Reconstruction from two views using approximate calibration

Richard Hartley and Chanop Silpa-Anan,
Department of Systems Engineering,
RSISE,
Australian National University,
ACT, 0200, AUSTRALIA.

Email: {hartley,chanop}@syseng.anu.edu.au

Abstract

The problem of Euclidean reconstruction from two perspective images is well studied for calibrated cameras, and good algorithms are known. On the other hand, if the cameras are known to have square pixels (no skew and unit aspect ratio) then the problem is also theoretically solvable, provided an estimate of the principal point is provided. The focal lengths of the cameras may be computed from the fundamental matrix, however, it is quite sensitive to the computed fundamental matrix and the assumed position of the principal point. In fact, sometimes the estimate of the focal length fails, and so Euclidean reconstruction is impossible using this method. In this paper, we investigate the cause of this problem, and suggest an algorithm that more reliably leads to a reconstruction. Weak bounds on the principal point locations, and the focal lengths of the cameras and the condition that points must lie in front of the cameras give enough constraint to compute a fundamental matrix that always leads to a plausible focal length estimate, and hence Euclidean reconstruction that suffers only a very small degradation in residual point-reprojection error.

1. Introduction

Scene reconstruction from a number of perspective images of a scene is one of the fundamental problems of computer vision, and reconstruction from two views is the simplest case of this problem. Earliest work in this area concentrated on calibrated cameras, from which it is in principal possible to obtain a Euclidean (sometimes called metric) reconstruction of the scene. Such a reconstruction is unique apart from a choice of the Euclidean coordinate frame, and an overall scale, which it is impossible to determine. Notable algorithms for obtaining a reconstruction from two calibrated views include [LH81, Hor90, Hor91].

With a shift of interest towards uncalibrated cameras, it was natural to ask whether Euclidean reconstruction was possible from uncalibrated cameras, leading to a well-known negative result ([HGC92, Fau92]) showing that for arbitrary uncalibrated views, the best one can achieve is projective reconstruction. On the other hand, it was shown in [Har92] that given reasonable assumptions about the calibration (square pixels and known principal points), it is still possible to achieve Euclidean reconstruction. Under these partial-calibration conditions, the remaining camera calibration parameters (the focal lengths of the two cameras) may be computed. After this a calibrated reconstruction method may be used to obtain a Euclidean reconstruction.

The simplest implementation of this method involves a formula of Bougnoux ([Bou98]):

$$f_2^2 = -\frac{(\mathbf{p}_1{}^\top[\mathbf{e}_1]_\times \mathtt{I}\mathtt{F}^\top\mathbf{p}_2)(\mathbf{p}_1{}^\top\mathtt{F}^\top\mathbf{p}_2)}{\mathbf{p}_1{}^\top([\mathbf{e}_1]_\times \mathtt{I}\mathtt{F}^\top \mathtt{I}\mathtt{F})\mathbf{p}_1} \qquad (1)$$

where $\mathtt{I}$ is the diagonal matrix $\mathrm{diag}(1,1,0)$. A similar formula for f_1^2 is obtained by interchanging the roles of the two images.

This algorithm, however, suffers from various deficiencies.

1. Computation of the focal lengths is not possible if the principal axes of the two cameras intersect ([NHBP96]).
2. The algorithm requires the position of the principal point to be known. Furthermore, the computation of the focal lengths can be very sensitive to the assumed positions of the principal points, and also to the fundamental matrix.
3. In the worst case, the algorithm will fail completely, because the value of f^2 computed from (1) is negative, resulting in an imaginary value for the focal length. It is our observation that a fundamental matrix (even those computed by the best known methods) may in this sense be incompatible with *any* reasonable choice of the principal point.

Indeed, the sensitivity is so severe that computation of the focal lengths can not be relied upon at all, and this algorithm is of doubtful practical value.

