Outlier identification in stereo correspondences using quadrics

Richard J.M. den Hollander and Alan Hanjalic
Information and Communication Theory Group
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Mekelweg 4, 2628 CD Delft, The Netherlands
{r.j.m.denhollander,a.hanjalic}@ewi.tudelft.nl

Abstract

Two images containing the same rigid object or scene are related according to the epipolar geometry. Using local image correspondences, the fundamental matrix describing the epipolar geometry can be estimated. Due to the presence of outliers in these correspondences, robust estimation methods need to be used for the computation of the fundamental matrix. The ratio of outliers among the correspondences largely determines the complexity of these robust estimators. We propose an algorithm that is able to decrease the effective outlier ratio under a minimal computational complexity. The algorithm is based on the evaluation of the correspondences' positions with respect to a series of quadrics. It is generally applicable in the sense that it makes no assumptions about the cameras or imaged points. Experiments performed on both synthetic data and real images show the benefit of our approach.

1 Introduction

When an object is observed by cameras at two different positions, the resulting images will be related by the epipolar geometry [4]. The fundamental matrix describes the epipolar geometry, and can be computed from point correspondences between the images. These correspondences are found by local feature detectors [10, 6, 5]. In general, a certain amount of correspondences established by these methods is erroneous. When such *outliers* are present among the correspondences, robust methods need to be used to find the fundamental matrix [9]. Usually these methods impose a large complexity on the computational process.

One of the most frequently used robust estimators is RANSAC [2]. It is based on the random selection of a small number of correspondences, so that the free parameters of the corresponding fundamental matrix are set. The remaining data is then searched for support for this model. The complexity of the algorithm depends on the outlier ratio and the number of correspondences needed, which is seven for the fundamental matrix.

A way of speeding up RANSAC is by decreasing the outlier ratio, which can be accomplished by adjusting the sampling probabilities of correspondences. In [7] a method in the context of motion estimation is proposed, that uses the matching score of a correspondence to alter the probability of selecting that correspondence. The distinction can

BMVC 2005 doi:10.5244/C.19.80

be further improved if spatial information of the correspondences is taken into account. A method based on using this information is the ROR (Rejection of Outliers by Rotations) algorithm [1]. Here it is shown that there exists a 3D rotation of the points in one image, which makes the directions of all good correspondences in the joint image plane equal. The directions that do not conform to the most prominent direction, which is found by random trials, originate from outlying correspondences. Eliminating these correspondences will decrease the effective outlier ratio. The method shows very good results, but a drawback is that it makes a few assumptions about the data, like a dominant depth value of the world points and the absence of significant camera rotation about the principal axis. It also assumes the camera's focal length to be known.

In this paper we propose a method to decrease the outlier ratio by using spatial correspondence information, without making any assumptions about the cameras or imaged points. The method is based on the comparison of the point correspondences to a series of quadrics. Outlying correspondences tend to have an arbitrary position with respect to a quadric. When we gather statistics about the correspondences' positions, we can use this fact to decrease the effective outlier ratio. The complexity of the proposed algorithm is low, since only a small number of algebraic distances for each correspondence is computed.

In Section 2 we describe the principle of fundamental matrix estimation. In Section 3 the proposed algorithm for outlier identification is explained. A short discussion of the complexity of the proposed algorithm and ROR is given in Section 4. The evaluation of both algorithms for synthetic and real data is given in Section 5. Finally, in Section 6, we give some concluding remarks about the paper.

2 The epipolar geometry and robust estimation

Consider a set of points $\mathbf{X}_i$ for $i = 1, \ldots, n$ in $\mathbb{R}^3$. Each point $\mathbf{X}$ is indicated by homogeneous coordinates $\mathbf{X} = (X, Y, Z, 1)^\top$. When viewing this point through a camera, the point is projected onto the image plane. The 3D point $\mathbf{X}$ is then transformed into the 2D point $\mathbf{x} = (x, y, 1)^\top$ where x and y are the image plane coordinates. The camera matrix $\mathbf{P}$, which contains the camera parameters such as focal length and relative 3D position, relates the points by $\mathbf{x} = \mathbf{PX}$.

