

Minimizing Algebraic Error *

Richard I. Hartley,
G.E. Corporate Research and Development
1 Research Circle
Niskayuna, NY 12309

Abstract

This paper gives a widely applicable technique for solving many of the parameter estimation problems encountered in geometric computer vision. A commonly used approach in such parameter minimization is to minimize an algebraic error function instead of a possibly preferable geometric error function. It is claimed in this paper, however, that minimizing algebraic error will usually give excellent results, and in fact the main problem with most algorithms minimizing algebraic distance is that they do not take account of mathematical constraints that should be imposed on the quantity being estimated. This paper gives an efficient method of minimizing algebraic distance while taking account of the constraints. This provides new algorithms for the problems of resectioning a pinhole camera, computing the fundamental matrix, and computing the tri-focal tensor.

1 Introduction

For many problems related to camera calibration and scene reconstruction, linear algorithms are known for solving for the entity required. In the sort of problem that will be addressed in this paper, a set of data (such as point correspondences) is used to construct a set of linear equations, and solution of these equations provides an estimate of the entity being computed. As examples of such problems we have :

1. The DLT algorithm for computing a camera matrix given a set of points in space, and corresponding points in the image. Provided at least 6 corre-

This work was sponsored by DARPA contract F33615-94-C-1549, monitored by Wright Patterson Airforce Base, Dayton, OH. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the United States Government, or General Electric.

pondences are give (more precisely $5\frac{1}{2}$ correspondences), one can solve for the camera matrix.

2. Computation of the Fundamental Matrix. From 8 point correspondences $\mathbf{u}_i \leftrightarrow \mathbf{u}'_i$ between two images one can construct the fundamental matrix using equations $\mathbf{u}'_i{}^T \mathbf{F} \mathbf{u}_i = 0$.
3. Computation of the trifocal tensor given a set of feature correspondences across three views.

In these three examples, and many others, a linear algorithm exists. However, the linear algorithm will lead to a solution that does not satisfy certain constraints that the estimated quantity must satisfy. In the cases considered here, the constraints are

1. The *skew* parameter of a camera matrix estimated using the DLT method will not generally be zero. This constraint, meaning the pixels are rectangular, should be enforced in cases where it is known to hold.
2. The fundamental matrix must satisfy a constraint $\det F = 0$.
3. The trifocal tensor must satisfy 8 non-linear constraints. The form of these constraints is not easily determined, but it is essential to constrain the tensor to correspond to a valid set of camera matrices.

These constraints are not in general linear constraints, and in general, it will be necessary to resort to iterative techniques to enforce them. Since iterative techniques are slow and potentially unstable, it is important to use them sparingly. Further, the smaller the dimension of the minimization problem, the faster and generally more stable the solution will be. In this paper an iterative algorithm is used to solve the three problems posed above. In each case the algorithms are based on a common technique of data reduction, whereby the input data is condensed into a *reduced measurement matrix*. The size of the iteration problem is then independent on the size of the input set. In the case of estimation of the fundamental matrix, only three homogeneous parameters are

used to parametrize the minimization problem, whereas for the trifocal tensor, just six parameters are used.

The problem of camera calibration solved using the DLT algorithm will be treated first. It will be used to illustrate the techniques that apply to the other problems.

2 Computing the Camera Matrix

We consider a set of point correspondences $\mathbf{x}_i \leftrightarrow \mathbf{u}_i$ between 3D points \mathbf{x}_i and image points \mathbf{u}_i . Our problem is to compute a 3×4 matrix \mathbf{P} such that $\mathbf{P}\mathbf{x}_i = \mathbf{u}_i$ for each i .

2.1 The Direct Linear Transformation (DLT) algorithm

We begin with a simple linear algorithm for determining \mathbf{P} given a set of 3D to 2D point correspondences, $\mathbf{x}_i \leftrightarrow \mathbf{u}_i$. The correspondence is given by the equation $\mathbf{u}_i = \mathbf{P}\mathbf{x}_i$. Note that this is an equation involving homogeneous vectors, thus \mathbf{u}_i and $\mathbf{P}\mathbf{x}_i$ may differ by a non-zero scale factor. One may, however write the equation in terms of the vector cross product as $\mathbf{u}_i \times \mathbf{P}\mathbf{x}_i = \mathbf{0}$.

If the j -th row of the matrix \mathbf{P} is denoted by \mathbf{p}^j , then we may write $\mathbf{P}\mathbf{x}_i = (\mathbf{p}^1 \mathbf{x}_i, \mathbf{p}^2 \mathbf{x}_i, \mathbf{p}^3 \mathbf{x}_i)^\top$. Writing $\mathbf{u}_i = (u_i, v_i, w_i)^\top$, The cross product may then be given explicitly as

$$\mathbf{u}_i \times \mathbf{P}\mathbf{x}_i = \begin{pmatrix} v_i \mathbf{p}^3 \mathbf{x}_i - w_i \mathbf{p}^2 \mathbf{x}_i \\ w_i \mathbf{p}^1 \mathbf{x}_i - u_i \mathbf{p}^3 \mathbf{x}_i \\ u_i \mathbf{p}^2 \mathbf{x}_i - v_i \mathbf{p}^1 \mathbf{x}_i \end{pmatrix}.$$

Since $\mathbf{p}^j \mathbf{x}_i = \mathbf{x}_i^\top \mathbf{p}^j$ for $j = 1, \dots, 3$, this gives a set of three equations, in the entries of \mathbf{P} , which may be written in the form

$$\begin{bmatrix} 0 & -w_i \mathbf{x}_i^\top & v_i \mathbf{x}_i^\top \\ w_i \mathbf{x}_i^\top & 0 & -u_i \mathbf{x}_i^\top \\ -v_i \mathbf{x}_i^\top & u_i \mathbf{x}_i^\top & 0 \end{bmatrix} \begin{pmatrix} \mathbf{p}^1 \\ \mathbf{p}^2 \\ \mathbf{p}^3 \end{pmatrix} = \mathbf{0}. \quad (1)$$

Note that $(\mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3)^\top$ which appears in (1) is a 12-vector made up of the entries of the matrix \mathbf{P} . Although there are three equations, only two of them are linearly independent. Thus each point correspondence gives two equations in the entries of \mathbf{P} . One may choose to omit the third equation, or else include all three equations, which may sometimes give a better conditioned set of equations. In future, we will assume that only the first two equations are used, namely

$$\begin{bmatrix} 0 & -w_i \mathbf{x}_i^\top & v_i \mathbf{x}_i^\top \\ w_i \mathbf{x}_i^\top & 0 & -u_i \mathbf{x}_i^\top \end{bmatrix} \begin{pmatrix} \mathbf{p}^1 \\ \mathbf{p}^2 \\ \mathbf{p}^3 \end{pmatrix} = \mathbf{0}. \quad (2)$$

