, \mathbf{n}'_n$ being mutually independent—an assumption that is part of the model—is not compatible with the requirement that, for each N , there exists a genuine random fundamental matrix \mathbf{F}_N such that $\mathbf{n}_n^T \mathbf{F}_N \mathbf{n}_n = 0$ for all $n = 1, \dots, N$. However, the $\mathbf{n}_n, \mathbf{n}'_n$ being independent is consistent with the $\mathbf{m}_n, \mathbf{m}'_n$ being independent. As the latter property underpins the statistical model presented in the previous section, one that we aim to extend, we accept the independence of the $\mathbf{n}_n, \mathbf{n}'_n$ as the

leading feature. Let

$$\bar{\mathbf{m}}_N = \frac{1}{N} \sum_{n=1}^N \mathbf{m}_n, \quad \bar{\mathbf{m}}'_N = \frac{1}{N} \sum_{n=1}^N \mathbf{m}'_n$$

and

$$\begin{aligned} s_{x,N} &= \left(\frac{1}{N} \sum_{n=1}^N (m_{x,n} - \bar{m}_{x,N})^2 \right)^{1/2}, \\ s_{y,N} &= \left(\frac{1}{N} \sum_{n=1}^N (m_{y,n} - \bar{m}_{y,N})^2 \right)^{1/2}, \\ s'_{x,N} &= \left(\frac{1}{N} \sum_{n=1}^N (m'_{x,n} - \bar{m}'_{x,N})^2 \right)^{1/2}, \\ s'_{y,N} &= \left(\frac{1}{N} \sum_{n=1}^N (m'_{y,n} - \bar{m}'_{y,N})^2 \right)^{1/2}. \end{aligned}$$

Let

$$\begin{aligned} \mathbf{T}_N &= \begin{bmatrix} s_{x,N}^{-1} & 0 & -s_{x,N}^{-1} \bar{m}_{x,N} \\ 0 & s_{y,N}^{-1} & -s_{y,N}^{-1} \bar{m}_{y,N} \\ 0 & 0 & 1 \end{bmatrix}, \\ \mathbf{T}'_N &= \begin{bmatrix} s'_{x,N}{}^{-1} & 0 & -s'_{x,N}{}^{-1} \bar{m}'_{x,N} \\ 0 & s'_{y,N}{}^{-1} & -s'_{y,N}{}^{-1} \bar{m}'_{y,N} \\ 0 & 0 & 1 \end{bmatrix}. \end{aligned}$$

Here and in the following, the additional subscript N serves to emphasize the dependence on N of various quantities involved. Using the strong law of large numbers one can show that, with N increasing to infinity, the sequences $\{\mathbf{T}_N\}_{N=1}^\infty$ and $\{\mathbf{T}'_N\}_{N=1}^\infty$ converge with probability 1 to certain nonrandom matrices \mathbf{T} and \mathbf{T}' explicitly expressible in terms of $\sigma_x, \sigma_y, \tau_x, \tau_y$ and $\sigma_x', \sigma_y', \tau_x', \tau_y'$. Now, for each $n = 1, \dots, N$, $\tilde{\mathbf{m}}_{n,N} = \mathbf{T}_N \mathbf{m}_n$ and $\tilde{\mathbf{m}}'_{n,N} = \mathbf{T}'_N \mathbf{m}'_n$ are the normalized noisy image points derived from \mathbf{m}_n and \mathbf{m}'_n , with the normalization based on $\{(\mathbf{m}_n, \mathbf{m}'_n)\}_{n=1}^N$. We define $\tilde{\mathbf{n}}_n = \mathbf{T} \mathbf{n}_n$, $\tilde{\mathbf{n}}'_n = \mathbf{T}' \mathbf{n}'_n$ for every n ranging from 1 to infinity. Let $\mathbf{x}_n = [m_{n,x}, m_{n,y}, m'_{n,x}, m'_{n,y}]^T$ and $\mathbf{y}_n = [n_{n,x}, n_{n,y}, n'_{n,x}, n'_{n,y}]^T$ be the joint image points derived from $(\mathbf{m}_n, \mathbf{m}'_n)$ and $(\mathbf{n}_n, \mathbf{n}'_n)$, respectively. Let $\tilde{\mathbf{x}}_{n,N} = [\tilde{m}_{n,N,x}, \tilde{m}_{n,N,y}, \tilde{m}'_{n,N,x}, \tilde{m}'_{n,N,y}]^T$ and $\tilde{\mathbf{y}}_{n,N} = [\tilde{n}_{n,N,x}, \tilde{n}_{n,N,y}, \tilde{n}'_{n,N,x}, \tilde{n}'_{n,N,y}]^T$ be the corresponding transformed joint image points. Let

$$\begin{aligned} \tilde{\mathbf{A}}_N &= \frac{1}{N} \sum_{n=1}^N \mathbf{u}(\tilde{\mathbf{x}}_{n,N}) \mathbf{u}(\tilde{\mathbf{x}}_{n,N})^T, \\ \tilde{\mathbf{B}}_N &= \frac{1}{N} \sum_{n=1}^N \mathbf{u}(\tilde{\mathbf{y}}_{n,N}) \mathbf{u}(\tilde{\mathbf{y}}_{n,N})^T. \end{aligned}$$