In reconstruction from two views, the camera calibration is usually not completely unknown, but often not completely known either. The principal point is usually near the centre of the image, and the focal length is at least known within some reasonable bounds. Our goal is to show that Euclidean reconstruction is possible from two views, even without exact knowledge of the principal point, provided some assumptions are admitted concerning the position of the principal points, and the focal lengths of the images. As part of this reconstruction process, a fundamental matrix is found that is compatible with reasonable estimates of the principal point. Thus, incorporation of this *a priori* knowledge of the probable principal points and focal lengths leads to an improved estimate of the fundamental matrix.

The constraints imposed by the a priori knowledge of the principal points and focal lengths can be very weak, and yet lead to a considerably improved results. For instance, in the examples discussed below, an assumption that the principal point is weakly constrained to be near the image centre (let us say plus or minus half the image radius), and that the focal lengths of the two cameras are approximately equal (maybe within 5%), is sufficient for Euclidean reconstruction to succeed. Other assumptions are possible, as will be seen. Please refer to section 5 for details of the weights used in our experiments.

Oliensis [Oli00] has undertaken a study of reconstruction methods, and argues that sometimes calibrated camera algorithms give better results than using uncalibrated algorithms. This can only be true when the camera calibration is accurately known, however. The algorithm given in this paper gives a solution for the intermediate case where the camera calibration parameters are known only approximately – with a degree of uncertainty.

2. Impossible fundamental matrices

In this section, it will be shown that some fundamental matrices, even those computed from good quality data by the best algorithms (such as a Maximum Likelihood algorithm) are nevertheless incompatible with the data, and hence wrong. This is actually quite a common phenomenon, as the examples will show.

For a fundamental matrix to be accepted as being compatible with a set of matched points $\mathbf{x}_{2i} \leftrightarrow \mathbf{x}_{1i}$, three conditions are necessary.[1]

1. The point correspondences must satisfy the coplanarity condition ${\mathbf{x}_{2i}}^\top \mathtt{F} \mathbf{x}_{1i}$, with a small residual error.
2. For at least some assumed locations of the principal points $\mathbf{p}_1$ and $\mathbf{p}_2$, the values of f_1^2 and f_2^2 computed from (1) are positive.
3. Given such assumed principal point positions, and corresponding focal length values, a calibrated reconstruction is possible, for which the reconstructed 3D points (apart from a small percentage of possible outliers) lie in front of the reconstructed cameras.

This last condition is related to the concept of "cheirality" discussed in [Har98], where it is shown that satisfying the coplanarity condition for *some* fundamental matrix is not sufficient for the set of matches to be realizable. Points incorrectly placed behind the cameras in a Euclidean reconstruction severely degrade the quality of that reconstruction.

This discussion is illustrated by the examples given in Fig 1. There a fundamental matrix is used, computed using a bundle-adjustment method (the Gold-standard method of [HZ00]), which is as good as any known method. The set of matched points used for this computation are of high quality, outliers having been previously removed, and the residual error from estimation of F being very small (see table 3). Nevertheless, it is seen that the possible positions for the principal point are very constrained. In fact, the most likely location of the principal point (the centre of the image) is not compatible with the computed fundamental matrix. This problem occurs with most of the other algorithms tried (see table 4 and table 5).

3. A cost function

Given a set of point correspondences $\mathbf{x}_{2i} \leftrightarrow \mathbf{x}_{1i}$, our object is to estimate the fundamental matrix subject to prior assumptions about the distribution of the focal lengths and principal points of the two cameras. The method is applicable easily to any assumed distributions, but most commonly a normal (Gaussian) distribution will be assumed. An iterative (Levenberg-Marquardt) method is used to minimize a cost function of the following form:

$$\text{Cost}(\mathtt{F}, f_1^2, f_2^2, \mathbf{p}_1, \mathbf{p}_2) = C_{\mathtt{F}}(\mathtt{F}) + C_f(f_1^2, f_2^2) + C_p(\mathbf{p}_1, \mathbf{p}_2) \quad . \tag{2}$$

Thus the total cost is the sum of three cost functions, measuring the cost of estimates of the fundamental matrix F, the squared focal lengths f_1^2 and f_2^2, and the principal points respectively. The reason for expressing the second cost function in terms of the *squared* focal lengths is because of the form of Bougnoux's formula (1). We want to define a cost (very high) for a negative value of f^2.