Solving the Equations. The equations (2) may be denoted by $\mathbf{M}_i \mathbf{p} = \mathbf{0}$, where the vector \mathbf{p} is a 12-vector, corresponding to the 12 entries of \mathbf{P} . The set of all equations derived from several point correspondences may be

written $\mathbf{M}\mathbf{p} = \mathbf{0}$ where \mathbf{M} is the matrix of equation coefficients. This matrix \mathbf{M} will be called the *measurement matrix*. The obvious solution $\mathbf{p} = \mathbf{0}$ is of no interest to us, so we seek a non-zero solution \mathbf{p} .

2.2 Scaling

One of the most important things to do in implementing an algorithm of this sort is to prenormalize the data. This type of data normalization was discussed in the paper [2]. Without this normalization, all these algorithms are guaranteed to perform extremely poorly.

Data normalization is designed to improve the conditioning of the measurement matrix \mathbf{M} . The appropriate scaling is to translate all data points so that their centroid is at the origin. Then the data should be scaled so that the average distance of any data point from the origin is equal to $\sqrt{2}$ for image points and $\sqrt{3}$ for 3D points. The algorithms are then carried out with the normalized data, and final transformations are applied to the result to compensate for the normalizing transforms.

2.3 Algebraic Error

In the presence of noise, one can not expect to obtain an exact solution to an overconstrained set of equations of the form $\mathbf{M}\mathbf{p} = \mathbf{0}$ such as those that arise in the DLT method.

The DLT algorithm instead finds the unit-norm vector \mathbf{p} that minimizes $\|\mathbf{M}\mathbf{p}\|$. The vector $\boldsymbol{\epsilon} = \mathbf{M}\mathbf{p}$ is the error vector and it is this error vector that is minimized. The solution is the unit singular vector corresponding to the smallest singular value of \mathbf{M} .

Define a vector $(\hat{u}_i, \hat{v}_i, \hat{w}_i)^\top = \hat{\mathbf{u}}_i = \mathbf{P}\mathbf{x}_i$. Using this notation, we may write

$$\mathbf{M}_i \mathbf{p} = \boldsymbol{\epsilon}_i = \begin{pmatrix} v_i \hat{w}_i - w_i \hat{v}_i \\ w_i \hat{u}_i - u_i \hat{w}_i \end{pmatrix} = \mathbf{0}. \quad (3)$$

This vector is the *algebraic error vector* associated with the point correspondence $\mathbf{u}_i \leftrightarrow \mathbf{x}_i$ and the camera mapping \mathbf{P} . Thus,

$$d_{\text{alg}}(\mathbf{u}_i, \hat{\mathbf{u}}_i)^2 = (v_i \hat{w}_i - w_i \hat{v}_i)^2 + (w_i \hat{u}_i - u_i \hat{w}_i)^2. \quad (4)$$

Given several point correspondences, the quantity $\boldsymbol{\epsilon} = \mathbf{M}\mathbf{p}$ is the algebraic error vector for the complete set, and one sees that

$$\sum_i d_{\text{alg}}(\mathbf{u}_i, \hat{\mathbf{u}}_i)^2 = \|\mathbf{M}\mathbf{p}\|^2 = \|\boldsymbol{\epsilon}\|^2 \quad (5)$$

The main lesson that we want to keep from this discussion is :

Proposition 2.1. *Given any set of 3D to image correspondences $\mathbf{u}_i \leftrightarrow \mathbf{x}_i$, let \mathbf{M} be the measurement matrix as in (2). For any camera matrix \mathbf{P} the vector $\mathbf{M}\mathbf{p}$ is the algebraic error vector, where \mathbf{p} is the vector of entries of \mathbf{P} .*

2.4 Geometric Distance

Under the assumption that measurement error is confined to image measurements, and an assumption of a gaussian error model for the measurement of 2D image coordinates, the optimal estimate for the camera matrix P is the one that minimizes the error function

$$\sum_i d(\mathbf{u}_i, \hat{\mathbf{u}}_i)^2 \quad (6)$$

where $d(\cdot, \cdot)$ represents Euclidean distance in the image. The quantity $d(\mathbf{u}_i, \hat{\mathbf{u}}_i)$ is known as the *geometric distance* between \mathbf{u}_i and $\hat{\mathbf{u}}_i$. Thus the error to be minimized is the sum of squares of geometric distances between measured and projected points.

For points $\mathbf{u}_i = (u_i, v_i, w_i)^\top$ and $\hat{\mathbf{u}}_i = (\hat{u}_i, \hat{v}_i, \hat{w}_i)^\top$, the geometric distance is

$$\begin{aligned} d(\mathbf{u}_i, \hat{\mathbf{u}}_i) &= ((u_i/w_i - \hat{u}_i/\hat{w}_i)^2 + (v_i/w_i - \hat{v}_i/\hat{w}_i)^2)^{1/2} \\ &= d_{\text{alg}}(\mathbf{u}_i, \hat{\mathbf{u}}_i)/w_i\hat{w}_i \end{aligned} \quad (7)$$

Thus, geometric distance is related to, but not quite the same as algebraic distance. Nevertheless, it will turn out that minimizing algebraic distance gives very good results in general.

2.5 The Reduced Measurement Matrix

Let $\mathbf{u}_i \leftrightarrow \mathbf{x}_i$ be a set of correspondences, and let \mathbf{M} be the corresponding measurement matrix. Let P be any camera matrix, and let \mathbf{p} be the vector containing its entries. The algebraic error vector corresponding to \mathbf{P} is $\mathbf{M}\mathbf{p}$, and its norm satisfies $\|\mathbf{M}\mathbf{p}\|^2 = \mathbf{p}^\top \mathbf{M}^\top \mathbf{M} \mathbf{p}$.