Here, the common normalizing factor $1/N$ is introduced to maintain an overall statistical balance—a quality of importance in performing relevant calculations. Crucially, this factor does not affect the eigenvectors of the matrices involved. For each $j = 1, \dots, 9$, let $\mathbf{a}_{j,N}$ be the j th normalized eigenvector of $\hat{\mathbf{A}}_N$ (defined uniquely up to a sign) with a corresponding eigenvalue $\alpha_{j,N}$, and let $\mathbf{b}_{j,N}$ be the j th normalized eigenvector of $\hat{\mathbf{B}}_N$ with a corresponding eigenvalue $\beta_{j,N}$. The eigenvalues are assumed to be arranged in descending order. The random eigenvectors $\mathbf{a}_{9,N}$ and $\mathbf{b}_{9,N}$ can be restructured to form random matrices $\hat{\mathbf{F}}_{\text{ALS},N}$ and $\hat{\mathbf{E}}_N$ such that

$$\mathbf{a}_{9,N} = \text{vec}(\hat{\mathbf{F}}_{\text{ALS},N}^T), \quad \mathbf{b}_{9,N} = \text{vec}(\hat{\mathbf{E}}_N^T).$$

Then, clearly,

$$\hat{\mathbf{F}}_{\text{NALS},N} = \mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T}_N$$

is the NALS estimate based on $\{\mathbf{x}_n\}_{n=1}^N$. The matrix

$$\mathbf{F}_N = \mathbf{T}'^T \tilde{\mathbf{E}}_N \mathbf{T}$$

will in what follows constitute a reference value against which to compare $\hat{\mathbf{F}}_{\text{NALS},N}$. Obviously, if the $\{(\mathbf{n}_n, \mathbf{n}'_n)\}_{n=1}^N$ were genuine correspondences bound by a fundamental matrix \mathbf{F}_N , then \mathbf{F}_N would be a natural “true” value with which to gauge $\hat{\mathbf{F}}_{\text{NALS},N}$. Since, generally, no such \mathbf{F}_N exists, we use \mathbf{F}_N instead. The significance of this choice can only be fully appreciated once the whole argument is complete.

With $\mathcal{A}_N = \{\mathbf{n}_n, \mathbf{n}'_n\}_{n=1}^N$, let $E[\hat{\mathbf{F}}_{\text{NALS},N} | \mathcal{A}_N]$ be the conditional expectation of $\hat{\mathbf{F}}_{\text{NALS},N}$ given \mathcal{A}_N . We adopt

$$\Delta_N = \|E[\hat{\mathbf{F}}_{\text{NALS},N} | \mathcal{A}_N] - \mathbf{F}_N\|$$

as a measure of *bias* in the $\hat{\mathbf{F}}_{\text{NALS},N}$ estimates. The random variable Δ_N captures the bias conditional on any particular arrangement of the “true” data points $\{\mathbf{n}_n, \mathbf{n}'_n\}_{n=1}^N$. Assuming that all the $\{\mathbf{n}_n, \mathbf{n}'_n, \mathbf{m}_n, \mathbf{m}'_n\}_{n=1}^\infty$ are uniformly bounded random variables carried on a single probability space, we shall outline an argument to show that a certain *approximation* to Δ_N tends to 0 as $N \rightarrow \infty$ with probability 1. This asymptotic approximate unbiasedness will be referred to as *approximate consistency*. It can be viewed as one of advantageous effects of data normalization. The establishment of the property parallels similar findings related to other methods [10, 14, 18].

