Singular Vector Methods for Fundamental Matrix Computation

Ferran Espuny¹ and Pascal Monasse²

¹ School of Environmental Sciences, University of Liverpool
Ferran.Espuny-Pujol@liverpool.ac.uk

² Université Paris-Est, LIGM (UMR CNRS 8049),
Center for Visual Computing, ENPC, F-77455 Marne-la-Vallée
pascal.monasse@enpc.fr

Abstract. The normalized eight-point algorithm is broadly used for the computation of the fundamental matrix between two images given a set of correspondences. However, it performs poorly for low-size datasets due to the way in which the rank-two constraint is imposed on the fundamental matrix. We propose two new algorithms to enforce the rank-two constraint on the fundamental matrix in closed form. The first one restricts the projection on the manifold of fundamental matrices along the most favorable direction with respect to algebraic error. Its complexity is akin to the classical seven point algorithm. The second algorithm relaxes the search to the best plane with respect to the algebraic error. The minimization of this error amounts to finding the intersection of two bivariate cubic polynomial curves. These methods are based on the minimization of the algebraic error and perform equally well for large datasets. However, we show through synthetic and real experiments that the proposed algorithms compare favorably with the normalized eight-point algorithm for low-size datasets.

Keywords: 3D Reconstruction, Fundamental Matrix, Closed Form.

1 Introduction

1.1 Overview

The *fundamental matrix* is a two-view tensor which encodes the relative geometry between two images, named *epipolar geometry* [4,7]. Concisely, for each point in one image the fundamental matrix determines a line of possible corresponding points on the other image, called the *epipolar line* associated to the point. Since all epipolar lines in one image pass through the projection of the other camera center (the *epipole*), the fundamental matrix is constrained to have rank two.

The fundamental matrix plays a central role in the theory of 3D reconstruction from stereo images. It is used at initial stages to achieve a sparse projective reconstruction of a scene, and to extract the camera motion either from intrinsically calibrated cameras or by a self-calibration approach. It is also used to

rectify the images as a preliminary step in order to achieve a dense reconstruction of a scene. The rank-two constraint is crucial for motion estimation and for image rectification algorithms requiring the existence of epipoles [8,6].

For its robust computation in presence of outlier point correspondences, the fundamental matrix needs to be estimated from minimal subsets of data in random sampling methods. This is a requirement of the RANSAC [5] algorithm and variants. It can also be estimated using non-minimal but relatively small sets of data in the local random sampling method of Chum *et al.* [3]. Finally, it is estimated using all the available inlier correspondences, yielding an initial solution for an overall error optimization, known as *bundle adjustment* [15]; even in this context, the size of the available dataset can be low due to typical problems in feature detection and matching. In this paper we will focus on the closed-form and fast computation of the fundamental matrix given noisy inlier data; our final goal is to improve the accuracy of the existing approaches in this context.

The normalized eight-point algorithm introduced by Hartley [9] allows computing in closed form an estimate of the fundamental matrix given at least eight correspondences between two images. A closed-form variant using seven correspondences exists which benefits from the rank-two constraint (see e.g. Zhang [16]). This broadly used method is an improvement over previous ones achieved through a prior data normalization step.

The *rank-two constraint* on the fundamental matrix is imposed by the normalized eight-point algorithm in an *a posteriori* step which does not take into account the error to minimize. Iterative strategies have been proposed to overcome this drawback [10], and modern optimization approaches exist that take into account the rank-two constraint at some computational expense [2,11,17]. In contrast, we propose two closed-form solutions for the direct computation of a fundamental matrix satisfying the rank-two constraint.

Fig. 1. Geometric Interpretation of the Singular Vector Methods. We represent in black the hypersurface of $\mathbb{P}^8$ with equation $\det(F) = 0$ and, in front of it, the three matrices F_1, F_2, F_3 corresponding to the singular values $s_1 \leq s_2 \leq s_3$ of A in (4); in cyan, the level sets of the algebraic error function. The solution given by the 8-point algorithm (red) is obtained as the closest matrix in Froebenius norm to F_1 on the hypersurface. Our first proposed solution (green) lies on the intersection of the line F_1, F_2 with the hypersurface. Our second proposal (blue) lies on the intersection of the plane spanned by F_1, F_2, F_3 with the hypersurface.

