Robust Parameterization and Computation of the Trifocal Tensor

P H S Torr, A Zisserman
Robotics Research Group
Department of Engineering Science
Oxford University, Parks Road, Oxford, OX1 3PJ, UK.

Abstract

The constraint that rigid motion places on the image positions of points and lines over three views is captured by the trifocal tensor. This paper demonstrates a novel robust estimator of the trifocal tensor, based on a minimum number of correspondences across an image triplet. In addition, it is shown how the robust estimate can be used to find a minimal parameterization that enforces the constraints between the elements of the tensor.

The matching techniques used to estimate the tensor are both robust (detecting and discarding mismatches) and fully automatic. Results are given for real image sequences.

1 Introduction

The trifocal tensor plays a similar role for three views to that played by the fundamental matrix for two. It encapsulates all the (projective) geometric constraints between three views that are independent of scene structure. The tensor only depends on the motion between views and the internal parameters of the cameras, but it can be computed from image correspondences alone without requiring knowledge of the motion or calibration. It is the culmination of developments by a number of researchers including [4, 7, 12, 13, 17, 19]. Recently the trifocal tensor has been used for applications in structure from motion including tracking [2], camera calibration [1], and motion segmentation [16].

Given correspondences for points in two images, the tensor determines the position of the point in the third (this is known as *transfer*). Similarly, given correspondences for lines in two images, the same tensor determines the position of the line in the third image. Unlike epipolar transfer [5, 20] transfer based on the trifocal tensor does not fail for 3D points lying on, or close to, the trifocal plane (the plane defined by the three optical centres of the cameras), or when the three optical centres are collinear. The algebraic properties of the tensor are reviewed in Section 2.

This paper extends the state of the art in two ways. First, a robust estimator is developed for the trifocal tensor based on 6 point correspondences over three views. Second, a simple parameterisation is described which enforces the algebraic constraints existing between the elements of the tensor.

BMVC 1996 doi:10.5244/C.10.61

Notation For a triplet of images the image of a 3D point $\mathbf{X}$ is $\mathbf{x}$, $\mathbf{x}^{'}$ and $\mathbf{x}^{''}$ in the first, second and third images respectively, and similarly the image of a line is $\mathbf{l}$, $\mathbf{l}^{'}$ and $\mathbf{l}^{''}$. Where $\mathbf{x} = (x_1, x_2, x_3)^{\top}$ and $\mathbf{l} = (l_1, l_2, l_3)^{\top}$ are homogeneous three vectors. The 3×4 camera projection matrix which relates a 3D point $\mathbf{X}$ to its image $\mathbf{x}$ is denoted as $\mathtt{P}$, i.e. $\mathbf{x} = \mathtt{P}\mathbf{X}$, where $\mathbf{X}$ is an homogeneous four vector.

2 Trifocal tensor - review

Corresponding points in three images, and corresponding lines in three images, satisfy trilinear relations which are encapsulated in the trifocal tensor.

The trifocal tensor, $\mathtt{T}$, is a $3 \times 3 \times 3$ homogeneous tensor. Using the tensor a point can be transferred to a third image from correspondences in the first and second:

$$x_l^{''} = x_i^{'} \sum_{k=1}^{k=3} x_k T_{kjl} - x_j^{'} \sum_{k=1}^{k=3} x_k T_{kil} \quad , \tag{1}$$

for all $i, j = 1 \ldots 3$—nine expressions. Similarly, a line can be transferred as $l_i = \sum_{j=1}^{j=3} \sum_{k=1}^{k=3} l_j^{'} l_k^{''} T_{ijk}$ i.e. the same tensor can be used to transfer both points and lines.

If the first camera matrix is chosen as $P = [I|\mathbf{0}]$ (where I is the 3×3 identity matrix, and $\mathbf{0}$ is a null three-vector), then the trifocal tensor can be computed from:

$$T_{ijk} = P_{ji}^2 P_{k4}^3 - P_{j4}^2 P_{ki}^3 \ , \tag{2}$$

where P_{ij}^w is the ijth element of the camera matrix of the wth image.