5. Approximate Consistency

Let $\mathcal{A}_\infty = \bigcup_{N=1}^\infty \mathcal{A}_N = \{\mathbf{n}_n, \mathbf{n}'_n\}_{n=1}^\infty$. Since both sequences $\{\mathbf{T}_N\}_{N=1}^\infty$, $\{\mathbf{T}'_N\}_{N=1}^\infty$ are uniformly bounded

(given that the family $\{\mathbf{m}_n, \mathbf{m}'_n\}_{n=1}^\infty$ is uniformly bounded) and pointwise converge to \mathbf{T} , \mathbf{T}' almost surely¹ (a.s.), an application of the conditional form of the dominated convergence theorem [17, Sec. 2, §7, Thm. 2] implies that

$$\lim_{N \rightarrow \infty} E[\mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T}_N - \mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T} | \mathcal{A}_\infty] = \mathbf{0}.$$

On the other hand,

$$\begin{aligned} E[\mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T}_N | \mathcal{A}_\infty] &= E[\mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T}_N | \mathcal{A}_N] \\ E[\mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T} | \mathcal{A}_\infty] &= \mathbf{T}'^T E[\hat{\mathbf{F}}_{\text{ALS},N} | \mathcal{A}_N] \mathbf{T} \end{aligned}$$

because \mathbf{T}_N , \mathbf{T}'_N , and $\hat{\mathbf{F}}_{\text{ALS},N}$ being expressible in terms of $\{\mathbf{m}_n, \mathbf{m}'_n, \mathbf{n}_n, \mathbf{n}'_n\}_{n=1}^N$ are jointly independent of $\{\mathbf{n}_n, \mathbf{n}'_n\}_{n=N+1}^\infty$. Therefore the difference

$$\begin{aligned} E[\hat{\mathbf{F}}_{\text{NALS},N} | \mathcal{A}_N] - \mathbf{F}_N \\ = E[\mathbf{T}'^T \hat{\mathbf{F}}_{\text{ALS},N} \mathbf{T}_N | \mathcal{A}_N] - \mathbf{T}'^T \tilde{\mathbf{E}}_N \mathbf{T} \end{aligned}$$

differs from

$$\mathbf{T}'^T E[\hat{\mathbf{F}}_{\text{ALS},N} | \mathcal{A}_N] \mathbf{T} - \mathbf{T}'^T \tilde{\mathbf{E}}_N \mathbf{T}$$

by an expression tending to zero almost surely. We see that the asymptotic behavior of Δ_N is fully controlled by the bias δ_N defined by

$$\delta_N = \|E[\hat{\mathbf{F}}_{\text{ALS},N} | \mathcal{A}_N] - \tilde{\mathbf{E}}_N\|$$

or, equivalently, by

$$\delta_N = \|E[\mathbf{a}_{9,N} | \mathcal{A}_N] - \mathbf{b}_{9,N}\|.$$

The value of δ_N is difficult to estimate because of the lack of a simple expression for $\mathbf{a}_{9,N}$. We shall get around this problem by considering a first-order approximation to $\mathbf{a}_{9,N}$. Specifically, we shall use the following approximate formula for $\mathbf{a}_{j,N}$

$$\mathbf{a}_{j,N} = \mathbf{b}_{j,N} + \sum_{\substack{k=1 \\ k \neq j}}^9 \frac{\mathbf{b}_{k,N}^T (\tilde{\mathbf{A}}_N - \tilde{\mathbf{B}}_N) \mathbf{b}_{j,N}}{\beta_{j,N} - \beta_{k,N}} \mathbf{b}_{k,N}$$

holding for each $1 \leq j \leq 9$ [6, Sec. 7.2]. We shift the focus from the exact bias δ_N to an approximate bias δ_N^* defined by

$$\delta_N^* = \left\| \sum_{k=1}^8 \frac{\mathbf{b}_{k,N}^T (E[\tilde{\mathbf{A}}_N | \mathcal{A}_N] - \tilde{\mathbf{B}}_N) \mathbf{b}_{9,N}}{\beta_{9,N} - \beta_{k,N}} \mathbf{b}_{k,N} \right\|.$$

¹Recall that the phrase “almost surely” is a concise, precise way to state “except on an event of probability zero”. It is equivalent to “with probability 1”.

Now, it can be shown that

$$\gamma = \lim_{N \rightarrow \infty} (\beta_{8,N} - \beta_{9,N})^{-1} < \infty \quad \text{a.s.}$$

for a generic choice of $\sigma_x, \sigma_y, \tau_x, \tau_y$ and $\sigma'_x, \sigma'_y, \tau'_x, \tau'_y$. This condition automatically implies that

$$\limsup_{N \rightarrow \infty} (\beta_{j,N} - \beta_{9,N})^{-1} \leq \gamma < \infty \quad \text{a.s.}$$

for all $1 \leq j \leq 8$. It can be also shown—and this is the main computational step, details of which we omit due to space limitations—that

$$\lim_{n \rightarrow \infty} \mathbb{E}[\tilde{\mathbf{A}}_N | \mathcal{A}_N] = \mathbf{I}_9 \quad \text{a.s.},$$

where, for each $n = 1, 2, \dots$, \mathbf{I}_n denotes the $n \times n$ identity matrix. Hence, with probability 1,

$$\limsup_{N \rightarrow \infty} \delta_N^* \leq \gamma \limsup_{N \rightarrow \infty} \sum_{k=1}^8 \|\mathbf{b}_{k,N}^T (\mathbf{I}_9 - \tilde{\mathbf{B}}_N) \mathbf{b}_{9,N}\| \|\mathbf{b}_{k,N}\|.$$