1.2 The Rank-Two Constraint

The fundamental 3×3 matrix F between two images can be expressed as

$$F = K^{-T} [t]_{\times} R K'^{-1} , \quad (1)$$

with R the relative rotation matrix and t the translation vector between the images, K and K' being the camera intrinsic calibration matrices. $[t]_{\times}$ denotes the 3×3 matrix corresponding to the linear operator mapping a vector to its vector product with t on the left: $[t]_{\times} x = t \times x$.

Due to the factor $[t]_{\times}$ in (1), the matrix F has rank 2 (unless $t = 0$, meaning no parallax). Conversely, any rank-two matrix accepts such a decomposition. The epipoles are then $e = K t$ and $e' = K' R^T t$, respectively left and right null vectors of F : $e^T F = 0$ and $F e' = 0$. For several applications, it is mandatory to get a rank-two solution F .

For a point correspondence $x_i \leftrightarrow x'_i$ between the two images (expressed in homogeneous coordinates), the *epipolar constraint* is written

$$x_i^T F x'_i = 0 . \quad (2)$$

The *geometric error* of F at a point correspondence $x_i \leftrightarrow x'_i$ is the point to line distance $\text{dist}(x_i, Fx'_i)$. We take as geometric error of F the root mean square error (RMSE) of the geometric error of F over all the available correspondences:

$$\text{GeomErr}(F) = \sqrt{\frac{1}{n} \sum_{i=1}^n \text{dist}(x_i, Fx'_i)^2} . \quad (3)$$

The simplest existing method for the computation of the fundamental matrix proceeds as follows. For $n \geq 8$ point correspondences between two images, the data coordinates are normalized linearly (we use the *isotropic scaling* normalization described e.g. in [7]), and then the epipolar constraint (2) on the transformed coordinates is expressed as a linear system of the type

$$A f = 0 , \quad (4)$$

f being a 9×1 vector representing the coefficients of F . The system is homogeneous due to a scale ambiguity in the fundamental matrix.

We denote by f_i the i -th eigenvector of $A^T A$ with associated eigenvalue s_i^2 , $s_i \geq 0$ being the singular values of A sorted by *increasing value*, $s_1 \leq s_2 \leq \dots \leq s_9$, and we denote by F_i the 3×3 matrix associated to f_i (the right singular vectors of A). The homogeneous least-squares problem associated to (4) is the minimization of the so-called *algebraic error*:

$$\arg \min \|A f\| \text{ s.t. } \|f\| = 1 . \quad (5)$$

Its solution is given by f_1 , the eigenvector of $A^T A$ corresponding to the smallest singular value of A . The solution provided by the *Direct Linear Transform (DLT)*

is the de-normalization of the matrix F_1 , which does not satisfy the rank-two constraint. Setting to zero the smallest singular value of the SVD decomposition of F_1 , we obtain the rank-two matrix closest in Frobenius norm to the minimum of the algebraic error. This corresponds to an orthogonal projection of F_1 on the manifold of rank-two matrices, see Figure 1. The de-normalization of that matrix is the final solution given by the *normalized eight-point algorithm*.

In general, one is more interested in the minimization of the geometric error (3), the Gold standard error or its approximation, the Sampson error [7]. However, there is no known closed-form solution for the minimization of such errors. Therefore, there is still interest in simple closed-form methods with low computational cost, even if based on minimizing the algebraic error.

This paper proposes two alternative algorithms to enforce the rank-two constraint (see Figure 1). The solutions we propose intersect a line or a plane with the manifold of rank-two matrices. These are constructed by taking into account F_2 and F_3 , the next right-singular vectors of A , in order to have a better control of the increasing error when computing F . We show that this has a low computational burden, being closed-form, and generally provides a precision gain, especially with few correspondences or in near-degenerate cases, such as low b/h (ratio baseline/distance to scene).

In next section, we show through simulation the effect of the classical projection, its poor performance for low-size datasets and the relation with the singular vectors of A . Then, we present our alternative proposals based on two and three right-singular vectors of A instead of just one for the standard eight-point algorithm. We report the results on synthetic and real data before concluding.