2.1 Degrees of Freedom

The trifocal tensor has 27 elements, but only their ratios are significant, leaving 26 that must be specified. Each triplet of point correspondences provides four independent linear equations for the elements of the tensor, and each triplet of line correspondences provides two linear equations. Therefore provided that $2n_l + 4n_p \geq 26$ (where n_l is the number of lines, and n_p is the number of points), the tensor can be determined uniquely (up to scale) using a linear algorithm. Consequently, the tensor can be computed linearly from a minimum of 7 points or 13 lines or a combination of the two. As in the case of the fundamental matrix [3, 15], the tensor can be estimated from more than the minimum number of correspondences by minimising a cost function. A solution may be computed by eigenvector methods by finding the eigenvector with least eigenvalue of a 27×27 matrix.

However, the tensor has only 18 independent degrees of freedom (this can be seen by considering three 3×4 projection matrices, less 15 projective degrees of freedom, i.e. $3 \times 11 - 15 = 18$). Consequently the 27 elements of the tensor satisfy a number of (polynomial) constraints. This is similar to the situation with the fundamental matrix — there are 9 elements in the 3×3 matrix, but only 7 degrees of freedom (again this follows from $2 \times 11 - 15 = 7$). In the case of the fundamental matrix only the ratio of the elements is significant – 8 ratios – and there is one cubic constraint that the determinant is zero. For the trifocal tensor the 8 constraints have not been as thoroughly explored.

In the case of the fundamental matrix if the constraint is not imposed then the epipolar lines do not all intersect in a single epipole [10]. Similarly, if the

1. Repeat for $m = 500$ samplings.
   (a) Select a random sample of the minimum number of six feature correspondences to estimate the trifocal tensor T. This provides 1 or 3 solutions.
   (b) For each of these solutions:
      i. Calculate the error e_i^p for the ith point correspondence.
      ii. Calculate the error e_j^l for the jth line correspondence.
      iii. Calculate the total number of inliers over the two sets (i.e errors below a user specified threshold).
   (c) Select the best solution over all the samples i.e. that with the highest number of inliers.
2. Re-estimate the parameters using all the data that has been identified as consistent, using the non-linear method described in Section 4.

Table 1: *A brief summary of random sampling algorithm*

constraints are not imposed on the trifocal tensor elements, the various linear methods of transfer in equation (1) will give different results for the transferred point position. The two central results of this paper are:

1. An algorithm is given for computing the trifocal tensor from 6 point correspondences over three views. This is the minimum number of correspondences possible for points in general position. The algorithm results in the solution of a cubic and there are one or three real solutions. This algorithm is given in appendix A.

2. A parameterisation of the trifocal tensor is given (18 DOF) which guarantees that the tensor elements satisfy all the constraints. This parameterisation is used to minimise a least squares estimate of the tensor. This parameterisation is discussed in section 4.

3 Robust Computation of the tri-focal tensor

In Torr [14] a comprehensive survey of robust estimators is reported, with comparisons of the results made on large scale synthetic tests as well as on real imagery. It was found that random sampling techniques provide a good first guess at a solution. This solution can be further improved by iterative least squares. In Torr *et al* [16] such a scheme was described in detail for the trifocal constraint using seven point sampling. Here just the improvements are presented using six point sampling and robust non-linear minimization. The algorithm is based on RANSAC and is summarized in Table 3.

For random sampling methods, such as RANSAC it is extremely important that the minimum number of correspondences are used, so as to reduce the probability of a mismatch being included in the random sample of correspondences. This is why the novel six point solution used here is important. It is not a problem that there are three solutions, since the correct solution of the three is identified when measuring support for the solution from the full set of putative matches. To determine support two error criteria are defined: one each for lines and points. Given the trifocal constraint and the location of a scene point in two images its location in the third image may be predicted. From empirical tests the best results have been obtained by the minimization of the distance of a predicted point from its actual location in the image plane. The error criterion that is used is the average of this distance over the three images. A trio of feature correspondences is deemed consistent with a given constraint if the error criterion is below a threshold of 1.96σ, where σ is the estimated standard deviation of the error in locating a point. For lines, the root mean square of the distances of the end points of the actual line to the predicted line is used as the error criterion.

The method for finding the trifocal tensor from six points uses the theory of Quan [11] for computing an invariant of six points from 3 views, and is described in Appendix A.