When viewing the same 3D point $\mathbf{X}$ through two cameras with matrices $\mathbf{P}$ and $\mathbf{P}'$, the relation between the projected 2D points $\mathbf{x} = \mathbf{PX}$ and $\mathbf{x}' = \mathbf{P}'\mathbf{X}$ is governed by the fundamental matrix $\mathbf{F}$ [4] according to

$$\mathbf{x}'^\top \mathbf{F}\mathbf{x} = 0 \tag{1}$$

The 3×3 matrix $\mathbf{F}$ is completely determined by the camera matrices $\mathbf{P}$ and $\mathbf{P}'$. Without knowledge of $\mathbf{P}$ and $\mathbf{P}'$, it can also be determined from a set of 7 corresponding points $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$.

We call a correspondence $\mathbf{x} \leftrightarrow \mathbf{x}'$ an outlier when the points $\mathbf{x}$ and $\mathbf{x}'$ are not projections of the same world point $\mathbf{X}$. In practice, the points in the images are derived from interest point detectors and local image neighborhoods are used for matching. When searching for correspondences between the images, errors may arise if neighborhoods differ due to large viewpoint changes.

Randomly selecting seven correspondences in the presence of outliers may not yield the correct fundamental matrix; if there are one or more outliers among the seven correspondences they will corrupt the estimate. A simple method to achieve robustness against outliers is the RANSAC [2] algorithm. Here one randomly picks seven correspondences and computes **F**. The remaining $n-7$ correspondences are then examined to see how close they are to the hypothesized **F**. If sufficient correspondences are close enough, the matrix **F** will be accepted as the solution. If not, another set of seven correspondences is picked and the process is repeated.

The complexity of RANSAC is given by the number of times M that a set of 7 correspondences needs to be picked. We determine M by requiring that with high probability p, say 0.99, a set of 7 inliers is picked at least once in M trials. If the outlier ratio is denoted by ε, we have the following relation between p, ε and M:

$$p = 1 - (1 - (1 - \varepsilon)^7)^M \tag{2}$$

Methods that are able to decrease the outlier ratio, prior to the execution of RANSAC, can therefore provide significant computational savings during **F** estimation.

3 Quadric comparisons

The fundamental matrix essentially defines a quadric relation on the points **X** in $\mathbb{R}^3$. This can be seen by writing

$$\mathbf{x}'^{\top}\mathbf{F}\mathbf{x} = \mathbf{X}^{\top}\mathbf{P}'^{\top}\mathbf{F}\mathbf{P}\mathbf{X} = 0 \tag{3}$$

where the quadric is given by $\mathbf{P}'^{\top}\mathbf{F}\mathbf{P}$. This quadric is special in the sense that the locus, i.e. the points **X** obeying (3), consists of all world points **X**.

Replacing **F** with an arbitrary 3×3 matrix **Q** will yield a quadric $\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}$, for which the locus in general consists of points **X** on a two-dimensional variety in $\mathbb{R}^3$. The quadric will split the points **X** into two sets: those with positive and those with negative algebraic distance to the quadric, that is

$$\mathbf{X}^{\top}\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}\mathbf{X} \gtrless 0 \tag{4}$$

If the quadric is shaped in such a way that the space with positive (or negative) algebraic distance is very small, few points **X** will yield $\mathbf{X}^{\top}\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}\mathbf{X} > 0$ (or $\mathbf{X}^{\top}\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}\mathbf{X} < 0$). This is equivalent to having few image points with $\mathbf{x}'^{\top}\mathbf{Q}\mathbf{x} > 0$ (or $\mathbf{x}'^{\top}\mathbf{Q}\mathbf{x} < 0$). In the case of outliers, however, there are not necessarily few points like that. Since the outliers do not correspond to real world points, their algebraic distances do not follow the same subdivision as for the inliers. Ideally, the outliers have an equal probability of yielding a positive or negative algebraic distance. The idea is that by examining statistics of the signs for several different quadrics, we are able to distinguish between inliers and outliers.