In general, the matrix \mathbf{M} may have a very large number of rows. It is possible to replace \mathbf{M} by a square matrix $\hat{\mathbf{M}}$ such that $\|\mathbf{M}\mathbf{p}\| = \|\hat{\mathbf{M}}\mathbf{p}\|$ for any vector \mathbf{p} . Such a matrix $\hat{\mathbf{M}}$ is called a *reduced measurement matrix*. One way to do this is using the Singular Value Decomposition (SVD). Let $\mathbf{M} = \mathbf{U}\mathbf{D}\mathbf{V}^\top$ be the SVD of \mathbf{M} , and define $\hat{\mathbf{M}} = \mathbf{D}\mathbf{V}^\top$. Then

$$\mathbf{M}^\top \mathbf{M} = (\mathbf{V}\mathbf{D}\mathbf{U}^\top)(\mathbf{U}\mathbf{D}\mathbf{V}^\top) = (\mathbf{V}\mathbf{D})(\mathbf{D}\mathbf{V}^\top) = \hat{\mathbf{M}}^\top \hat{\mathbf{M}}$$

as required. Another way of obtaining $\hat{\mathbf{M}}$ is to use the QR decomposition $\mathbf{M} = \mathbf{Q}\hat{\mathbf{M}}$, where \mathbf{Q} has orthogonal columns and $\hat{\mathbf{M}}$ is upper-triangular and square. This shows the following result.

Theorem 2.2. *Let $\mathbf{u}_i \leftrightarrow \mathbf{x}_i$ be a set of n world to image correspondences. Let \mathbf{M} be the measurement matrix derived from the point correspondences. Let $\hat{\mathbf{M}}$ be a reduced measurement matrix. Then, for any 3D to 2D projective transform \mathbf{P} and corresponding 3-vector \mathbf{p} , one has*

$$\sum_i d_{\text{alg}}(\mathbf{u}_i, \mathbf{P}\mathbf{x}_i)^2 = \|\hat{\mathbf{M}}\mathbf{p}\|^2$$

In this way, all the information we need to keep about the set of matched points $\mathbf{u}_i \leftrightarrow \mathbf{x}_i$ is contained in the

single 12×12 matrix \mathbf{M} . If we wish to minimize algebraic error as \mathbf{P} varies over some restricted set of transforms, then this is equivalent to minimizing the norm of the 12-vector $\|\hat{\mathbf{M}}\mathbf{p}\|$.

2.6 Restricted Camera Mappings

The camera mapping expressed by a general 3D projective transformation is in some respects too general. A non-singular 3×4 matrix P with center at a finite point may be decomposed as $P = K[R \mid -Rt]$ where R is a 3×3 rotation matrix and

$$K = \begin{bmatrix} \alpha_u & s & u_0 \\ & \alpha_v & v_0 \\ & & 1 \end{bmatrix}. \quad (8)$$

The non-zero entries of K are geometrically meaningful quantities, the internal calibration parameters of P . A common assumption is that $s = 0$, while for a true pinhole camera, $\alpha_u = \alpha_v$.

Given a set of world to image correspondences, one may wish to find a matrix P that minimizes algebraic error, subject to a set of constraints on P . Usually, this will require an iterative solution. For instance, suppose we wish to enforce the constraints $s = 0$ and $\alpha_u = \alpha_v$. One can parametrize the camera matrix using the remaining 9 parameters (p_u, p_v, α plus 6 parameters representing the orientation R and location \mathbf{t} of the camera). Let this set of parameters be denoted collectively by \mathbf{q} . Then, one has a map $\mathbf{p} = g(\mathbf{q})$, where \mathbf{p} is as before the vector of entries of the matrix \mathbf{P} . According to Theorem 2.2, minimizing algebraic error over all point matches is equivalent to minimizing $\|Mg(\mathbf{q})\|$. Note that the mapping $\mathbf{q} \mapsto Mg(\mathbf{q})$ is a mapping from R^9 to R^{12} . This is a simple parameter-minimization problem that may be solved using the Levenberg-Marquardt method. The important point to note is the following :

Given a set of n world-to-image correspondences, $\mathbf{x}_i \leftrightarrow \mathbf{u}_i$, the problem of finding a constrained camera matrix \mathbf{P} that minimizes the sum of algebraic distances $\sum_i d_{\text{alg}}(\mathbf{u}_i, \mathbf{P}\mathbf{x}_i)^2$ reduces to the minimization of a function $R^9 \rightarrow R^{12}$, independent of the number n of correspondences.

If this problem is solved using the Levenberg-Marquardt (LM) method, then an initial estimate of the parameters may be obtained by decomposing a camera matrix P found using the DLT algorithm. A central step in the LM method is the computation of the derivative matrix (Jacobian matrix) of the function being minimized, in this case $Mg(\mathbf{q})$. Note that $\partial Mg/\partial \mathbf{q} = M\partial g/\partial \mathbf{q}$. Thus, computation of the Jacobian reduces to computation of the Jacobian matrix of g , and subsequent multiplication by M .

Minimization of $\|Mg(\mathbf{q})\|$ takes place over all values of the parameters \mathbf{q} . Note, however, that if $P = K[R \mid$

$-Rt]$ with K as in (8) then P satisfies the condition $p_{31}^2 + p_{32}^2 + p_3^2 = 1$, since these entries are the same as the last row of the rotation matrix R . Thus, minimizing $Mg(\mathbf{q})$ will lead to a matrix P satisfying the constraints $s = 0$ and $k_u = k_v$ and scaled such that $p_{31}^2 + p_{32}^2 + p_3^2 = 1$, and which in addition minimizes the algebraic error for all point correspondences.

2.7 Experimental Evaluation

Experiments were carried out with synthetic data to evaluate the performance of this algorithm. The data were created to simulate a standard 35mm camera with a 35mm focal length lens. A set of points were synthesized inside a sphere of radius 1m, and the camera was located at a distance of about 2.5m from the centre of the sphere. The image is sampled so that the magnification factors are $\alpha_u = \alpha_v = 1000.0$, the same in each direction. This corresponds to a pixel size of $35\mu\text{m}$ for a 35mm camera.

Experiments were carried out to find the camera matrix with four different assumptions on known camera parameters.

1. The pixels are square : $s = 0$. The number of remaining degrees of freedom d for the camera matrix is equal to 10.
2. The pixels are square and the pixels are square : $s = 0$ and $\alpha_u = \alpha_v$. This corresponds to the situation for a true pinhole camera where image coordinates are measured in a Euclidean coordinate frame. In this case, $d = 9$.
3. In addition to the above assumptions, the principal point (u_0, v_0) is assumed to be known. There remain $d = 7$ degrees of freedom.
4. The complete internal calibration matrix K in (8) is assumed to be known. However, the pose of the camera is unknown. Thus $d = 6$.