The cost function $C_{\mathtt{F}}(\mathtt{F})$ is also a function of the point correspondences. Although various cost functions are possible, we prefer to use the Sampson cost function

$$\sum_i \frac{({\mathbf{x}_{2i}}^\top \mathtt{F} \mathbf{x}_{1i})^2}{(\mathtt{F}\mathbf{x}_{1i})_1^2 + (\mathtt{F}\mathbf{x}_{1i})_2^2 + (\mathtt{F}^\top \mathbf{x}_{2i})_1^2 + (\mathtt{F}^\top \mathbf{x}_{2i})_2^2}. \tag{3}$$

This cost function is a first-order approximation to the correct geometric (or Maximum Likelihood) cost function.

[1] Recall that we are making an assumption of square pixels (that is zero skew and unit aspect ratio) for both cameras.

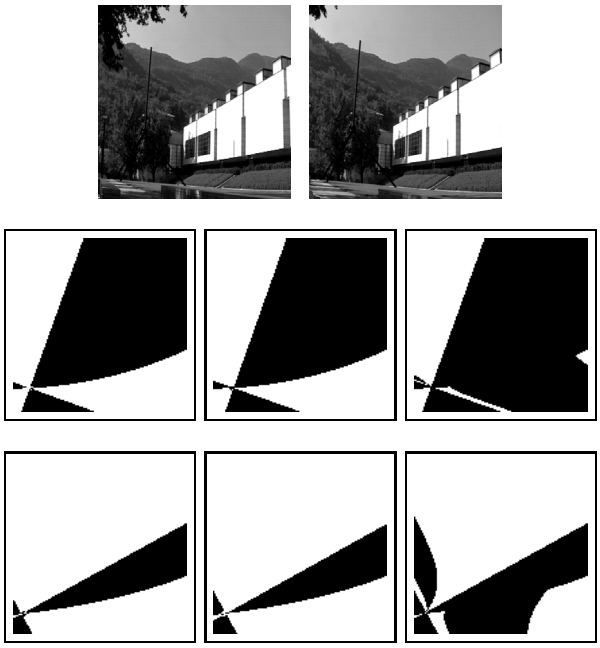

Figure 1: *At the top are a pair of images used for computing a fundamental matrix. The fundamental matrix was computed using the Gold-standard algorithm of [HZ00]. For this example it is assumed that the principal point is the same in both images. In the middle row of figures, the possible positions for the principal point are shown. White shows possible positions of the principal point, and black impossible. The criteria used for the three diagrams are: f_1^2 positive (left), f_2^2 positive (centre) and $< 10\%$ points behind cameras. The graph on the right shows those positions for the principal points where all three conditions are satisfied. As seen, there are very few possible positions for the principal points consistent with this fundamental matrix – certainly not the centre of the image.*

The third row of figures shows the same results for the fundamental matrix computed using the method described in this paper. The obtained fundamental matrix is consistent with the assumption of principal point near the centre of the image.

The total cost function is minimized over a set of parameters by Levenberg-Marquardt algorithm. The set of parameters is divided into two parts.

1. A set of parameters parametrizing the fundamental matrix F, and
2. A set of 4 (or 2) parameters defining the positions of the principal points. Two parameters may be used if it is assumed (and enforced) that the principal point is the same in both images.

The minimization algorithm is as follows.

1. Given F and $\mathbf{p}_i$, compute f_1^2 and f_2^2 using Bougnoux's formula (1).
2. Compute the cost function Cost$(\mathtt{F}, f_i^2, \mathbf{p}_i^2)$.
3. Vary the parameters of F and the $\mathbf{p}_i$ to minimize the cost function.