4 Parameterization of trifocal tensor

Once outliers have been eliminated using RANSAC a non-linear iterative method is used to further minimize the error. Such minimizations have been plagued in the past by a lack of a good parameterization of the trifocal tensor. In this case the six points associated with the best solution provided by the RANSAC method described in the previous section provide a very natural minimal parameterization. By fixing $x, y, x^{'}$ for each of these correspondences and varying $y^{'}, x^{''}, y^{''}$ the elements of the trilinear tensor may be parameterized in terms of the 6 correspondences and a standard gradient descent algorithm [6] conducted over $y^{'}, x^{''}, y^{''}$.

Note, this parameterisation has exactly 18 DOF (3 variables for each of the 6 correspondences) and the solution for the trifocal tensor is guaranteed to be consistent, i.e. the elements of the matrix satisfy the necessary constraints. Indeed the 6 point correspondences for the parameterisation need not be real point correspondences at all, but can be *virtual* points.

The error minimized is

$$D = \sum_i \gamma(e_i^p) + \sum_l \gamma(e_j^l) \tag{3}$$

where e_p and e_l are point and line errors and $\gamma(d)$ is a robust Huber function [8]:

$$\gamma(d) = \begin{cases} d & d < 1.96\sigma \\ 1.96\sigma & d \geq 1.96\sigma \end{cases} \tag{4}$$

σ being the standard deviation of the error. The Huber cost function (4) allows the minimization to be conducted on all features whether they are outliers or inliers. Use of the Huber function limits the effects of outliers on the minimization. (If it is not used the results significantly degrade.) Typically as the minimization progresses many outliers are redesignated inliers.

Given six correspondences one or three trifocal tensors are obtained, the solution that minimizes (3) is always selected, and in practise the solutions have been found to be very stable, i.e small perturbations of the minimal point set do not cause a switch to a radically different solution.

The non-linear minimization and RANSAC take about the same amount of computation time.

4.1 Other Parameterisations

Existing parameterisations which enforce the constraints between the elements generally proceed by initially computing the three 3×4 projection matrices, and computing the tensor from these [9]. This over parameterises the tensor. However, the main detrimental effect of this over parameterisation is likely to be the cost of the numerical solution. In the following section the 6 point parameterisation is compared to a 24 DOF parameterisation based on $P = [I|\ \mathbf{0}]$ for the first projection matrix, and two 3×4 projection matrices for the other images, each with 12 elements; and a 27 DOF parameterisation.

5 Results

We have rigorously tested the various methods presented on real and synthetic data. First experiments were made on synthetic data randomly generated in three space; 100 sets of 100 corresponding triples were generated. The image data was perturbed by Gaussian noise, standard deviation 1.0, and then quantized to the nearest 0.1 pixel. We then introduced mismatched features to make a given percentage of the total, between 10 and 50 percent. With synthetic data the estimate can be compared with the ground truth as follows: The standard deviation of the distance of the *actual* noise free projections of the synthetic world points to their transferred points (using two images to predict the third) is measured. This gives a good measure of the validity of each method in terms of the ground truth.

The first test compared the 7 point and 6 point methods using RANSAC to generate the trilinear tensor. The 7 point method finds the solution as an eigenvector of a 27×27 matrix [16]. It was found that the 6 point method performed better than the 7 point giving a standard deviation around 40% lower.

This is for two reasons, the first being that that the six point algorithm requires fewer correspondences to estimate and so has less chance of including an outlier; the second and perhaps more important is that the six point algorithm exactly encodes the constraints on the parameters of the trifocal constraint. The seven point algorithm on the other hand has too many degrees of freedoms, 27 when there should only be 18.

The six point algorithm is also considerably faster, in the case of the seven point algorithm the eigenvector of a 27×27 matrix must be found, which is slower than the solution of the cubic described in Appendix A. Furthermore far fewer six point samples need to be taken to get a given degree of confidence in the result, given a certain proportion of outliers.

The second test compared the 18 degree of freedom (DOF) parameterization described with two other suggested parameterizations which could be used when minimizing the trifocal tensor. One is a 27 DOF using all the coefficients of the tensor. Another is a 24 DOF of the tensor in terms of the camera matrices.

The 18 DOF parameterization gave the best standard deviation with standard deviation 0.3789, the 24 DOF had standard deviation 0.5271 and the 27 DOF standard deviation 0.789. Thus the minimal parameterization is clearly most accurate. Next two typical real examples: Figure 1 demonstrates the improvement effected by the non-linear routine; Figures 2 shows three consecutive views of a star ship as it looms towards the viewer. The point matches are shown over the three images, as are the line matches, and the inlying and outlying point and line matches. The line matches across the three images are all displayed on the third image.