Fig. 2. Performance of the DLT (cyan) and normalized eight-point (red) algorithms for variable size datasets and noisy data. Median over 5000 runs of the obtained geometric (3) and real (6) errors in dashed and continuous lines, respectively.

2 Performance of the Normalized Eight-Point Algorithm for Low-Size Datasets

We carried out the following simulation with the initial purpose of showing the performance of the normalized eight-point algorithm for low-size datasets, and later to compare this algorithm with our proposed improvements. We simulated

a camera with calibration matrix $K = \begin{pmatrix} 900 & 0 & 320 \\ 0 & 900 & 240 \\ 0 & 0 & 1 \end{pmatrix}$ with a 640×480 image, corresponding to an horizontal angular field of view of approximately 40 degrees. A variable number n of 3D points was generated at a minimum distance $h = 1$ in the direction of the principal axis inside a cuboid of depth ΔZ ; the width and height of the cuboid were chosen so that its image was inside the domain $[0, 640] \times [0, 480]$. A second camera position was generated at a distance b of the first one. The obtained image points were perturbed with uniform noise with standard deviation $\sigma = 1$ pixel. The experiment was repeated 5000 times for each value of n , and the geometric errors (3) of the obtained fundamental matrices were computed using the data correspondences.

In order to have an unbiased error measure of each fundamental matrix, we randomly generated a dense cloud of 1000 points inside the simulation cuboid and computed at each iteration as *real error* the geometric error of the exact correspondences $y_i \leftrightarrow y'_i$ obtained by projecting that cloud:

$$RealErr(F) = \sqrt{\frac{1}{1000} \sum_{i=1}^{1000} \text{dist}(y_i, Fy'_i)^2}. \quad (6)$$

Given that we are not considering outlier data, the geometric error (3) gives us a measure of the goodness of the fundamental matrix for reconstructing the given matches; in constrast, the real error (6) serves us to validate the goodness of the fundamental matrix for reconstructing new correspondences or for motion estimation (e.g. in self-calibration). This error is expected to be worse for low-size datasets, for which many points in the dense cloud used for evaluation are far from the few given data points used for model fitting. Notice that in absence of ground truth correspondences $y_i \leftrightarrow y'_i$, the real error is not computable.

We show in Figure 2 the results of a simulation attempting to cover a broad range of realistic scenes and camera motions: for each iteration, the scene depth variation ΔZ was randomly chosen between $10^{-4}h$ (satellite images) and h (close scenes), and the baseline size b was randomly chosen between $0.01h$ and h .

Observing the real error, a first remark is that for low-size datasets the enforcement of the rank-two constraint (eight-point algorithm) improves the solution given by F_1 , the minimum of the algebraic error. However, the geometric error (as well as the algebraic error, and the Sampson error [7]) increases with that enforcement. Observe that for low-size datasets the geometric errors corresponding to the noisy data are not necessarily correlated with the real errors of the models (which we estimate using a dense cloud of exact points). This does not hold for bigger size datasets, which in our simulation tend to be dense inside the cuboid. When dealing with real images, the errors can only be computed using the available data, which can be low-size due to typical problems in feature detection and matching.

A second remark is that the normalized eight-point algorithm behaves quite poorly for very low-size datasets, in comparison with its performance with general datasets. In our simulation, the real error (6) is divided by a factor of 3 when

considering datasets of $n = 12$ correspondences instead of $n = 8$. This behaviour was reported before in [3], where the normalized eight-point algorithm was compared to its variant for 7 correspondences, which takes as solution $F = F_1 + \alpha F_2$, for α such that $\det(F_1 + \alpha F_2) = 0$ [7].

It can be observed in Figure 3 that the groundtruth fundamental matrix, when decomposed into the orthonormal basis $\{F_i\}_{i=1,\dots,9}$, has significant coefficients not only in the right-singular vector F_1 , but also in the right-singular vector F_2 and, in minor measure, F_3 . This is because the matrices F_i are computed using noisy data. Based on this observation, we propose two solutions for a better closed-form computation of the fundamental matrix for low-size datasets.