To evaluate the performance of the algorithm, the result was compared with the optimal estimate with different degrees of noise. Thus, gaussian noise with a given variance was added to the image coordinates of each point, and the camera matrix was estimated. The residual error was then computed, that is the difference between the measured and projected pixel. It is known (see [?]) that the expected lower bound on the root mean squared residual error is equal to $\sigma(1 - d/N)^{1/2}$ where σ is the standard deviation of the input noise, N is the number of measurements (in this case $2 \times$ number of points), and d is the number of degrees of freedom of the object being estimated. This represents the performance of an optimal estimation technique, and we can not do better. Since the residual error appeared to grow proportionally to injected noise (at least for noise levels less than about 10 pixels), a value of $v = 1$ pixel was

used in the experiments. For each level of noise σ , the camera matrix was estimated 100 times with different noise. The residual error was averaged over all 100 runs and compared to the optimal value.

Results of the experiments are shown in Fig 1. The results show that minimizing algebraic error gives an almost optimal estimate of the camera matrix. In fact, the residual error is scarcely distinguishable from the optimal value. This is true in each of the four calibration problem types tried.

3 Computation of the Fundamental Matrix

We now turn to the computation of the Fundamental Matrix. It will turn out that very similar methods apply to its computation as were used in the DLT algorithm.

Given a set of correspondences $\mathbf{u}_i \leftrightarrow \mathbf{u}'_i$ between two images, the fundamental matrix is defined by the relation $\mathbf{u}'_i{}^T \mathbf{F} \mathbf{u}_i = 0$ for all i . In the presence of noise, this relation will not hold precisely, and so one seeks a least-squares solution. Note that the equation $\mathbf{u}'_i{}^T \mathbf{F} \mathbf{u}_i = 0$ is linear in the entries of \mathbf{F} . From 8 or more point matches, one may solve for the entries of \mathbf{F} by finding the least-squares solution to a set of linear equations ([2]). Let the set of equations be denoted by $M\mathbf{f} = 0$. The vector $M\mathbf{f}$ has components equal to $\mathbf{u}'_i{}^T \mathbf{F} \mathbf{u}_i$, and $\|M\mathbf{f}\|^2 = \sum_i (\mathbf{u}'_i{}^T \mathbf{F} \mathbf{u}_i)^2$. Thus, in this case, as before with the DLT algorithm, $M\mathbf{f}$ represents the algebraic error vector. Matrix \mathbf{F} is found by minimizing $\|M\mathbf{f}\|$ subject to $\|\mathbf{f}\| = 1$.

The fundamental matrix \mathbf{F} must, however satisfy a constraint $\det \mathbf{F} = 0$, and this constraint will not generally be satisfied by the matrix \mathbf{F} found by this linear algorithm. One would therefore like to minimize the algebraic error $\|M\hat{\mathbf{f}}\|$ over all vectors $\hat{\mathbf{f}}$ corresponding to singular matrices $\hat{\mathbf{F}}$.

In [2], the matrix $\hat{\mathbf{F}}$ was taken to be the closest singular matrix to \mathbf{F} under Frobenius norm, where \mathbf{F} is the linear solution. This is not an especially good way of proceeding, since it weights errors in each of the entries of \mathbf{F} equally. A preferable method is to proceed as with the DLT. One parametrizes the matrix $\hat{\mathbf{F}}$ by a set of parameters \mathbf{q} in a way so as to ensure it is singular. Then letting $\hat{\mathbf{f}} = g(\mathbf{q})$, one uses an iterative algorithm to minimize $\|Mg(\mathbf{q})\|$. This is the general scheme which will be followed, but there are details to be filled out, and a new twist will arise, which allows a parametrization with only three parameters.

3.1 Parametrization of the Fundamental Matrix

Consider the fundamental matrix \mathbf{F} , which can be written as a product $\mathbf{F} = \mathbf{Q}[\mathbf{e}]_{\times}$ where \mathbf{Q} is a non-singular matrix, and \mathbf{e} is the epipole in the first image.

Suppose we wish to compute the fundamental matrix \mathbf{F} of the form $\mathbf{F} = \mathbf{Q}[\mathbf{e}]_{\times}$ that minimizes the algebraic error $\|M\mathbf{f}\|$ subject to the condition $\|\mathbf{f}\| = 1$. The vector