6 Conclusions and Future Work

The point parameterisation works so well because the 6 points initially selected by RANSAC are known to provide a good estimate of the trifocal tensor (because there is a lot of support for this solution). RANSAC filters out "poor" 6 point bases — for example those arising from 6 nearby image points, or 6 points whose pre-images are poorly conditioned (for example near coplanar). However, it is still worth investigating whether improvements can be obtained by "freezing" other combinations of point coordinates. For example, instead of fixing $x, y, x^{'}$ and varying $y^{'}, x^{''}, y^{''}$, fixing $x, x', y^{'}$ or other combinations differing over the points. The general method (of minimal parameterization in terms of basis points found from RANSAC) could be used for any other estimation problem in vision, for instance estimating the fundamental matrix, projectivities, camera matrices etc.

The error measures used in this work (distance to transferred point) are not ideal for several reasons: they are not Gaussian, do not have constant variance across the image and provide a sub optimal estimate.

Acknowledgements We thank Stephen Laveau for discussions, and ACTS Project Vanguard for financial support.

A Computation of the trifocal tensor from six point correspondences

In this appendix the method for finding the trilinear constraint from six points is determined, the method is inspired by that of Quan [11] the derivation follows that given in [18].

Six space points in general position, otherwise the trifocal tensor cannot be uniquely determined, can be assigned canonical projective coordinates as follows: $(1,0,0,0)^\top$, $(0,1,0,0)^\top$, $(0,0,1,0)^\top$, $(0,0,0,1)^\top$, $(1,1,1,1)^\top$ and $(X,Y,Z,W)^\top$ where X,Y,Z,W are unknown. Similarly, and without loss of generality the image coordinates of the first four points in each image are assigned to the projective basis in each image, i.e. the coordinates of the six image points are $(1,0,0)^\top$, $(0,1,0)^\top$, $(0,0,1)^\top$, $(1,1,1)^\top$, $(x_5^{(i)}, y_5^{(i)}, w_5^{(i)})^\top$, and $(x_6^{(i)}, y_6^{(i)}, w_6^{(i)})^\top$; and $\mathbf{B}^{(i)}$ is the 3×3 projectivity that takes the image coordinates into this canonical frame.

Once the canonical system is set up,

$$\begin{bmatrix} 1 & 0 & 0 & 1 & x_5^{(i)} & x_6^{(i)} \\ 0 & 1 & 0 & 1 & y_5^{(i)} & y_6^{(i)} \\ 0 & 0 & 1 & 1 & w_5^{(i)} & w_6^{(i)} \end{bmatrix} = \begin{bmatrix} \alpha^{(i)} & 0 & 0 & \delta^{(i)} \\ 0 & \beta^{(i)} & 0 & \delta^{(i)} \\ 0 & 0 & \gamma^{(i)} & \delta^{(i)} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & X \\ 0 & 1 & 0 & 0 & 1 & Y \\ 0 & 0 & 1 & 0 & 1 & Z \\ 0 & 0 & 0 & 1 & 1 & W \end{bmatrix} \quad (5)$$

Figure 1: **Top 3 images** *Three images from a sequence of a chapel, acquired by a hand-held camcorder. Camera motion is lateral and a few centimetres between frames. The image size is* 760×550 *pixels. There are 341 corners tracked over all three images (from about 500 tracked pairwise).* **Middle 3 images** *Matches are shown for all three images, with the end points joined together and superimposed on the third image. The first shows all the matches, the middle shows the inliers after RANSAC and the right outliers. There are 146 inliers. After RANSAC the r.m.s. error between transferred points (using the trifocal tensor) and actual points in the third image is 2.5 pixels.* **Bottom 3 images** *Showing the results after the non-linear minimization. The left shows the six point basis selected by RANSAC. The middle shows the inliers after the non linear scheme, the right outliers. After the non-linear scheme the number of inliers has increased to 184. The r.m.s. error has decreased to 1.6 pixels. Note all the outliers on the tree to the left have been removed.*