Fig. 3. Significance of Singular Vectors. For low-size datasets, we show the average squared coefficients of the groundtruth fundamental matrix with respect to the three right-singular vectors F_1, F_2, F_3 (top), and only for F_2, F_3 (bottom).

3 The Two Singular Vector Algorithm

The normalized eight-point algorithm imposes the rank-two constraint on the fundamental matrix ignoring the fact that the algebraic error may increase drastically for certain directions. We propose a first alternative to this strategy which consists in selecting the rank-two matrix close to the algebraic minimum moving in the best direction for minimizing the algebraic error.

Whereas the solution of the normalized eight-point algorithm is given by a modification of F_1 , our first proposal, the *two singular vector algorithm*, is to compute the fundamental matrix as (see Figure 1):

$$\tilde{F} = F_1 + \alpha F_2, \text{ for } \alpha \text{ s.t. } \det(F_1 + \alpha F_2) = 0. \quad (7)$$

This proposal is commonly used for computing three possible fundamental matrices given 7 correspondences, and has been previously used in a non-minimal context without further validation [14]. The computation of α in (7) can be done in closed form as in those cases, consisting in the resolution of the univariate cubic equation in α given by (7). In case that we obtain three possible real values α when solving (7), we choose the one that has minimum geometric error (other error choices gave similar performance).

The algebraic error associated to a solution of (7) can be obtained straightforwardly using the singular values of A :

$$\|A \tilde{f}\|^2 = s_1^2 + \alpha^2 s_2^2. \quad (8)$$

Although the value of this error is clearly bigger than $\|A \tilde{f}_1\|^2 = s_1^2$, it may be smaller than the algebraic error of the normalized eight-point algorithm.

Using the general simulation proposed in Section 2, we show in Figure 4 the median (over 5000 runs) of the real and geometric errors associated to the two singular vector algorithm for noisy data ($\sigma = 1$ pixel). In Section 6 we will show using both simulation and real images that the performance of this method is superior to what could be concluded from the results hereby.

Fig. 4. Performance of the normalized eight-point algorithm (8pt, in red), the two singular vector algorithm (2sv, in green), and the three singular vector algorithm (3sv, in blue). We show the median over 5000 runs of the geometric (3) and real (6) errors.

4 The Three Singular Vector Algorithm

A plausible criticism to the previous algorithm is that, even if it imposes the rank constraint on the fundamental matrix, the increase in algebraic error is known but remains uncontrolled. Our second proposal addresses this issue by imposing the rank-two constraint while minimizing the algebraic error. Concisely, we compute the fundamental matrix as

$$\hat{F} = F_1 + \alpha F_2 + \beta F_3, \quad (9)$$

with

$$(\alpha, \beta) = \arg \min s_1^2 + \alpha^2 s_2^2 + \beta^2 s_3^2 \quad (10)$$

$$\text{s.t.} \quad G(\alpha, \beta) = 0$$

where $G(\alpha, \beta) = \det(F_1 + \alpha F_2 + \beta F_3)$.

Using Lagrange multipliers, the solution to the problem (10) can be searched inside the set of real values α, β, μ that satisfy:

$$2s_2^2 \alpha + \mu \partial_\alpha G(\alpha, \beta) = 0, \quad (11)$$

$$2s_3^2 \beta + \mu \partial_\beta G(\alpha, \beta) = 0, \quad (12)$$

$$G(\alpha, \beta) = 0. \quad (13)$$

Assuming for the moment that $\partial_\beta G(\alpha, \beta) \neq 0$, the equations (11) and (12) have a compatible solution μ if the following cubic equation in α, β holds

$$s_2^2 \alpha \partial_\beta G(\alpha, \beta) = s_3^2 \beta \partial_\alpha G(\alpha, \beta) . \quad (14)$$

Otherwise, if $\partial_\beta G(\alpha, \beta) = 0$, by (12) it holds $\beta = 0$, which can be obtained too by imposing equation (14).

Therefore, the solution(s) of (10) will lie in the intersection of two plane cubics, given by equations (13) and (14). In conclusion, theoretically there are at most 9 solutions of (10). In general, the minimum of (10) is reached at a unique value (α, β) .