The type of quadric formed by **Q** depends on the rank of the symmetric part of $\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}$, which is given by

$$\frac{\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P} + \left(\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}\right)^{\top}}{2} \tag{5}$$

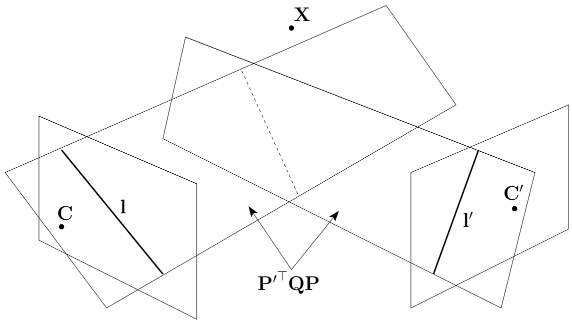

Figure 1: The quadric $\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}$ formed by selection of lines $\mathbf{l}$ and $\mathbf{l}'$ in the image planes. Both camera centers $\mathbf{C}$ and $\mathbf{C}'$ lie on a plane of the quadric, which subdivides the points $\mathbf{X}$ in the space.

The antisymmetric part, which equals (5) for a minus instead of a plus sign, always yields zero in the quadric relation and therefore does not contribute to $\mathbf{X}^{\top}\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}\mathbf{X}$. If we choose $\mathbf{Q}$ to have rank 1, then both $\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}$ and $(\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P})^{\top}$ from (5) have rank 1. As a result, the sum of these matrices in the symmetric part will in general have rank 2. In this case the locus corresponds to two planes. The camera centers $\mathbf{C}$ and $\mathbf{C}'$ each lie on one of the planes, so that this plane is visible as a line in the image. Therefore, the quadric of rank 2 can be defined by choosing appropriate lines $\mathbf{l}$ and $\mathbf{l}'$ in the image planes, see Fig. 1. It is possible to choose a $\mathbf{Q}$ with higher rank, but the abovementioned method seems to be the easiest way of letting the quadric pass through the means of the projected points. This is to make sure that there is enough variation in the sign of the algebraic distances. We therefore select the lines $\mathbf{l}$ and $\mathbf{l}'$ such that they pass through the means $\overline{\mathbf{x}} = (\overline{x}, \overline{y}, 1)^{\top}$ and $\overline{\mathbf{x}}' = (\overline{x}', \overline{y}', 1)^{\top}$ in the first and second image, respectively. If the angles of the lines with the image x-axis are denoted by θ and θ', the lines are given by

$$\mathbf{l} = (-\sin(\theta), \cos(\theta), \overline{x}\sin(\theta) - \overline{y}\cos(\theta))^{\top} \tag{6}$$

in the first image and

$$\mathbf{l}' = (-\sin(\theta'), \cos(\theta'), \overline{x}'\sin(\theta') - \overline{y}'\cos(\theta'))^{\top} \tag{7}$$

in the second. Then the matrix $\mathbf{Q}$ is constructed by

$$\mathbf{Q} = \mathbf{l}'\mathbf{l}^{\top} \tag{8}$$

which is a rank 1 matrix.

Note that we do not know the precise position of the planes of the quadric $\mathbf{P}'^{\top}\mathbf{Q}\mathbf{P}$, since the camera matrices $\mathbf{P}$ and $\mathbf{P}'$ are unknown. Yet, we can pick an arbitrary θ and θ' and thus $\mathbf{Q}$, and subsequently examine the sign of the algebraic distances of all correspondences $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$. If the sign of the outliers is more or less random, the inliers should determine the dominant sign. Therefore, a point belonging to the set with the dominant sign is likely to be an inlier. Counting these occurrences for several different $\mathbf{Q}$'s, will gather statistics about the probability of dealing with either an inlier or an outlier. The counts will be put in variables c_i for $i = 1, \ldots, n$, which are being updated after each quadric used. We select L equally spaced angles from $[0,\pi]$ for both θ and θ', and for every pair of θ and θ' a quadric is generated. It is also possible to select θ and θ' randomly for each $\mathbf{Q}$, which gives almost comparable results. However, due to the randomness we should run it several times for averaging and this is more costly. The resulting algorithm for equally spaced angles is given in Fig. 2.