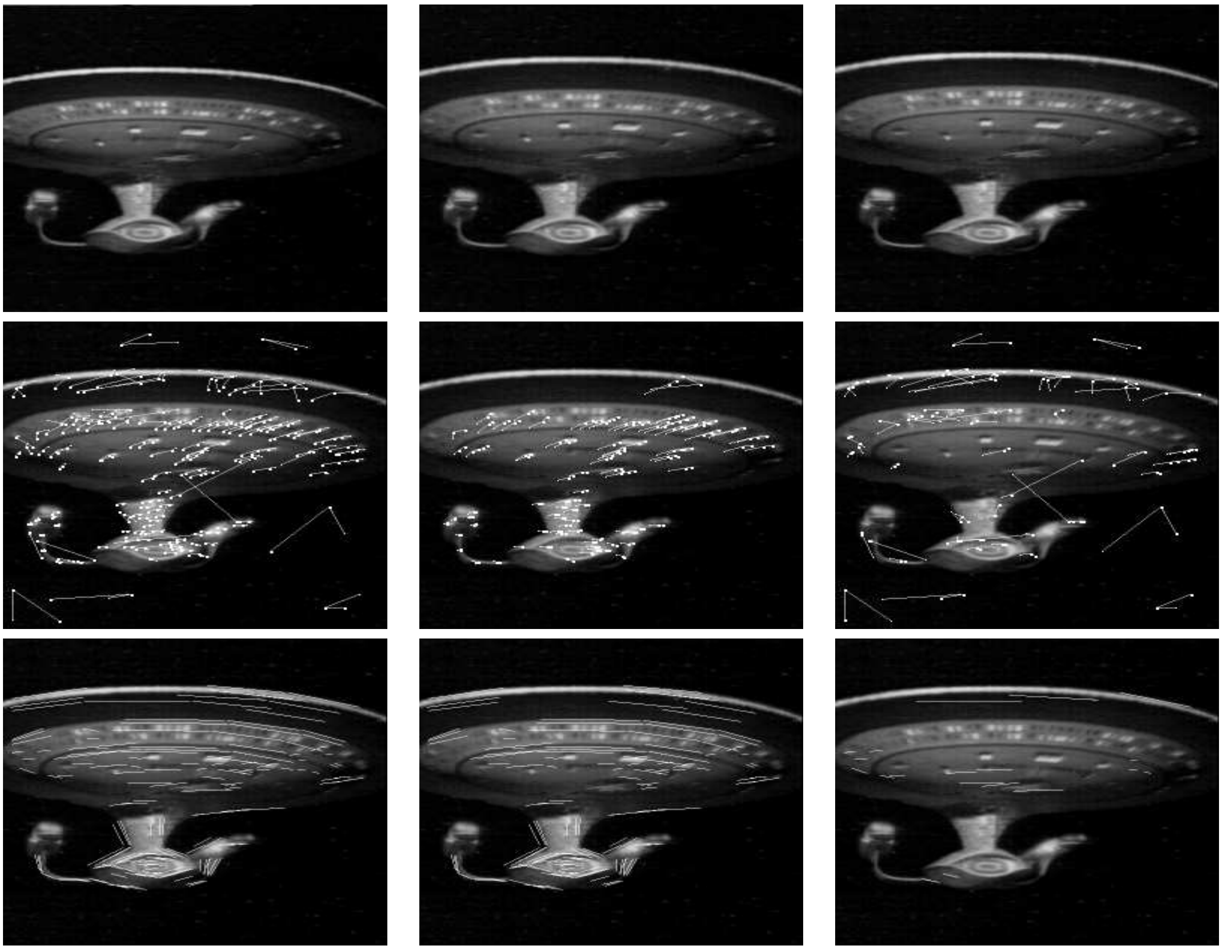

Figure 2: **Top 3 images** *three consecutive views of a starship looming.* **Middle 3 images** *Left: the point correspondences over the three images, shown superimposed onto the last image. Small squares indicate the previous positions. Middle: inliers. Right outliers.* **Bottom 3 images** *Left line matches from all three images displayed onto the last image, line correspondence is based on proximity and relative orientation. Middle: inliers. Right outliers. The efficacy of using the trifocal constraint to guide line matching can be evaluated by examining the inliers and outliers. The line matcher works well and there are few outliers.*

for each image $i = 1, 2, 3$. Thus recovery of the trifocal tensor is equivalent to recovering the coordinates of $(X, Y, Z, W)^\top$ and $(\alpha^{(i)}, \beta^{(i)}, \gamma^{(i)}, \delta^{(i)})$ for each camera. From (5) the values of the sixth space point and camera parameters may be obtained in terms of the fifth and sixth image coordinate as follows:

$$\begin{bmatrix} (-x_5^{(1)} y_6^{(1)} + x_5^{(1)} w_6^{(1)}) & (x_6^{(1)} y_5^{(1)} - y_5^{(1)} w_6^{(1)}) & (-x_6^{(1)} w_5^{(1)} + y_6^{(1)} w_5^{(1)}) \\ (-x_5^{(2)} y_6^{(2)} + x_5^{(2)} w_6^{(2)}) & (x_6^{(2)} y_5^{(2)} - y_5^{(2)} w_6^{(2)}) & (-x_6^{(2)} w_5^{(2)} + y_6^{(2)} w_5^{(2)}) \\ (-x_5^{(3)} y_6^{(3)} + x_5^{(3)} w_6^{(3)}) & (x_6^{(3)} y_5^{(3)} - y_5^{(3)} w_6^{(3)}) & (-x_6^{(3)} w_5^{(3)} + y_6^{(3)} w_5^{(3)}) \end{bmatrix} \tag{6}$$

$$\begin{bmatrix} (-x_5^{(1)} w_6^{(1)} + y_5^{(1)} w_6^{(1)}) & (x_5^{(1)} y_6^{(1)} - y_6^{(1)} w_5^{(1)}) \\ (-x_5^{(2)} w_6^{(2)} + y_5^{(2)} w_6^{(2)}) & (x_5^{(2)} y_6^{(2)} - y_6^{(2)} w_5^{(2)}) \\ (-x_5^{(3)} w_6^{(3)} + y_5^{(3)} w_6^{(3)}) & (x_5^{(3)} y_6^{(3)} - y_6^{(3)} w_5^{(3)}) \end{bmatrix} \mathbf{t} = 0 \tag{7}$$

$$\tag{8}$$

therefore the vector $\mathbf{t} = (WX - YZ, WY - YZ, WZ - YZ, XY - YZ, XZ - YZ)$ lies in the null space of the matrix on the left, which, if the points are in general position has rank 3. The two dimensional null space of the matrix in (8) may be recovered by a singular value decomposition, let $\mathbf{t}_1$ and $\mathbf{t}_2$ be the two vectors spanning this null space. Then up to a scale factor $\mathbf{t} = \mathbf{t}_1 + \alpha \mathbf{t}_2$. There is a cubic constraint on the elements of $\mathbf{t}$, let $\mathbf{t} = (t_1, t_2, t_3, t_4, t_5)$ then $t_1 t_2 t_5 - t_2 t_3 t_5 - t_2 t_4 t_5 = t_1 t_3 t_4 - t_2 t_3 t_4 - t_3 t_4 t_5$ which can be proven by substituting the X, Y, Z, W values of the t_i. Imposing this cubic constraint on $\mathbf{t} = \mathbf{t}_1 + \alpha \mathbf{t}_2$ leads to a cubic equation in α with one or three real solutions for $\mathbf{t}$. Given $\mathbf{t}$ then $(X, Y, Z, W)^\top$ may be recovered as follows:

$$\frac{X}{W} = \frac{t_4 - t_5}{t_2 - t_3} \quad \frac{Y}{W} = \frac{t_4}{t_1 - t_3} \quad \frac{Z}{W} = \frac{t_5}{t_1 - t_2}$$

assuming that $W \neq 0$, if $W = 0$ then the sixth point is on the plane at infinity it is trivial to use an alternate set of equations to recover $(X, Y, Z, W)^\top$. Given $(X, Y, Z, W)^\top$ the parameters of the camera matrices $(\alpha^{(i)}, \beta^{(i)}, \gamma^{(i)}, \delta^{(i)})$ may be recovered in a linear manner from (5).

From the camera matrices the structure may be initialised directly in the original coordinate system the camera matrices are:

$$\mathbf{P}^{(i)} = \mathbf{B}^{-1} \begin{bmatrix} \alpha^{(i)} & 0 & 0 & \delta^{(i)} \\ 0 & \beta^{(i)} & 0 & \delta^{(i)} \\ 0 & 0 & \gamma^{(i)} & \delta^{(i)} \end{bmatrix}. \tag{9}$$

To recover the trifocal tensor the first camera is set to $[\mathbf{I}|\mathbf{0}]$ (effected by a simple transformation of the coordinates) then the trifocal tensor's coefficients are given by (2).