The exact intersection of the two cubics (13) and (14) can be obtained by solving their resultant with respect to β , which is a univariate polynomial in α of degree 9. In practice, for each real root α of this polynomial, we compute the corresponding β using (13), and finally select as solution the (α, β) with minimum geometric error (the Sampson error gave similar performance). We obtained experimentally better results using the geometric error than using the algebraic error, which however we need to generate the set of candidate solutions in closed form.

Using again the general simulation proposed in Section 2, Figure 4 shows the median real and geometric errors associated to the three singular vector algorithm for noisy data ($\sigma = 1$ pixel). From the observed results we conclude that for general computations of the fundamental matrix (general camera motions and scenarios), the closed-form three singular vector algorithm outperforms the normalized eight-point algorithm for datasets of sizes between 8 and 12. Other applications are discussed in next section.

5 Applicability to General-Size Datasets

We have already shown, using the general scenario described in Section 2, that the use of more than one right-singular vector can be useful for low-size datasets (Figure 4). The reason is that the smallest singular values of the matrix A in (4) are close to each other for those datasets.

We can imagine two degenerate scenarios where this is also the case: small baseline motions and distant scenes. In fact, these problems can be solved by computing a planar homography instead of the fundamental matrix; the configurations close to them, however, cannot be modelled using a homography.

An example is given in satellite imagery, where the distance from the camera to the scene is much bigger than both the baseline size and the scene depth. We show in Figure 5 the simulation results (median over 5000 runs) corresponding to the *Pleiades* system (<http://smc.cnes.fr/PLEIADES/>), where, using the notation introduced in Section 2, we approximately have $b/h = 0.2$, $\Delta Z/h = 0.00014$. For any data size, the two singular vector algorithm outperforms the other closed-form methods.

Fig. 5. Performance with simulated satellite data. Geometric (3) and real (6) errors.

6 Experimental Results

6.1 Numerical Stability Evaluation

We plot in Figure 6 the errors as a function of noise standard deviation by the different methods, at various datasizes, for general camera motions and scenarios simulated as in Section 2; see also Figure 4 for a clear picture when $\sigma = 1$. A first remark is that all methods exhibit a linear dependence on noise level, showing that the proposed methods have correct numerical stability. In all cases, the 3sv method outperforms the 2sv method, which is to be expected concerning the algebraic error, since solutions of (7) are inside the minimization space of (9).

It can be also observed that for high noise levels the geometric error increases as the dataset size increases, the behaviour of the normalized eight-point algorithm for datasets of size $n = 8$ and close values being particularly critical. For the low dataset sizes, the 3sv outperforms the 8pt algorithm, whereas for the 2sv this seems only true for dataset sizes $n = 8, 9$. For bigger datasets, there is a very slight advantage for the normalized eight-point algorithm over 3sv.

Fig. 6. Errors against noise for different datasizes

6.2 Real Data

We first test the algorithms on stereo pairs from the SyntIm INRIA database (<http://perso.lcpc.fr/tarel.jean-philippe/syntim/paires.html>) (see Figure 7), which are small size images (mostly 512×512) with a moderate number of SIFT matches [12] obtained by using demanding parameters in

Fig. 7. Image pairs used for the validation with real images. From top to bottom and from left to right: Baballe, BalMouss, BatInria, BatSynt, Billet, Color, GrRub, Sabine, Serfaty, Sport; next, images from Hartley-Zisserman’s book [7], and then images from Strecha’s website; the last two row images are from Aanæs *et al.* [1].

Table 1. Results on 10 SyntIm stereo pairs, 5 pairs from Hartley-Zisserman’s book [7] and 3 pairs from Strecha’s site. The first column is the pair name, n the number of inlier correspondences after RANSAC. In each cell are indicated the root mean/max square geometric error over the data correspondences. In each test, the least error over the three constraint-enforcing methods (8pt, 2sv, and 3sv) is in **bold**.

