# A Framework for Understanding Renormalisation-Type Methods in Computer Vision

Wojciech Chojnacki Michael J. Brooks Anton van den Hengel
Department of Computer Science
University of Adelaide
Adelaide, SA 5005, Australia

**Abstract**

*The renormalisation scheme of Kanatani is intended to iteratively minimise a cost function of certain form while avoiding systematic bias inherent in Sampson's method of minimisation. This paper is concerned with enhancing our understanding of Kanatani's complex scheme by expressing it within a novel framework common to several methods. This approach enables us to demonstrate that the renormalisation scheme does not have as its theoretical limit the desired minimiser, in contrast with an alternative, simpler approach presented.*

## 1 Introduction

A wide class of computer vision problems may be couched in terms of an equation of 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 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 such a way that: (i) each component is a quadratic form in the compound vector $[\boldsymbol{x}^T, 1]^T$, (ii) one component of $\boldsymbol{u}(\boldsymbol{x})$ is equal to 1. An *ancillary constraint* may also apply that does not involve the data, and this can be expressed as $\psi(\boldsymbol{\theta}) = 0$ for some scalar-valued function $\psi$. We consider the following estimation problem: Given a collection $(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ of image data, determine $\boldsymbol{\theta} \neq \mathbf{0}$ satisfying the ancillary equation (if applicable) together with the system of equations obtained from (1) by substituting $\boldsymbol{x}_i$ into $\boldsymbol{x}$ for each $1 \leq i \leq n$. When $n > l$ and noise is present, the corresponding system of equations is overdetermined and as such may fail to have a non-zero solution. In this situation, we are concerned with finding $\boldsymbol{\theta}$ that best fits the data in some sense.

*Conic fitting* is one problem of this kind (Bookstein, 1979; Zhang, 1997). Two other conformant problems are estimating coefficients of the *epipolar equation* (Faugeras, 1993), and estimating coefficients of the *differential epipolar equation* (Brooks et al., 1997; Viéville and

Faugeras, 1995), each involving an ancillary *cubic constraint*. The precise way in which these example problems accord with the our problem form is described in a companion work (Chojnacki et al., 1999).

# 2 Cost Functions and Estimators

A vast class of techniques for solving our problem rest upon the use of *cost functions* measuring the extent to which the data and candidate estimates fail to satisfy (1). If—for simplicity—one sets aside the ancillary constraint, then, given a cost function $J = J(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$, a corresponding estimate $\widehat{\boldsymbol{\theta}}$ is defined by $J(\widehat{\boldsymbol{\theta}}) = \min_{\boldsymbol{\theta} \neq 0} J(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$. Since (1) does not change if $\boldsymbol{\theta}$ is multiplied by a non-zero scalar, we consider only cost functions satisfying $J(t\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = J(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ for every non-zero scalar $t$. The assignment of $\widehat{\boldsymbol{\theta}}$, uniquely defined up to a scalar factor, to $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n$ will be termed the $J$-based *estimator* of $\boldsymbol{\theta}$.

## 2.1 Ordinary Least Squares Estimator

A straightforward estimator is derived from the cost function

$$J_{\mathrm{OLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \|\boldsymbol{\theta}\|^{-2} \sum_{i=1}^{n} \boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta},$$

where $\boldsymbol{A}_i = \boldsymbol{u}(\boldsymbol{x}_i)\boldsymbol{u}(\boldsymbol{x}_i)^T$ and $\|\boldsymbol{\theta}\| = (\theta_1^2 + \cdots + \theta_l^2)^{1/2}$. Here each summand $\boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}$ is the square of the *algebraic distance* $\boldsymbol{\theta}^T \boldsymbol{u}(\boldsymbol{x}_i)$. Accordingly, the $J_{\mathrm{OLS}}$-based estimate of $\boldsymbol{\theta}$ is termed the *ordinary least squares* (OLS) *estimate* and is denoted $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$. It is uniquely determined, up to a scalar factor, by an eigenvector of $\sum_{i=1}^{n} \boldsymbol{A}_i$ associated with the smallest eigenvalue.

## 2.2 Weighted Least Squares Estimator

The OLS estimator treats all data as being equally valuable. When information characterising the measurement errors is available, it is desirable that better data be weighted more heavily than poorer data during the estimation process. We assume that the data come equipped with a collection $(\boldsymbol{\Sigma}_{\boldsymbol{x}_1}, \ldots, \boldsymbol{\Sigma}_{\boldsymbol{x}_n})$ of positive definite $k \times k$ *covariance matrices*. These matrices constitute repositories of prior information about the uncertainty of the data.