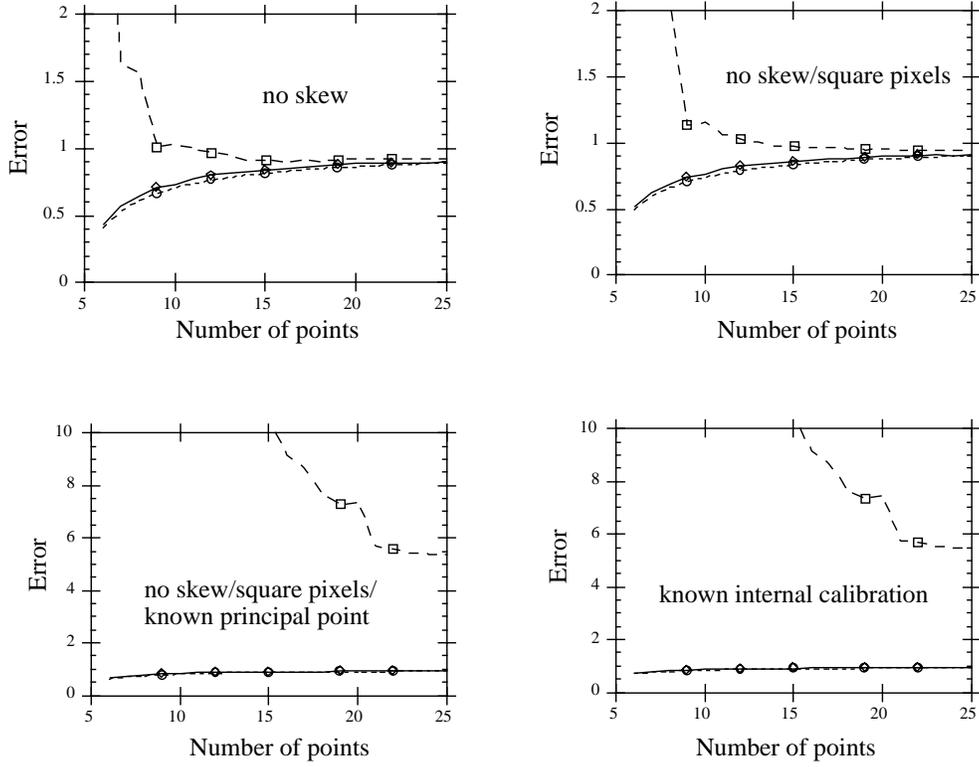


Figure 1: *The residual error was averaged over 100 runs for each $n = 6, \dots, 25$, where n is the number of points used to estimate the camera matrix. Four different levels of knowledge of the internal camera matrix were tried, corresponding to the four different graphs. In each of the graphs, the solid line represents the result of our iterative DLT algorithm, and the almost identical dotted line is the optimal estimate. In all four graphs, these two lines are barely distinguishable. For comparison, the results of a further method are also plotted. In this method, the complete calibration matrix K in (8) is estimated using the QR decomposition, and the known internal parameters are subsequently set to their known values. This method performs very poorly for small numbers of points, lying well off the graph, and is markedly inferior to the optimal estimation method even for larger numbers of points.*

\mathbf{f} is the 9-vector containing the entries of \mathbf{F} . It has been seen that the 8-point algorithm finds such an \mathbf{f} , without the condition that $\mathbf{F} = \mathbf{Q}[\mathbf{e}]_{\times}$. We now wish to enforce that condition.

Let us assume for now that the epipole \mathbf{e} is known. Later we will let \mathbf{e} vary, but for now it is fixed. The equation $\mathbf{F} = \mathbf{Q}[\mathbf{e}]_{\times}$ can be written in terms of the vectors \mathbf{f} and \mathbf{q} comprising the entries of \mathbf{F} and \mathbf{Q} as an equation $\mathbf{f} = \mathbf{E}\mathbf{q}$ where \mathbf{E} is a 9×9 matrix. Supposing that \mathbf{f} and \mathbf{q} contain the entries of the corresponding matrices in row-major order, then it can be verified that \mathbf{E} has the form

$$\mathbf{E} = \begin{bmatrix} [\mathbf{e}]_{\times} & & \\ & [\mathbf{e}]_{\times} & \\ & & [\mathbf{e}]_{\times} \end{bmatrix}. \quad (9)$$

Now, our minimization problem is : minimize $\|\mathbf{M}\mathbf{E}\mathbf{q}\|$ subject to the condition $\|\mathbf{E}\mathbf{q}\| = 1$.¹ This problem is solved as follows. Let the Singular Value Decomposition of \mathbf{E} be $\mathbf{E} = \mathbf{U}\mathbf{D}\mathbf{V}^{\top}$. It is easily seen that the matrix \mathbf{E} has rank 6, since each of the diagonal blocks has rank 2. It follows that \mathbf{D} has 6 non-zero diagonal entries. Let \mathbf{U}' be the 9×6 matrix consisting of the first 6 columns of \mathbf{U} , and let \mathbf{V}' consist of the first 6 columns of \mathbf{V} and let \mathbf{D}' be the top-left 6×6 minor of \mathbf{D} , containing the non-zero diagonal entries. The minimization problem then becomes : minimize $\|\mathbf{M}\mathbf{U}'\mathbf{D}'\mathbf{V}'^{\top}\mathbf{q}\|$ subject to $\|\mathbf{U}'\mathbf{D}'\mathbf{V}'^{\top}\mathbf{q}\| = 1$. This last condition is equivalent to $\|\mathbf{D}'\mathbf{V}'^{\top}\mathbf{q}\| = 1$, since \mathbf{U}' has orthogonal columns. Now writing $\mathbf{q}' = \mathbf{D}'\mathbf{V}'^{\top}\mathbf{q}$, the problem becomes : minimize $\|\mathbf{M}\mathbf{U}'\mathbf{q}'\|$ subject to $\|\mathbf{q}'\| = 1$, which is our standard minimization problem. The solution \mathbf{q}' is the singular vector corresponding to the smallest singular value of $\mathbf{M}\mathbf{U}'$. Subsequently, we can compute $\mathbf{f} = \mathbf{E}\mathbf{q} = \mathbf{U}'\mathbf{D}'\mathbf{V}'^{\top}\mathbf{q} = \mathbf{U}'\mathbf{q}'$, and the algebraic error is $\mathbf{M}\mathbf{f} = \mathbf{M}\mathbf{U}'\mathbf{q}'$.

The complete algorithm is :

Algorithm 3.3. Given the epipole \mathbf{e} , find the fundamental matrix \mathbf{F} of the form $\mathbf{F} = \mathbf{Q}[\mathbf{e}]_{\times}$ that minimizes the algebraic error $\|\mathbf{M}\mathbf{f}\|$ subject to $\|\mathbf{f}\| = 1$.

Solution :

1. Compute the SVD $\mathbf{E} = \mathbf{U}\mathbf{D}\mathbf{V}^{\top}$, where \mathbf{E} is given in (9).
2. Let \mathbf{U}' be the matrix comprising the first 6 columns of \mathbf{U}
3. Find the unit vector \mathbf{q}' that minimizes $\|\mathbf{M}\mathbf{U}'\mathbf{q}'\|$.
4. The required matrix \mathbf{F} corresponds to the vector $\mathbf{f} = \mathbf{U}'\mathbf{q}'$, and the minimum algebraic error is $\mathbf{M}\mathbf{f}$.

¹It does not do to minimize $\|\mathbf{M}\mathbf{E}\mathbf{q}\|$ subject to the condition $\|\mathbf{q}\| = 1$, since a solution to this occurs when \mathbf{q} is a unit vector in the right null-space of \mathbf{E} . In this case, $\mathbf{E}\mathbf{q} = 0$, and hence $\|\mathbf{M}\mathbf{E}\mathbf{q}\| = 0$.

3.2 Iterative Estimation

The algorithm of the last section gives a way of computing an algebraic error vector $\mathbf{M}\mathbf{f}$ given a value for the epipole \mathbf{e} . This mapping $\mathbf{e} \mapsto \mathbf{M}\mathbf{f}$ is a map from R^3 to R^9 . Note that the value of $\mathbf{M}\mathbf{f}$ is unaffected by scaling \mathbf{e} . Starting from a value of \mathbf{e} derived as the generator of the right null-space of an initial estimate of \mathbf{F} , one may iterate to find the final \mathbf{F} that minimizes algebraic error. The initial estimate of \mathbf{F} may be obtained from the 8-point algorithm, or any other simple algorithm.