Pair	n	8pt	2sv	3sv
Baballe	9	0.65/1.79	0.21/0.48	0.20/0.43
BalMouss	24	0.49/1.41	0.92/3.20	0.41/0.93
BatInria	46	0.59/ 1.87	0.82/2.52	0.58/1.92
BatSynt	26	0.52/1.84	0.46/1.74	0.59/ 1.31
Billet	11	0.39/0.72	0.16/0.38	0.15/0.28
Color	15	0.61/1.49	0.31/0.76	0.30/0.66
GrRub	24	0.66/1.54	0.66/1.45	0.65/1.52
Sabine	16	0.38/1.01	0.60/1.39	0.51/ 0.97
Serfaty	39	0.75/1.97	0.59/ 1.72	0.59/1.83
Sport	115	0.37/1.24	1.16/3.65	0.39/1.52
bt.000, bt.002	62	0.50/1.45	0.50/1.44	0.50/ 1.44
bt.000, bt.004	41	0.72/1.82	0.72/1.86	0.73/1.86
bt.000, bt.006	28	0.86/ 2.79	0.84/2.87	0.83/2.87
chapel	31	0.69/ 1.69	0.68/1.73	0.68/1.72
keble	50	0.33/1.01	0.33/ 0.95	0.32/1.02
Brussels	671	0.35/1.34	0.35/1.34	0.35/1.34
Dresden	1615	0.22/0.79	0.23/1.03	0.23/1.01
Leuven	904	0.48/1.69	0.55/1.89	0.54/2.13

order to avoid outliers. The SIFT matches are filtered by a non-parametric variant of RANSAC [13], and the remaining outliers are manually discarded to obtain a real noisy inlier dataset for our algorithms. We report in Table 1 the results. In all but two datasets, one of our proposed algorithms outperforms the 8-point algorithm, usually the three singular vector algorithm. We also tested on image pairs from Hartley-Zisserman’s book [7]

(<http://www.robots.ox.ac.uk/~vgg/hzbook/code/>) and from Strecha's website (<http://cvlab.epfl.ch/data/strechamvs/>). Results are reported in Table 1. Even though the number n of inliers is higher, we can see that 2sv and 3sv yield comparable precision.

Aanaes *et al.* [1] propose a large dataset of calibrated image sequences (<http://roboimagedata.imm.dtu.dk/index.html>). We take some $1/4$ size sequences (300×400) and consider as stereo pair the first image of a sequence and successive images in that sequence. Statistics are reported in Table 2. As we compare image 1 with further images in the sequence, the number n of inlier correspondences tends to decrease as the b/h ratio increases. Images 1 and 2 have $b/h \sim 0.03$, and images 1 and 6 have $b/h = 0.17$. In all sequences, the most accurate method is either 2sv or 3sv. Many times, one of both outperforms the eight-point algorithm, especially for small b/h . Notice for example the dramatic failure of the latter for pairs (1, 5) and (1, 6) of SET037. Even for relatively high n a significant improvement can be obtained with 3sv.

Table 2. Results with DTU data sets (see Table 1 caption). The data error (mean/max) is measured using available groundtruth calibration and depicted in column 'data error'.

Set	pair	n	data error	8pt	2sv	3sv
037	1,2	15	0.48/1.06	0.51/1.41	0.44/ 0.92	0.44 /0.94
037	1,3	13	0.67/1.83	0.51/ 0.41	0.50 /1.01	0.50/0.99
037	1,4	8	0.60/1.39	0.07/0.14	0.02/0.03	0.01/0.03
037	1,5	8	0.97/2.41	6.50/12.7	0.15/0.25	0.19/0.38
037	1,6	8	0.37/0.75	9.00/23.5	0.21/0.54	0.12/0.26
040	1,2	16	0.43/1.19	0.41/1.11	0.18 /0.43	0.19/ 0.35
040	1,3	17	0.63/1.34	0.32/0.86	0.30 /0.86	0.30/ 0.82
040	1,4	18	0.68/1.65	0.40/0.84	0.37/0.84	0.37/0.78
040	1,5	16	0.75/1.39	0.53/1.08	0.50/ 1.04	0.50 /1.06
040	1,6	15	1.04/2.59	1.67/4.41	0.54/ 0.90	0.54 /0.97
041	1,2	31	0.17/0.39	0.26/1.07	0.16 /0.38	0.16/ 0.36
041	1,3	27	0.32/1.21	0.77/2.06	0.29/ 0.88	0.25 /1.02
041	1,4	25	0.30/0.98	0.39/1.05	0.18/0.46	0.14/0.42
041	1,5	25	0.34/1.09	0.37/0.90	0.21/0.47	0.21/0.48
041	1,6	21	0.24/0.63	0.20 /0.50	0.61/1.82	0.20/ 0.49
042	1,2	21	0.46/1.87	0.50/1.46	0.23/ 0.42	0.23 /0.42
042	1,3	15	0.19/0.45	0.27/0.63	0.18/0.34	0.16/0.33
042	1,4	17	0.77/2.55	0.62/1.22	1.08/2.09	0.48/0.87
042	1,5	15	0.96/3.15	1.40/3.54	0.93/1.90	0.69/1.70
042	1,6	14	0.66/2.00	1.14/2.71	1.04/1.97	0.52/1.03