Since $\|\mathbf{b}_{k,N}\| = 1$ and

$$\mathbf{b}_{k,N}^T (\mathbf{I}_9 - \tilde{\mathbf{B}}_N) \mathbf{b}_{9,N} = (1 - \beta_{9,N}) \mathbf{b}_{k,N}^T \mathbf{b}_{9,N} = 0$$

for $1 \leq k \leq 8$, it follows that $\lim_{N \rightarrow \infty} \delta_N^* = 0$ with probability 1. We thus see that, when quantified with $\lim_{N \rightarrow \infty} \delta_N^*$, the asymptotic approximate bias in the normalized eight-point algorithm is zero almost surely. This is our sought-after approximate consistency result.

6. Conclusion

A new argument has been proposed to explain the improvement in performance of the eight-point algorithm that results from using normalized data. The argument relies upon adoption of a structured model for statistical data distribution. In the framework of this model, the advantage of using normalized data is attributed to a systematic increase in the consistency of estimates as the number of data points mounts. The proposed approach merges the previous statistical argument explaining the advantageous character of data normalization with the approach of Mühlich and Mester for enhancing total least squares estimation methods via equilibration. As such it contributes to a wider effort to gain a deeper understanding of the rationale underlying foundational methods in vision parameter estimation.

Acknowledgments

The support of the Australian Research Council is gratefully acknowledged.

References

- [1] W. Chojnacki, M. J. Brooks, and A. van den Hengel. Rationalising the renormalisation method of Kanatani. *J. Math. Imaging Vision*, 14(1):21–38, 2001.
- [2] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. On the fitting of surfaces to data with covariances. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(11):1294–1303, 2000.
- [3] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. Revisiting Hartley’s normalized eight-point algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(9):1172–1177, 2003.
- [4] W. Chojnacki, M. J. Brooks, A. van den Hengel, and D. Gawley. FNS, CFNS and HEIV: A unifying approach. *J. Math. Imaging Vision*, 23(2):175–183, 2005.
- [5] O. Faugeras, Q.-T. Luong, and T. Papadopoulos. *The Geometry of Multiple Images*. MIT Press, Cambridge, Mass., 2001.
- [6] G. H. Golub and C. F. van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, Maryland, 2nd edition, 1989.
- [7] R. Hartley. In defense of the eight-point algorithm. *IEEE Trans. Pattern Anal. Mach. Intell.*, 19(6):580–593, 1997.
- [8] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, 2000.
- [9] K. Kanatani. *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, Amsterdam, 1996.
- [10] A. Kukush, I. Markovsky, and S. van Huffel. Consistent fundamental matrix estimation in a quadratic measurement error model arising in motion analysis. *Comput. Statist. Data Anal.*, 41(1):3–18, 2002.
- [11] Y. Leedan and P. Meer. Heteroscedastic regression in computer vision: problems with bilinear constraint. *Int. J. Comput. Vision*, 37(2):127–150, 2000.
- [12] H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293(10):133–135, 1981.
- [13] H. Lütkepohl. *Handbook of Matrices*. John Wiley & Sons, Chichester, 1996.
- [14] W. J. MacLean. Removal of translation bias when using subspace methods. In *Proc. Seventh Int. Conf. Computer Vision, Corfu, Greece*, volume 2, pages 753–758, 1999.
- [15] M. Mühlich and R. Mester. The role of total least squares in motion analysis. In *Proc. Fifth European Conf. Computer Vision, Freiburg, Germany*, volume 1407 of *Lecture Notes in Computer Science*, pages 305–321, 1998.
- [16] M. Mühlich and R. Mester. Subspace methods and equilibration in computer vision. In *Proc. 12th Scandinavian Conf. Image Analysis, Stavanger, Norway*, pages 415–422, 2001.
- [17] A. N. Shiryaev. *Probability*. Springer-Verlag, New York, 2 edition, 1996.
- [18] T. Zhang and C. Tomasi. On the consistency of instantaneous rigid motion estimation. *Int. J. Comput. Vision*, 46(1):51–79, 2002.