It emerges that a strong case may be mounted for adoption of the cost function given by

$$J_{\mathrm{WLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \sum_{i=1}^{n} \frac{\boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}}{\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta}},$$

where $\boldsymbol{B}_i = \partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{x}_i) \boldsymbol{\Sigma}_{\boldsymbol{x}_i} \partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{x}_i)^T$ and $\partial_{\boldsymbol{x}} \boldsymbol{u}(\boldsymbol{y}) = [(\partial u_i / \partial x_j)(\boldsymbol{y})]_{1 \leq i \leq l, 1 \leq j \leq k}$ (Chojnacki et al., 1999; Kanatani, 1996). The $J_{\mathrm{WLS}}$-based estimate of $\boldsymbol{\theta}$ will be called the *weighted least squares* (WLS) *estimate* and will be denoted $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$.

### 2.3 Variational Equation

As a minimiser of $J_{\mathrm{WLS}}$, $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ satisfies $\partial_{\boldsymbol{\theta}} J_{\mathrm{WLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \mathbf{0}^T$, where $\partial_{\boldsymbol{\theta}} J_{\mathrm{WLS}}$ denotes the row vector of the partial derivatives of $J_{\mathrm{WLS}}$ with respect to $\boldsymbol{\theta}$. We term this the *variational equation*. Direct computation shows that $[\partial_{\boldsymbol{\theta}} J_{\mathrm{WLS}}(\boldsymbol{\theta}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)]^T = 2\boldsymbol{X}_{\boldsymbol{\theta}}\boldsymbol{\theta}$, where $\boldsymbol{X}_{\boldsymbol{\theta}}$ is the $l \times l$ symmetric matrix

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

Thus the variational equation can be rephrased as

$$\boldsymbol{X}_{\boldsymbol{\theta}}\boldsymbol{\theta} = \mathbf{0}. \tag{2}$$

This is a non-linear equation and is unlikely to admit solutions in closed form.

Obviously, not every solution of the variational equation is a point at which the global minimum of $J_{\mathrm{WLS}}$ is attained. However, the solution set of the equation provides a severely restricted family of candidates for the global minimiser. Within this set, the minimiser is much easier to identify.

## 3 Numerical Schemes

### 3.1 Fundamental Numerical Scheme

Closed-form solutions of the variational equation may be infeasible, so in practice $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ has to be found numerically. Throughout we shall assume that $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ lies close to $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$. This assumption is to increase the chances that any candidate minimiser obtained via a numerical method seeded with $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$ coincides with $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$.

A vector $\boldsymbol{\theta}$ satisfies (2) if and only if it falls into the null space of the matrix $\boldsymbol{X}_{\boldsymbol{\theta}}$. Thus, if $\boldsymbol{\theta}_{k-1}$ is a tentative guess, then an improved guess 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. The *fundamental numerical scheme* implementing this idea is presented in Figure 1. The algorithm can be regarded as a variant of the Newton-Raphson method.

### 3.2 Iteratively Reweighted Least Squares Scheme

Let $\boldsymbol{M}_{\boldsymbol{\theta}} = \sum_{i=1}^{n} (\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta})^{-1} \boldsymbol{A}_i$ and let $J'_{\mathrm{WLS}}(\boldsymbol{\theta}, \boldsymbol{\xi}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = \boldsymbol{\theta}^T \boldsymbol{M}_{\boldsymbol{\xi}} \boldsymbol{\theta}$. For simplicity, we abbreviate $J'_{\mathrm{WLS}}(\boldsymbol{\theta}, \boldsymbol{\xi}; \boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ to $J'_{\mathrm{WLS}}(\boldsymbol{\theta}, \boldsymbol{\xi})$. Sampson (1982) was the first to propose a scheme aiming to minimise a function involving fractional expressions, such as $J_{\mathrm{WLS}}$. Modelled on Sampson's scheme, the *iteratively reweighted least squares* (IRWLS) *scheme* applicable to $J_{\mathrm{WLS}}$ takes $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$ for an initial guess $\boldsymbol{\theta}_0$, and given $\boldsymbol{\theta}_{k-1}$ generates an update $\boldsymbol{\theta}_k$ by minimising the cost function $\boldsymbol{\theta} \mapsto J'_{\mathrm{WLS}}(\boldsymbol{\theta}, \boldsymbol{\theta}_{k-1})$. Assuming that the sequence $\{\boldsymbol{\theta}_k\}$ converges, the *iteratively reweighted least squares estimate* is defined as $\widehat{\boldsymbol{\theta}}_{\mathrm{IRWLS}} = \lim_{k\to\infty} \boldsymbol{\theta}_k$. Note that each

1. Set $\boldsymbol{\theta}_0 = \widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$.
2. Assuming that $\boldsymbol{\theta}_{k-1}$ is known, compute the matrix $\boldsymbol{X}_{\boldsymbol{\theta}_{k-1}}$.
3. Compute a normalised eigenvector of $\boldsymbol{X}_{\boldsymbol{\theta}_{k-1}}$ corresponding to the eigenvalue closest to zero 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.