Note the advantage of this method of computing \mathbf{F} is that the iterative part of the algorithm consists of a very small parameter minimization problem, involving the estimation of only three parameters. Despite this, the algorithm finds the fundamental matrix that minimizes the algebraic error for all matched points. The matched points themselves do not come into the final iterative estimation.

Simplifying the computation Because of the simple form of the matrix \mathbf{E} , it is easy to compute its SVD without having to resort to a full SVD algorithm. This may be important in the iterative algorithm to achieve maximum speed, since this SVD is computed repeatedly during the minimization. As seen in (9), the matrix \mathbf{E} has a diagonal block structure consisting of three blocks $[\mathbf{e}]_{\times}$. The SVD consequently has a corresponding block-structure. Specifically, if $[\mathbf{e}]_{\times} = \hat{\mathbf{U}}\hat{\mathbf{D}}\hat{\mathbf{V}}'$, then the SVD of $\mathbf{E} = \text{diag}([\mathbf{e}]_{\times}, [\mathbf{e}]_{\times}, [\mathbf{e}]_{\times})$ is $\mathbf{E} = \mathbf{U}\mathbf{D}\mathbf{V}^{\top}$ where $\mathbf{U} = \text{diag}(\hat{\mathbf{U}}, \hat{\mathbf{U}}, \hat{\mathbf{U}})$, and similarly for \mathbf{D} and \mathbf{V} .

The SVD of $[\mathbf{e}]_{\times}$ itself can be computed easily as follows. Suppose that $\hat{\mathbf{U}}$ is an orthogonal matrix such that $\mathbf{e}\hat{\mathbf{U}} = (0, 0, 1)$. Such a matrix $\hat{\mathbf{U}}$ is a Householder transformation and is easily computed ([1]). Then one sees that $[\mathbf{e}]_{\times} = \pm\hat{\mathbf{U}}\hat{\mathbf{Z}}\hat{\mathbf{U}}^{\top} = \pm\hat{\mathbf{U}}\text{diag}(1, 1, 0)\hat{\mathbf{Z}}\hat{\mathbf{U}}^{\top} = \pm\hat{\mathbf{U}}\hat{\mathbf{D}}\hat{\mathbf{V}}^{\top}$ where

$$\mathbf{z} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad ; \quad \hat{\mathbf{z}} = \begin{bmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

This is easily verified by observing that both $[\mathbf{e}]_{\times}$ and $\pm\hat{\mathbf{U}}\hat{\mathbf{Z}}\hat{\mathbf{U}}^{\top}$ are skew-symmetric matrices with the same null-space, generated by \mathbf{e} in each case. We are interested in $\hat{\mathbf{U}}'$ consisting of the first two columns of $\hat{\mathbf{U}}$. Turning now to the SVD of $\mathbf{E} = \text{diag}([\mathbf{e}]_{\times}, [\mathbf{e}]_{\times}, [\mathbf{e}]_{\times})$, we see that $\mathbf{U}' = \text{diag}(\hat{\mathbf{U}}', \hat{\mathbf{U}}', \hat{\mathbf{U}}')$. If we partition the 9×9 matrix \mathbf{M} into blocks $\mathbf{M} = [\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3]$ where each \mathbf{M}_i has 3 columns, then one computes that $\mathbf{M}\mathbf{U}' = [\mathbf{M}_1\hat{\mathbf{U}}', \mathbf{M}_2\hat{\mathbf{U}}', \mathbf{M}_3\hat{\mathbf{U}}']$. Thus, the computation of $\mathbf{M}\mathbf{U}'$ required in Algorithm 3.3 has two simple steps

1. Compute the 3×3 Householder matrix $\hat{\mathbf{U}}$ such that $\mathbf{e}^{\top}\hat{\mathbf{U}} = (0, 0, 1)$, and let $\hat{\mathbf{U}}'$ comprise its first two columns.
2. Set $\mathbf{M}\mathbf{U}' = [\mathbf{M}_1\hat{\mathbf{U}}', \mathbf{M}_2\hat{\mathbf{U}}', \mathbf{M}_3\hat{\mathbf{U}}']$.

3.3 Experimental Evaluation of the Algorithm

A set of experiments were carried out similar to those in [2]. One image from each pair of images used is shown in Fig 2. These images contain a wide variation of measurement noise and placement of the epipoles. For each pair of images, a number n of matched points were chosen and the fundamental matrix was computed. The fundamental matrix computed was shown evaluated against the full set of all matched points, and the residual error was computed. This experiment was done 100 times for each value of n and each pair of images, and the average residual error was plotted against n . This gives an idea of how the different algorithms behave as the number of points is increased.

The results of these experiments are shown and explained in Fig 3. They show that minimizing algebraic error gives essentially indistinguishable results from minimizing the geometric error, but both perform better than the linear normalized 8-point algorithm ([2]).

4 Computation of the Trifocal Tensor

The trifocal tensor ([5, 3]), relates the coordinates of points or lines seen in three views in a similar way to that in which the fundamental matrix relates points in two views.

The basic formula relates a point \mathbf{u} in one image and a pair of lines λ' and λ'' in the other two images. Provided there is a point \mathbf{x} in space that maps to \mathbf{u} in the first image and a point on the lines λ' and λ'' in the other two images, the following identity is satisfied :

$$u^i \lambda_j^i \lambda_k^i T_i^{jk} = 0 . \quad (10)$$

Here we are using tensor notation, in which a repeated index appearing in covariant (lower) and contravariant (upper) positions implies summation over the range of indices (namely, $1, \dots, 3$).

This equation may be used to generate equations given either point or line correspondences across three images. In the case of a line correspondence, $\lambda \leftrightarrow \lambda' \leftrightarrow \lambda''$ one selects two points \mathbf{u}_0 and \mathbf{u}_1 on the line λ , and for each of these points one obtains an equation of the form (10). In the case of a point correspondence $\mathbf{u} \leftrightarrow \mathbf{u}' \leftrightarrow \mathbf{u}''$ one selects any lines λ' and λ'' passing through \mathbf{u}' and \mathbf{u}'' respectively. Then (10) provides one equation. Four equations are generated from a single 3-view point correspondence by choosing two lines through each of \mathbf{u}' and \mathbf{u}'' , each pair of lines giving rise to a single equation.