Form of $C_{\mathbf{p}}$ and C_f. The specific form of the cost functions for $\mathbf{p}_i$ and f_i^2 may be chosen in various ways. However, a natural form for $C_{\mathbf{p}}$ is the squared Euclidean distance of the estimated principal point from the nominal value. Thus

$$C_{\mathbf{p}}(\mathbf{p}_i) = w_p^2 d(\mathbf{p}_i, \bar{\mathbf{p}}_i)^2 \tag{4}$$

where $\bar{\mathbf{p}}_i$ represents the nominal position of the principal point, and w_p is a weight. There is a cost term of this form for each of the two principal points.

The form of the cost function C_f may be a little more complicated, since we want to ensure that the value of f^2 does not end up negative. In addition, it may be appropriate to enforce a condition that the two focal lengths are the same (or approximately so). Accordingly, the cost function $C_f(f_1^2, f_2^2)$ used in our implementation of the algorithm has several components:

$$\begin{aligned} C_f(f_1^2, f_2^2) \quad = \quad & w_1^2(f_1^2 - \bar{f}_1^2)^2 + w_2^2(f_2^2 - \bar{f}_2^2)^2 \\ & + w_d^2(f_1^2 - f_2^2) \\ & + w_{z1}^2(f_{\min}^2 - f_1^2)^2 \\ & + w_{z2}^2(f_{\min}^2 - f_2^2)^2 \end{aligned} \tag{5}$$

Recall here that f_i^2 is the value returned by Bougnoux's formula (1), and may be negative. The final two terms of this equation involve a "minimum" value $f_{\min}$ for the focal length, and are only included if $f_i^2 < f_{\min}^2$. (That is w_{z1} and w_{z2} are zero unless f_i^2 is small, or negative.). This term grows rapidly for small or negative values of f_i^2, and effectively prevent f_i^2 taking on negative values. A reasonable minimum value of $f_{\min}$ can be deduced from the size of the image. The field of view of the camera is equal to $2\arctan(\mathrm{dim}/f)$, where dim is the radius of the image. For small values of f, this becomes unrealistically large. Most images encountered (except for extreme wide angled views) do not have field of view exceeding 75°.

The choice of the weight values may be chosen according to taste. The values of w_{zi} are not critical, and normally it is sufficient to apply quite weak weights for the other values.

4. Initialization

The input data for the reconstruction problem includes an estimate of the focal length and principal point of the cameras. Thus, given a set of point correspondences, $\mathbf{x}_{2i} \leftrightarrow \mathbf{x}_{1i}$ and initial estimates $\bar{\mathbf{p}}_i$ and $\bar{f}_i$ of the principal points and focal

lengths, an initial value for the fundamental matrix is found as follows. Let $\mathtt{K}_1$ and $\mathtt{K}_2$ be the initial calibration matrices for the two images. Thus,

$$\mathtt{K}_1 = \begin{bmatrix} \bar{f}_1 & 0 & \bar{x}_1 \\ & \bar{f}_1 & \bar{y}_1 \\ & & 1 \end{bmatrix} \qquad \mathtt{K}_2 = \begin{bmatrix} \bar{f}_2 & 0 & \bar{x}_2 \\ & \bar{f}_2 & \bar{y}_2 \\ & & 1 \end{bmatrix} \tag{6}$$

where $\bar{\mathbf{p}}_i = (\bar{x}_i, \bar{y}_i)$ are the principal points. Let $\mathtt{F}$ be an estimate of the fundamental matrix computed from the point correspondences using any desired method. We used a method ([HZ00]) that minimizes algebraic error. The essential matrix $\mathtt{E}$ may then be computed as

$$\mathtt{E} = {\mathtt{K}_2}^\top \mathtt{F} \mathtt{K}_1 \ .$$

This value of the essential matrix will not generally be quite correct for the assumed calibration matrices, since it will not satisfy the necessary condition for a essential matrix, namely that it have two equal singular values. To correct this, the singular value decomposition of $\mathtt{E}$ is computed as $\mathtt{E} = \mathtt{U}\mathtt{D}\mathtt{V}^\top$, and a corrected essential matrix is computed, by setting $\widehat{\mathtt{E}} = \mathtt{U}\mathtt{I}\mathtt{V}^\top$, where $\mathtt{I} = \mathrm{diag}(1, 1, 0)$. Finally a corrected value of the fundamental matrix is computed by setting

$$\widehat{\mathtt{F}} = {\mathtt{K}_2}^{-\top} \widehat{\mathtt{E}} \mathtt{K}_1^{-1} \ .$$

The resulting fundamental matrix $\widehat{\mathtt{F}}$ is the best approximation to $\mathtt{F}$, compatible with the assumed calibration matrices.