Figure 1: Fundamental numerical scheme.

function $J'_{\mathrm{WLS}}(\boldsymbol{\theta}, \boldsymbol{\theta}_{k-1})$ is quadratic in $\boldsymbol{\theta}$. Finding a minimiser of such a function is straightforward. The minimiser $\boldsymbol{\theta}_k$ is an eigenvector of $\boldsymbol{M}_{\boldsymbol{\theta}_{k-1}}$ corresponding to the smallest eigenvalue; moreover, this eigenvalue is equal to $J'_{\mathrm{WLS}}(\boldsymbol{\theta}_k, \boldsymbol{\theta}_{k-1})$, so

$$\boldsymbol{M}_{\boldsymbol{\theta}_{k-1}} \boldsymbol{\theta}_k = J'_{\mathrm{WLS}}(\boldsymbol{\theta}_k, \boldsymbol{\theta}_{k-1}) \boldsymbol{\theta}_k. \tag{3}$$

The IRWLS scheme takes the same form as the fundamental numerical scheme given in Figure 1, except that $\boldsymbol{X}_{\boldsymbol{\theta}}$ is replaced by $\boldsymbol{M}_{\boldsymbol{\theta}}$. These two matrices are related by the formula $\boldsymbol{X}_{\boldsymbol{\theta}} = \boldsymbol{M}_{\boldsymbol{\theta}} - \boldsymbol{E}_{\boldsymbol{\theta}}$, where $\boldsymbol{E}_{\boldsymbol{\theta}} = \sum_{i=1}^{n} [(\boldsymbol{\theta}^T \boldsymbol{A}_i \boldsymbol{\theta}) / (\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta})^2] \boldsymbol{B}_i$. Letting $k \to \infty$ in (3) and taking into account the equality $J'_{\mathrm{WLS}}(\boldsymbol{\theta}, \boldsymbol{\theta}) = J_{\mathrm{WLS}}(\boldsymbol{\theta})$, we see that $\widehat{\boldsymbol{\theta}}_{\mathrm{IRWLS}}$ satisfies $[\boldsymbol{M}_{\boldsymbol{\theta}} - J_{\mathrm{WLS}}(\boldsymbol{\theta}) \boldsymbol{I}] \boldsymbol{\theta} = \boldsymbol{0}$, where $\boldsymbol{I}$ is the $l \times l$ identity matrix. We call this the *Sampson equation*. Note that it is different from the variational equation (2) and that, as a result, $\widehat{\boldsymbol{\theta}}_{\mathrm{IRWLS}}$ is not a genuine minimiser of $J_{\mathrm{WLS}}$.

## 3.3 Renormalisation Scheme

The matrices $\boldsymbol{M}_{\boldsymbol{\theta}} - \boldsymbol{E}_{\boldsymbol{\theta}}$ and $\boldsymbol{M}_{\boldsymbol{\theta}} - J_{\mathrm{WLS}}(\boldsymbol{\theta}) \boldsymbol{I}$ underlying the variational and Sampson equations can be viewed as modified or normalised forms of $\boldsymbol{M}_{\boldsymbol{\theta}}$. As first realised by Kanatani (1996), a different modification can be proposed based on statistical considerations. The requirement is that the modified or *renormalised* $\boldsymbol{M}_{\boldsymbol{\theta}}$ be *unbiased* in some sense. Using the renormalised $\boldsymbol{M}_{\boldsymbol{\theta}}$, one can formulate an equation analogous to both the variational and Sampson equations. This *renormalisation equation* can in turn be used to define an estimate of $\boldsymbol{\theta}$.