7 Conclusion

We have shown that a small modification of the eight-point algorithm provides in several cases a better precision. The two singular vector algorithm is quite similar to the seven-point algorithm. The three singular vector algorithm is slightly more

complex to implement, but it is more likely to improve the precision. Whatever the case, our recommendation is to compute the errors resulting from the three considered rank-two enforcement algorithms and keep the least of them.

Acknowledgments. Part of this work was funded by the Agence Nationale de la Recherche, Callisto project (ANR-09-CORD-003).

References

1. Aanæs, H., Dahl, A.L., Pedersen, K.S.: Interesting interest points. *Int. J. Comput. Vis.* 97, 18–35 (2012)
2. Chesi, G., Garulli, A., Vicino, A., Cipolla, R.: Estimating the fundamental matrix via constrained least-squares: a convex approach. *IEEE Trans. Pattern Anal. Mach. Intell.* 24(3), 397–401 (2002)
3. Chum, O., Matas, J., Kittler, J.: Locally optimized RANSAC. In: Michaelis, B., Krell, G. (eds.) *DAGM 2003. LNCS*, vol. 2781, pp. 236–243. Springer, Heidelberg (2003)
4. Faugeras, O., Luong, Q.-T., Papadopolou, T.: *The Geometry of Multiple Images: The Laws That Govern the Formation of Images of A Scene and Some of Their Applications*. MIT Press, Cambridge (2001) ISBN 0262062208
5. Fischler, M.A., Bolles, R.C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. ACM* 24(6), 381–395 (1981)
6. Gluckman, J., Nayar, S.K.: Rectifying transformations that minimize resampling effects. In: *Proc. CVPR* (2001)
7. Hartley, R., Zisserman, A.: *Multiple View Geometry in Computer Vision*, 2nd edn. Cambridge University Press (2004) ISBN 0521540518
8. Hartley, R.I.: Theory and practice of projective rectification. *Int. J. Comput. Vis.* 35, 115–127 (1999)
9. Hartley, R.I.: In defence of the 8-point algorithm. In: *Proc. ICCV* (1995)
10. Hartley, R.I.: Minimizing algebraic error in geometric estimation problems. In: *Proc. ICCV* (1998)
11. Kahl, F., Henrion, D.: Globally optimal estimates for geometric reconstruction problems. In: *Proc. ICCV* (2005)
12. Lowe, D.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vis.* 60, 91–110 (2004)
13. Moisan, L., Stival, B.: A probabilistic criterion to detect rigid point matches between two images and estimate the fundamental matrix. *Int. J. Comput. Vis.* 57(3), 201–218 (2004)
14. Schaffalitzky, F., Zisserman, A., Hartley, R.I., Torr, P.: A six point solution for structure and motion. In: Vernon, D. (ed.) *ECCV 2000. LNCS*, vol. 1842, pp. 632–648. Springer, Heidelberg (2000)
15. Triggs, B., McLauchlan, P., Hartley, R., Fitzgibbon, A.: Bundle adjustment – a modern synthesis. In: *Proc. VA Workshop* (1999)
16. Zhang, Z.: Determining the epipolar geometry and its uncertainty: a review. *Int. J. Comput. Vis.* 27(2), 161–195 (1998)
17. Zheng, Y., Sugimoto, S., Okumoti, M.: A branch and contract algorithm for globally optimal fundamental matrix estimation. In: *Proc. CVPR* (2011)