Iteration will start with this estimate $\widehat{\mathtt{F}}$ for the fundamental matrix, and the assumed values $\bar{\mathbf{p}}_1$ and $\bar{\mathbf{p}}_2$ for the two principal points. The cost function $C_{\mathtt{F}}(\widehat{\mathtt{F}})$ will be slightly greater than $C_{\mathtt{F}}(\mathtt{F})$, but usually the difference is not too great. However, because of the way $\widehat{\mathtt{F}}$ is defined, the cost functions $C_{\mathbf{p}}(\bar{\mathbf{p}}_1, \bar{\mathbf{p}}_2)$ and $C_f(f_1^2, f_2^2)$ will both have initial values zero.

5. Experiments

The algorithm was carried out on both synthetic and real data. The algorithms used for computing the fundamental matrix were as follows. For details of these algorithms, the reader is referred to [HZ00].

1. A normalized 8-point algorithm.
2. The gold-standard algorithm (bundle adjustment).
3. Iterative minimization of algebraic error (algebraic minimization algorithm).
4. Iterative minimization of Sampson error, (3).
5. The algorithm of this paper (denoted by *a-priori* in the graphs and tables).
6. Calibrated reconstruction, given fixed assumed values of the principal points and focal length.

The final algorithm is the one described in section 4, which is used as an initial point for the *a-priori* algorithm.

Noise Level	Cube		Shed	
	gold-standard	a-priori	gold-standard	a-priori
0.0	43.82	4.40	20.01	42.77
0.2	44.18	6.31	51.76	43.82
0.5	44.81	15.27	102.9	82.06
1.0	46.46	47.56	148.6	90.86
2.0	60.73	66.00	177.9	125.4
4.0	117.3	123.0	203.1	186.8

Table 1: **Mean errors for synthetic data.** *This table shows the average reconstruction error for the two synthetic data set, cube and shed images. For this example, two algorithms are used to estimated the fundamental matrix, the gold-standard algorithm and the the* a-priori *algorithm. An incorrect initial principal point (225.5,225.5) and focal length (590) are given to the algorithms.*

5.1. Synthetic image experiments

Evaluation of the method was carried out on two sets of data. In the first set, scattered points on the surface of a cube were used, and data were generated by projecting the points into a pair of nominal cameras. In the second set, a pair of images were used to build a realistic-looking model of a house (using the images in [HZ00] (page 250)) The model was then projected back into two images, approximating the original images. This was the synthetic data, corresponding to a realistic imaging setup, that was used for experiments. For each synthetic data set, noise was added in varying degrees, and Euclidean reconstruction was carried out.

The experiments were carried out with two different a priori estimates for the principal point and focal length. In the first set, the exact values were given. This of course gives a significant advantage to the calibrated reconstruction algorithm, since it is provided with exact values for these parameters. In the second set of experiments, the assured value of the principal point is shifted by 30 pixels in each dimension (in a 512×512 image), and a slightly modified value for the focal length is given. This is perhaps more realistic, since these parameter values may not be known exactly in advance.

Choice of weights. The *a-priori* algorithm was run with very weak weights, in order to place minimal constraints on principal point and focal length. The value of w_p was equal to 0.01, which means that a variation of 100 pixels (in a 512×512 image) was given as much weight as one pixel reprojection error for a single matched point. The principal point was, however, assumed to be the same in both images.

As for the focal length, the weights w_1 and w_2 were set to zero, which means that no constraint was placed on these parameters individually. However, a value of 0.001 was chosen for w_d. This means that a value of 1000 for $f_1^2 - f_2^2$ was equivalent to one pixel reprojection error. Since f_i was

around 500, this means approximately one pixel difference in $f_1 - f_2$.