The equations (10) give rise to a set of equations of the form $\mathbf{M}\mathbf{t} = 0$ in the 27 entries of the trifocal tensor. From these equations, one may solve for the entries of the tensor. As before, for any tensor T_i^{jk} the value of $\mathbf{M}\mathbf{t}$ is the algebraic error vector associated with the input data.

Consider the analogy with the 8-point algorithm for computing the fundamental matrix in the two-view case.

The fundamental matrix is not in general precisely satisfied by the solution found from linear algorithm. In the case of the trifocal tensor, there are 27 entries in the tensor, but the camera geometry that it encodes has only 18 degrees of freedom. This means that the trifocal tensor must satisfy 8 constraints, apart from scale ambiguity to make up the 27 degrees of freedom of a general $3 \times 3 \times 3$ tensor. The exact form of these constraints is not known precisely. Nevertheless, they must be enforced in order that the trifocal tensor should be well behaved. It will now be shown how this can be done, while minimizing algebraic error.

Formula for the Trifocal Tensor. We denote the three camera matrices P' and P'' by a_j^i and b_j^i respectively, instead of by p_j^i and $p_j^{i'}$. Thus, the three camera matrices P , P' and P'' may be written in the form $P = [I \mid 0]$, $P' = [a_j^i]$ and $P'' = [b_j^i]$.

In this notation, the formula for the entries of the trifocal tensor is :

$$T_i^{jk} = a_i^j b_4^k - a_4^j b_i^k . \quad (11)$$

Our task will be to compute a trifocal tensor T_i^{jk} of this form from a set of image correspondences. The tensor computed will minimize the algebraic error associated with the input data. The algorithm is quite similar to the one given for computation of the fundamental matrix. Just as with the fundamental matrix, the first step is the computation of the epipoles.

4.1 Retrieving the epipoles

We consider the task of retrieving the epipoles from the trifocal tensor. If the first camera has matrix $P = [I \mid \mathbf{0}]$, then the epipoles \mathbf{e}_{21} and \mathbf{e}_{31} are the last columns a_4^i and b_4^i of the two camera matrices $P' = [a_j^i]$ and $P'' = [b_j^i]$ respectively. These two epipoles may easily be computed from the tensor T_i^{jk} according to the following proposition.

Proposition 4.4. *For each $i = 1, \dots, 3$, the matrix T_i^{jk} is singular. Furthermore, the generators of the three left null-spaces have a common perpendicular, the epipole \mathbf{e}_{21} . Similarly epipole \mathbf{e}_{31} is the common perpendicular of the right nullspaces of the three matrices T_i^{jk} .*

This proposition translates easily into an algorithm for computing the epipoles ([5, 3]). This algorithm may be applied to the tensor T_i^{jk} obtained from the linear algorithm to obtain a reasonable approximation for the epipoles.

4.2 Constrained Estimation of the Trifocal Tensor

From the form (11) of the trifocal tensor, it may be seen that once the epipoles $\mathbf{e}_{21} = a_4^j$ and $\mathbf{e}_{31} = b_4^k$ are known, the trifocal tensor may be expressed linearly in terms of the remaining entries of the matrices a_j^i and b_j^i .