It emerges that the renormalisation equation takes the form $\boldsymbol{Y}_{\boldsymbol{\theta}} \boldsymbol{\theta} = \boldsymbol{0}$, where

$$\boldsymbol{Y}_{\boldsymbol{\theta}} = \boldsymbol{M}_{\boldsymbol{\theta}} - \frac{\boldsymbol{\theta}^T \boldsymbol{M}_{\boldsymbol{\theta}} \boldsymbol{\theta}}{\boldsymbol{\theta}^T \boldsymbol{N}_{\boldsymbol{\theta}} \boldsymbol{\theta}} \boldsymbol{N}_{\boldsymbol{\theta}}$$

and $\boldsymbol{N}_{\boldsymbol{\theta}} = \sum_{i=1}^{n} (\boldsymbol{\theta}^T \boldsymbol{B}_i \boldsymbol{\theta})^{-1} \boldsymbol{B}_i$ (Chojnacki et al., 1999). The equation is not naturally derived from any specific cost function, and, as a result, it is not clear whether it has any solution at all. A general belief is that in the close vicinity of $\widehat{\boldsymbol{\theta}}_{\mathrm{OLS}}$ there is a solution and only one. This

solution is termed the *renormalisation estimate* and is denoted $\widehat{\boldsymbol{\theta}}_{\mathrm{REN}}$. Since the renormalisation equation is different from the variational and Sampson equations, $\widehat{\boldsymbol{\theta}}_{\mathrm{REN}}$ is distinct both from $\widehat{\boldsymbol{\theta}}_{\mathrm{WLS}}$ and $\widehat{\boldsymbol{\theta}}_{\mathrm{IRWLS}}$.

A scheme for solving the renormalisation equation is readily generated. It follows precisely the pattern of the fundamental numerical scheme given in Figure 1, except that $\boldsymbol{X}_{\boldsymbol{\theta}}$ is replaced by $\boldsymbol{Y}_{\boldsymbol{\theta}}$. Many variations on this theme may be generated (Chojnacki et al., 1999).

# 4 Experimental Results

The previously derived algorithms were tested on the problem of conic fitting, which constitutes a classical benchmark problem in the literature (Bookstein, 1979; Ellis, 1992; Fitzgibbon et al., 1999; Fitzgibbon and Fisher, 1995; Gander et al., 1994; Kanatani, 1994; Porill, 1990; Rosin, 1993; Rosin and West, 1995; Sampson, 1982; Taubin, 1991; Zhang, 1997). Specifically, the fitting algorithms were applied to contaminated data arising from a portion of an ellipse.

Tests reported here are synthetic as they enable precise control of the nature of the data and their associated uncertainties. The tests proceeded as follows. A randomly oriented ellipse was generated such that the ratio of its major to minor axes was in the range $[2, 3]$, and its major axis was approximately 200 pixels in length. One third of the ellipse's boundary was chosen as the base curve, and this included the point of maximum curvature of the ellipse. A set of *true points* was then randomly selected from a distribution uniform along the length of the base curve.

For each of the true points, a covariance matrix was randomly generated in accordance with some chosen *average level of noise* (Chojnacki et al., 1999). The true points were then perturbed randomly in accordance with their associated covariance matrices, yielding the *data points*. In general, the noise conformed to an inhomogeneous and anisotropic distribution.

The various methods were then supplied with the data points. Those methods able to utilise uncertainty information were also supplied with the data points' covariance matrices. Each method was then challenged to determine the coefficients of the best fitting conic and a measure of the error of this estimate was then computed. (Note, therefore, that it was not assumed that the conic was an ellipse.) Testing was repeated many times using newly generated data points (with the covariance matrices and true data points remaining intact). The average errors were then displayed for each method.

The error measure chosen for the results displayed below was as follows. Assume that a particular method has estimated an ellipse. The error in this estimate is given by the sum of the shortest distances of each *true point* from the estimated ellipse. Note that this measure takes advantage of the fact that the underlying true points are known. Were these unknown, an alternative measure might be the sum of the so-called *Mahalanobis distances* from the data points to the estimated ellipses.

The methods tested were ordinary least squares scheme (OLS), iteratively reweighted least squares (IRWLS), renormalisation scheme (REN), and the fundamental numerical scheme (FNS). Figure 2 shows the average error obtained when each of the methods was applied to 200 sets of perturbed data for each of 10 different average levels of noise. Each set of perturbed

data contained 55 points generated by the method described earlier. As would be expected, the