As stated, the values of w_{zi} are not very critical. A value of $f_{\min} = 100$ was taken, and a weight $w_{zi} = 0.01$ was chosen.

Findings. Perhaps because of the controlled nature of the synthetic algorithms, in no case did the values of f_i^2 turn out to be negative. As a result, the reconstruction results were good for all the algorithms used. The *a-priori* algorithm performed significantly better than the calibrated reconstruction algorithm, except in the case where exact parameters were given, and for high noise values. See table 1 for a comparison of the *a-priori* and the gold-standard algorithms.

5.2. Real image experiments

Experiments were also carried out on real images, for which, however, no ground truth structure was known. The image pairs used were the ones used in [HZ00] (page 273) for evaluation of fundamental matrix computation. They may be taken as reasonably representative of real image sets. The experiments were carried out with the same weight values and a priori parameter estimates as the synthetic images. The results are given in the following tables 2 – 3.

Findings. Most of the algorithms used gave values for the fundamental matrices for which the focal length values f_i^2 were negative, under the reasonable assumption that the principal points were in the centre of the images. Thus, in this case it was not possible to proceed with Euclidean reconstruction. For each image, a calibrated reconstruction algorithm was also used to compute the fundamental matrix. Since the calibration was only guessed at, understandably the residual reprojection error in this case was much worse than for the present *a-priori* method. The consequence of this will be a significantly degraded reconstruction, since the 3D points will not be estimated accurately. In addition, in one case (the statue example, see table 5), the calibrated reconstruction resulted in points being found behind the camera. This also occured for some other algorithms and image sets, but not for the *a-priori* algorithm. The consequence of points ending up behind the camera will be a severely distorted Euclidean reconstruction.

It is important to note that the residual projection error is only very slightly greater for the *a-priori* algorithm than it is for the gold-standard algorithm (or the other algorithms). Despite this, the *a-priori* algorithm gives fundamental matrices that are much more realistic, in fact usable for subsequent 3D reconstruction.

Method / Image	a-priori
Basement	(255.5,255.5)
Calib	(255.9,254.6)
House	(254.2,234.9)
Museum	(232.0,229.8)
Statue	(248.5,206.9)

Table 2: **Estimated principal points.** *For the* a-priori *algorithm of this paper, the principal points were estimated to be at location (255.5,255.5). The algorithm has moved the points away from the initial estimate to the locations given in this table. As noted, the required changes in principal point are quite small. This change comes at a very small additional cost in point residual (see table 3). However, it results in possible (and plausible) values for the cameras' focal lengths (see table 4).*

Method / Image	gold standard	a-priori	calibrated
Basement	6.95×10^{-2}	9.82×10^{-2}	9.74×10^{-2}
Calibration	1.32×10^{-2}	1.43×10^{-2}	1.03×10^{-1}
House	1.83×10^{-1}	1.83×10^{-1}	3.47×10^{-1}
Museum	2.46×10^{-1}	2.9×10^{-1}	3.51×10^{-1}
Statue	2.77×10^{-1}	2.81×10^{-1}	1.26
Cube	1.26×10^{-6}	1.26×10^{-6}	3.05
Shed	1.31×10^{-6}	1.21×10^{-6}	3.19

Table 3: **Residuals.** *This table summarizes the residual RMS error of the estimated fundamental matrix obtained from 3 algorithms: the gold-standard algorithm, the* a-priori *algorithm of this paper, and calibrated fundamental matrix estimation using nominal values of of the focal length and principal point. Note that the extra cost (in terms of residual error) of using the* a-priori *algorithm, compared with the gold-standard algorithm is very small. However, it has the great advantage (see Tables 4 and 5) that it leads to a fundamental matrix compatible with cheirality and reasonable estimates of principal point and focal length. The* calibrated *algorithm gives larger residuals, presumably due to incorrect assumed values of focal length and principal point.*

6. Conclusion

This paper points out some of the difficulties involved in Euclidean reconstruction from two views. It is shown that most of the standard algorithms for computing the fundamental matrix may result in matrices that are wrong, and in fact unusable for 3D reconstruction. However, adding some weak constraint terms to the cost function used to compute the fundamental matrix can lead to vastly improved results, particularly as far as estimation of the focal lengths, and subsequent 3D reconstruction is concerned. The effect is far more noticeable for real images than synthetic ones. The cost of adding these constraint terms is very small in terms of increased reprojection residual.