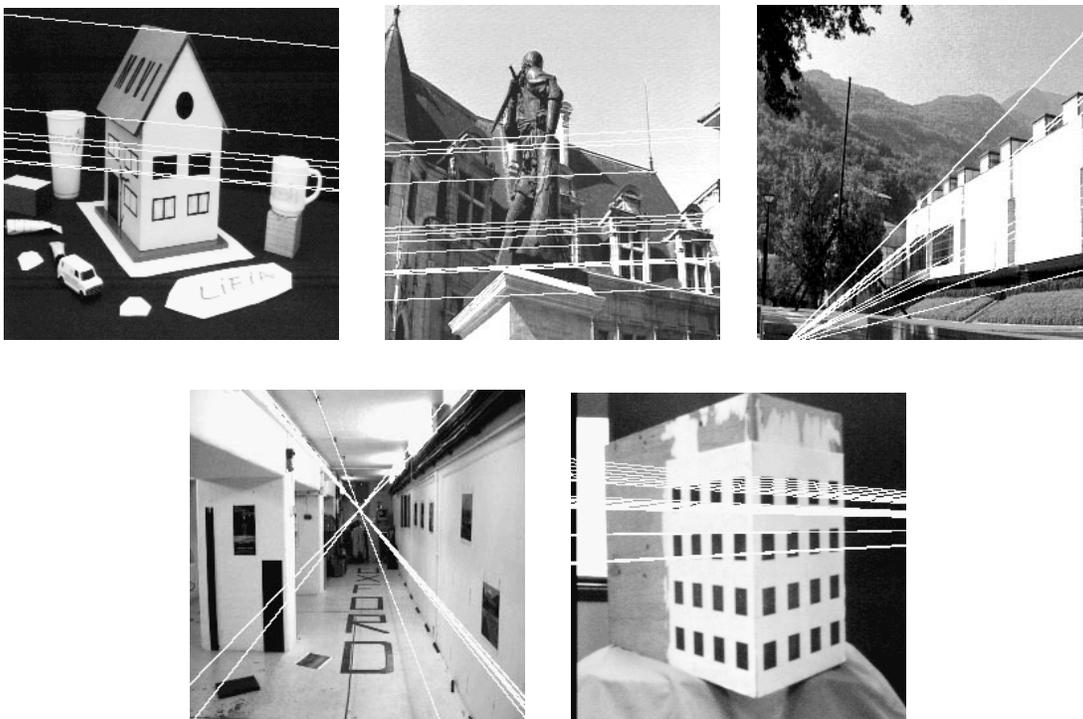


Figure 2: The images used in the experiments

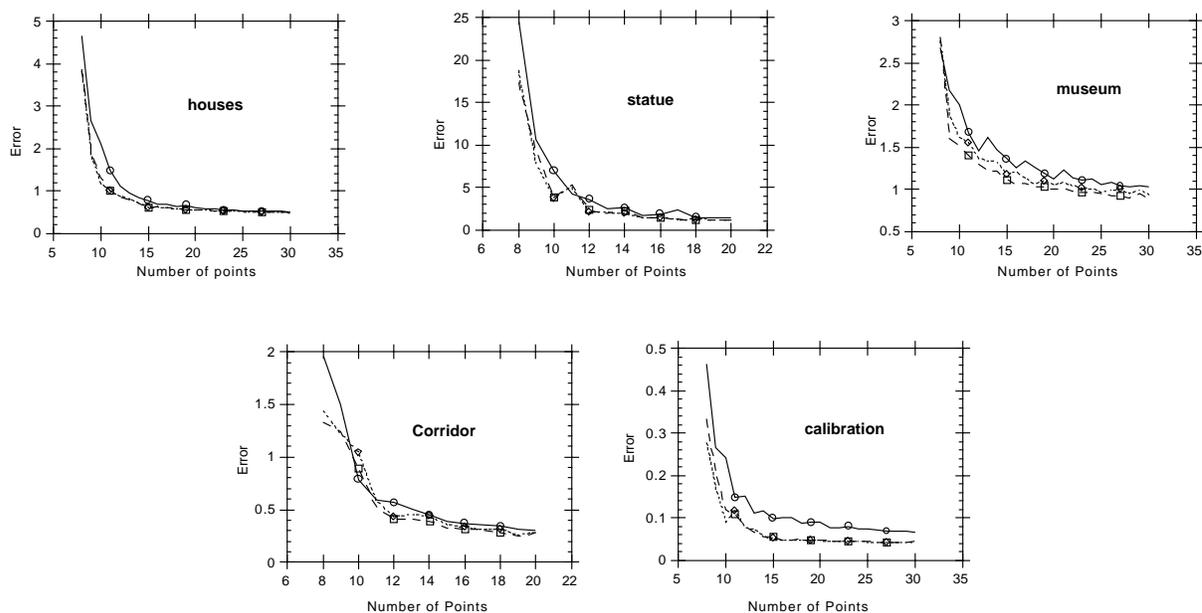


Figure 3: Results of the experimental evaluation of the algorithms. In each case, three methods of computing F were compared. In each graph, the top (solid) line shows the results of the normalized 8-point algorithm. Also shown are the results of minimizing geometric error and algebraic error, using the algorithm of this paper. In most cases, the result of minimizing algebraic error is almost indistinguishable from minimizing geometric error. Both are noticeably better than the non-iterative 8-point algorithm, though that algorithm gives reasonable results.

Assuming the epipoles a_4^i and b_4^i to be known, we may write $\mathbf{t} = H\mathbf{a}$ where \mathbf{a} is the vector of the remaining entries a_j^i and b_j^i , \mathbf{t} is the vector of entries of the trifocal tensor, and H is the linear relationship expressed by (11). We wish to minimize the algebraic error $\|M\mathbf{t}\| = \|MH\mathbf{a}\|$ over all choices of \mathbf{a} constrained such that $\|\mathbf{a}\| = 1$. The solution is the eigenvector corresponding to the least eigenvalue of $H^\top M^\top MH$.

In solving this set of equations to find \mathbf{a} it is advisable to restrict the dimensionality of the solutions set by applying the constraint that $\sum_i a_4^i a_j^i = 0$ for each $j = 1, \dots, 3$. This constraint is discussed in [5, 3]. Given that a_4^i is known, it is a linear constraint that may be expressed by a matrix equation $C\mathbf{a} = 0$. Thus, the minimization problem is to minimize $\|MH\mathbf{a}\|$ subject to the $\|\mathbf{a}\| = 1$ and the linear constraint $C\mathbf{a} = 0$. This may be done by the algorithm given in [5, 3].

Writing $\hat{\mathbf{t}} = H\mathbf{a}$ where \mathbf{a} is the solution vector, we see that $\hat{\mathbf{t}}$ minimizes algebraic error $\|M\hat{\mathbf{t}}\|$ subject to the condition that T_i^{jk} is of the correct form (11), for the given choice of epipoles.

Iterative Solution The two epipoles used to compute a correct constrained tensor T_i^{jk} are computed using the estimate of T_i^{jk} obtained from the linear algorithm. Analogous to the case of the fundamental matrix, the mapping $(\mathbf{e}_{21}, \mathbf{e}_{31}) \mapsto M\mathbf{H}\mathbf{a}$ is a mapping $R^6 \rightarrow R^{27}$. An application of the Levenberg-Marquardt algorithm to optimize the choice of the epipoles will result in an optimal (in terms of algebraic error) estimate of the trifocal tensor. Note that the iteration problem is of modest size, since only 6 parameters, the homogeneous coordinates of the epipoles, are involved in the iteration problem.

This contrasts with an iterative estimation of the optimal trifocal tensor in terms of geometric error. This latter problem would require estimating the three camera parameters, plus the coordinates of all the points, a large estimation problem.

5 Conclusion

Experimental evidence backs up the assertion that minimizing algebraic distance can usually give good results at a fraction of the computation cost associated with minimizing geometric distance. The great advantage of the method for minimizing algebraic error given in this paper is that even for problems that need an iterative solution the size of the iteration problem is very small. Consequently, the iteration is very rapid and there is reduced risk of falling into a local minimum, or otherwise failing to converge.

The method has been illustrated by applying it to three problems. For the computation of the fundamental matrix, iteration is over only three homogeneous parameters. For the trifocal tensor, iteration is over 6 parameters. This leads to more efficient methods than have

been previously known.

The general technique is applicable to problems other than those treated here. It may be applied in a straightforward manner to estimation of projective transformations between 2 or 3-dimensional point sets. In these problems, iteration is necessary if one restricts the class of available transformations to a subgroup of the projective group, such as planar homologies (used in [6]), or conjugates of rotations ([4]).

References

- [1] Gene H. Golub and Charles F. Van Loan. *Matrix Computations, Second edition*. The Johns Hopkins University Press, Baltimore, London, 1989.
- [2] R. I. Hartley. In defence of the 8-point algorithm. In *Proc. International Conference on Computer Vision*, pages 1064 – 1070, 1995.
- [3] R. I. Hartley. A linear method for reconstruction from lines and points. In *Proc. International Conference on Computer Vision*, pages 882 – 887, 1995.
- [4] Richard I. Hartley. Self-calibration from multiple views with a rotating camera. In *Computer Vision - ECCV '94, Volume I, LNCS-Series Vol. 800, Springer-Verlag*, pages 471–478, May 1994.
- [5] Richard I. Hartley. Lines and points in three views and the trifocal tensor. *International Journal of Computer Vision*, 22(2):125–140, March 1997.
- [6] A. Zisserman, D. A. Forsyth, J. L. Mundy, C. A. Rothwell, J. Liu, and N. Pillow. 3D object recognition using invariance. *AI Journal*, 78:239 – 288, 1995.