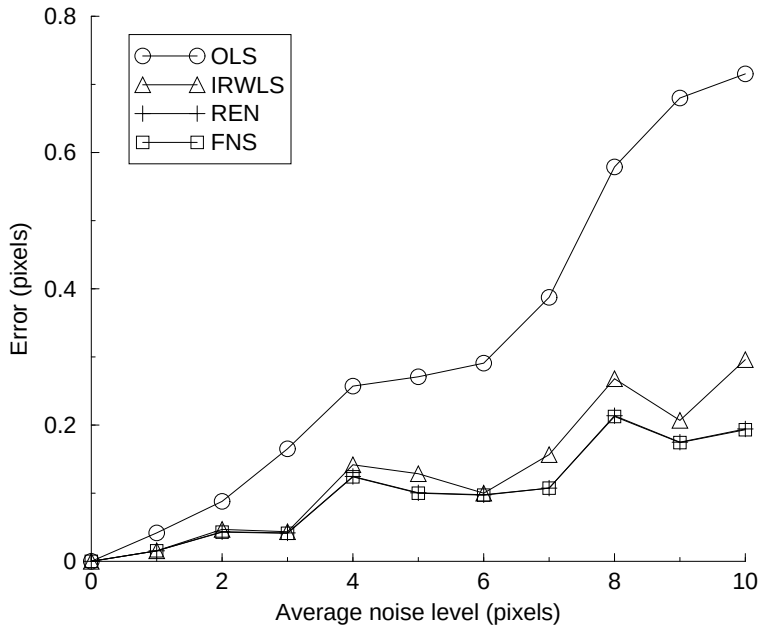


Figure 2: Error results for OLS, IRWLS, REN and FNS.

results show that in the face of data contaminated with inhomogeneous and anisotropic noise, the OLS method, which does not make use of uncertainty (covariance) information, performs markedly worse than the other schemes. The IRWLS scheme lags a little behind with extent determined by problem sensitivity. These results were typical across extensive testing.

# Acknowledgements

The authors are grateful for the insightful comments of Darren Gawley, Marino Ivancic, Kenichi Kanatani, Garry Newsam and Robyn Owens. This work was in part funded by the Australian Research Council.

# References

Bookstein, F. (1979). Fitting conic sections to scattered data. *Computer Vision, Graphics, and Image Processing*, 9(1):56–71.

Brooks, M. J., Chojnacki, W., and Baumela, L. (1997). Determining the egomotion of an uncalibrated camera from instantaneous optical flow. *Journal of the Optical Society of America A*, 14(10):2670–2677.

Chojnacki, W., Brooks, M. J., and van den Hengel, A. (1999). Fitting surfaces to data with covariance information: Fundamental methods applicable to computer vision. Technical Report TR 99-03, Department of Computer Science, University of Adelaide, South Australia. `http://www.cs.adelaide.edu.au/~mjb/mjb_abstracts.html`.

Ellis, T. (1992). Ellipse detection and matching with uncertainty. *Image and Vision Computing*, 10(2):271–276.

Faugeras, O. D. (1993). *Three-Dimensional Computer Vision: A Geometric Viewpoint*. The MIT Press, Cambridge, Mass.

Fitzgibbon, A., Pilu, M., and Fisher, R. B. (1999). Direct least square fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):476–480.

Fitzgibbon, A. W. and Fisher, R. B. (1995). A buyer's guide to conic fitting. In Pycock, D., editor, *Proceedings of the Sixth British Machine Vision Conference*, Birmingham, England, 11-14 September, 1995. BMVA.

Gander, W., Golub, G. H., and Strebel, R. (1994). Least-square fitting of circles and ellipses. *BIT*, 43:558–578.

Kanatani, K. (1994). Statistical bias of conic fitting and renormalisation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(3):320–326.

Kanatani, K. (1996). *Statistical Optimization for Geometric Computation: Theory and Practice*. Elsevier, Amsterdam.

Porill, J. (1990). Fitting ellipses and predicting confidence envelopes using a bias corrected kalman filter. *Image and Vision Computing*, 8(1):37–41.

Rosin, P. L. (1993). A note on the least square fitting of ellipses. *Pattern Recognition Letters*, 14(10):799–808.

Rosin, P. L. and West, G. A. (1995). Nonparametric segmentation of curves into various representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(12):1140–1153.

Sampson, P. D. (1982). Fitting conics sections to 'very scattered' data: An iterative refinement of the Bookstein algorithm. *Computer Graphics and Image Processing*, 18(1):97–108.

Taubin, G. (1991). Estimation of planar curves, surfaces and nonplanar space curves defined by implicit equations, with applications to edge and range image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(11):1115–1138.

Viéville, T. and Faugeras, O. D. (1995). Motion analysis with a camera with unknown, and possibly varying intrinsic parameters. In *Proceedings of the Fifth International Conference on Computer Vision*, pages 750–756, Cambridge, MA, June 20–23, 1995. IEEE Computer Society Press, Los Alamitos, CA.

Zhang, Z. (1997). Parameter estimation techniques: a tutorial with application to conic fitting. *Image and Vision Computing*, 15(1):57–76.