Method / Image	normalized 8-points	gold standard	algebraic distance	Sampson error	a-priori
Basement	X/X	X/X	X/X	X/X	503.3/503.4
Calibration	872.3/864.3	510.7/509.2	X/X	X/X	279.9/279.7
House	X/X	X/X	X/X	X/X	451.6/451.5
Museum	X/X	X/X	X/X	2964/3001	654.8/654.8
Statue	X/X	X/X	X/X	X/X	1167/1167

Table 4: **Estimated focal lengths.** *This table shows the estimated focal length for each image pair calculated from 5 algorithms. Note that some well known algorithms such as normalized 8-points algorithm and gold standard algorithm may give an impossible answer (imaginary focal length, denoted by X).*

Method / Image	normalized 8-points	gold standard	algebraic distance	Sampson error	a-priori	calibrated
Basement	0/0	0/0	0/100	0/0	100/100	100/100
Calibration	100/100	100/100	100/0	100/100	100/100	100/100
House	91.7/91.7	100/0	100/0	100/0	100/100	100/100
Museum	0/0	0/100	0/100	94.9/94.9	100/100	100/100
Statue	100/100	100/0	100/0	100/100	100/100	52.4/52.4

Table 5: **Percentage points in front of cameras.** *This table shows the percentage of points that lie in front of each camera in each image pair. It is clear that the algorithm in this paper gives a better result in that the estimated fundamental matrix yields estimated points in front of the camera. When wrong principal points and focal length are used in the calibrated reconstruction algorithm some points may end up behind the cameras, as seen for the statue pair.*

In our experiments, we did not attempt to add any extra cost terms to ensure that the reprojected points lie in front of the cameras. In all the cases shown here, the new algorithm resulted in points being reconstructed in front of the cameras. This is not guaranteed, however, and we occasionally found cases in which this condition was violated.

Further work in this area can include determination of the best weight functions for principal point and focal-length errors, and an investigation of adding terms to the cost function to ensure points lie in front of the cameras.

References

[Bou98] S. Bougnoux. From Projective to Euclidean space under any practical situation, a criticism of self-calibration. In *Proc. 6th International Conference on Computer Vision, Bombay, India*, pages 790–796, January 1998.

[Fau92] O. D. Faugeras. What can be seen in three dimensions with an uncalibrated stereo rig? In *Proc. European Conference on Computer Vision*, LNCS 588, pages 563–578. Springer-Verlag, 1992.

[Har92] R. I. Hartley. Estimation of relative camera positions for uncalibrated cameras. In *Proc. European Conference on Computer Vision*, LNCS 588, pages 579–587. Springer-Verlag, 1992.

[Har98] R. I. Hartley. Chirality. *International Journal of Computer Vision*, 26(1):41–61, 1998.

[HGC92] R. I. Hartley, R. Gupta, and T. Chang. Stereo from uncalibrated cameras. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1992.

[Hor90] B. K. P. Horn. Relative orientation. *International Journal of Computer Vision*, 4:59–78, 1990.

[Hor91] B. K. P. Horn. Relative orientation revisited. *Journal of the Optical Society of America*, 8(10):1630–1638, 1991.

[HZ00] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

[LH81] H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, September 1981.

[NHBP96] G. Newsam, D. Q. Huynh, M. Brooks, and H. P. Pan. Recovering unknown focal lengths in self-calibration: An essentially linear algorithm and degenerate configurations. In *Int. Arch. Photogrammetry & Remote Sensing*, volume XXXI-B3, pages 575–80, Vienna, 1996.

[Oli00] J. Oliensis. A critique of structure from motion algorithms. *Computer Vision and Image Understanding*, 80:172–214, 2000.