References

[1] M. Armstrong, A. Zisserman, and R. Hartley. Self-calibration from image triplets. In B. Buxton and Cipolla R., editors, *Proc. 4th European Conference on Computer Vision, LNCS 1064, Cambridge*, pages 3–16. Springer–Verlag, 1996.

[2] P. Beardsley, P. H. S. Torr, and A. Zisserman. 3d model aquisition from extended image sequences. In B. Buxton and Cipolla R., editors, *Proc. 4th European Conference on Computer Vision, LNCS 1065, Cambridge*, pages 683–695. Springer–Verlag, 1996.

[3] R. Deriche, Z. Zhang, Q. T. Luong, and O. Faugeras. Robust recovery of the epipolar geometry for an uncalibrated stereo rig. In J. O. Eckland, editor, *Proc. 3rd European Conference on Computer Vision, LNCS 800/801, Stockholm*, pages 567–576. Springer-Verlag, 1994.

[4] O.D. Faugeras and B. Mourrain. On the geometry and algebra of the point and line correspondences between n images. In *Proc. 5th Int'l Conf. on Computer Vision, Boston*, pages 951–962, 1995.

[5] O.D. Faugeras and L. Robert. What can two images tell us about a third one. In J. O. Eckland, editor, *Proc. 3rd European Conference on Computer Vision, LNCS 800/801, Stockholm*, pages 485–492. Springer-Verlag, 1994.

[6] Numerical Algorithms Group. *NAG Fortran Library vol 7*. NAG, 1988.

[7] R. I. Hartley. Lines and points in three views – a unified approach. In *ARPA Image Understanding Workshop, Monterey*, 1994.

[8] P. J. Huber. *Robust Statistics*. John Willey and Sons, 1981.

[9] S. Laveau. *Geometry of a system of N cameras. Theory, estimation and applications.* PhD thesis, INRIA, 1996.

[10] Q. T. Luong, R. Deriche, O. D. Faugeras, and T. Papadopoulo. On determining the fundamental matrix: analysis of different methods and experimental results. Technical Report 1894, INRIA (Sophia Antipolis), 1993.

[11] L. Quan. Invariants of 6 points from 3 uncalibrated images. In J. O. Eckland, editor, *Proc. 3rd European Conference on Computer Vision, LNCS 800/801, Stockholm*, pages 459–469. Springer-Verlag, 1994.

[12] A. Shashua. Trilinearity in visual recognition by alignment. In *Proc. 3rd European Conference on Computer Vision, LNCS 800/801, Stockholm*, volume 1, pages 479–484, May 1994.

[13] M. Spetsakis and J. Aloimonos. Structure from motion using line correspondences. *International Journal of Computer Vision*, 4(3):171–183, 1990.

[14] P. H. S. Torr. *Outlier Detection and Motion Segmentation*. PhD thesis, University of Oxford, Engineering Dept., 1995.

[15] P. H. S. Torr, P. A. Beardsley, and D. W. Murray. Robust vision. In J. Illingworth, editor, *Proc. 5th British Machine Vision Conference, York*, pages 145–155. BMVA Press, 1994.

[16] P. H. S. Torr, A. Zisserman, and D. W. Murray. Motion clustering using the trilinear constraint over three views. In R. Mohr and C. Wu, editors, *Europe-China Workshop on Geometrical Modelling and Invariants for Computer Vision*, pages 118–125. Springer–Verlag, 1995.

[17] B. Triggs. The geometry of projective reconstruction i: Matching constraints and the joint image. In *Proc. 5th Int'l Conf. on Computer Vision, Boston*, pages 338–343, 1995.

[18] D. Weinshall, M. Werman, and A. Shashua. Duality of multi-point and multi-frame geometry: Fundamental shape matrices and tensors. In B. Buxton and Cipolla R., editors, *Proc. 4th European Conference on Computer Vision, LNCS 1065, Cambridge*, pages 217–227. Springer–Verlag, 1996.

[19] J. Weng, N. Ahuja, and T. Huang. Optimal motion and structure estimation. *IEEE PAMI*, vol.15(9):864–884, 1993.

[20] A. Zisserman and Maybank S. A case against epipolar geometry. In J. Mundy, A. Zisserman, and D. Forsyth, editors, *Applications of Invariance in Computer Vision LNCS 825*. Springer-Verlag, 1994.