- $c_i = 0 \qquad i = 1,\ldots,n$
- **for all** $\theta = 0, \frac{\pi}{L}, \frac{2\pi}{L}, \ldots, \pi - \frac{\pi}{L}$ **do**
  - **for all** $\theta' = 0, \frac{\pi}{L}, \frac{2\pi}{L}, \ldots, \pi - \frac{\pi}{L}$ **do**
    - Construct $\mathbf{Q}$ according to (6), (7) and (8).
    - Find the index sets $I_{pos} = \{i|\mathbf{x}'_i\mathbf{Q}\mathbf{x}_i > 0\}$ and $I_{neg} = \{i|\mathbf{x}'_i\mathbf{Q}\mathbf{x}_i < 0\}$, and determine
    $$I = \arg\max_{I_{pos},I_{neg}}(|I_{pos}|,|I_{neg}|)$$
    - Set the new counts
    $$c_i = \begin{cases} c_i + 1 & \text{if } i \in I \\ c_i & \text{if } i \notin I \end{cases} \qquad i = 1,\ldots,n$$
  - **end for**
- **end for**

Figure 2: The quadric algorithm.

The merit of the algorithm is the distinction that can be made using the values of the counts. In particular, when the probability of selecting a correspondence $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ is the relative number of counts c_i, the effective outlier ratio ε_{quad} after applying the quadric algorithm has become

$$\varepsilon_{quad} = \frac{\sum\limits_{i\in I_{out}} c_i}{\sum\limits_{i} c_i} \tag{9}$$

where I_{out} denotes the index set of all outliers.

4 Complexity

The computational savings for a single iteration of RANSAC are the computation of $\mathbf{F}$ and one or three times n distance computations [4]. We have chosen for $L = 8$ in the experiments, which means that for the quadric algorithm we need $64n$ algebraic distance computations. This is a minor computational load that justifies the use of the algorithm prior to RANSAC in almost every situation.

The standard implementation of the ROR algorithm, however, uses 1,000 random rotations to reject the outliers [1]. The authors propose to take the majority vote over 10 runs, so that the algorithm eventually requires $10,000n$ computations of segment angles and 10,000 computations of the mode of an angle distribution. This makes ROR much more costly than the quadric algorithm.

5 Experimental results

We will evaluate the proposed quadric algorithm by comparing it to RANSAC without any preprocessing and the ROR algorithm [1]. In ROR possible outliers are completely rejected, so that the effective outlier ratio ε_{ROR} is the ratio of the number of retained outliers and the total number of retained points. We have used the standard implementation[1] of the algorithm without adjusting any parameters. In the quadric algorithm we have used $L = 8$ angles for θ and θ'.

5.1 Synthetic data

We have generated synthetic data by randomly positioning points in a cube in $\mathbb{R}^3$. The cameras with equal internal parameters are randomly positioned on a sphere around the cube. The radius of the sphere is twice the edge length of the cube. The points are between 20 and 60 focal lengths away from the cameras. The image coordinates of the inliers are perturbed by Gaussian distributed noise with a standard deviation of 0.3% of the image size. For the outliers we randomly select two different space points, and use their non-corresponding projections as a data pair. The experiment is run 100 times for a particular outlier ratio ε, and each run will have different point and camera positions. For the ROR algorithm we have scaled the coordinate values to resemble realistic pixel values.

The results for 50 and 200 points are shown in Table 1 and 2, respectively. In the tables we have shown the average ε_{quad} and ε_{ROR} and their standard deviations over the 100 runs. In addition, we have calculated according to (2) for each run the *theoretical* number of iterations M, M_{quad} and M_{ROR} corresponding to ε, ε_{quad} and ε_{ROR}, respectively. The quantities M_{quad} and M_{ROR} were then averaged over all runs and are shown in the tables.

ε	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
$\overline{\varepsilon_{quad}}$	0.078	0.158	0.245	0.333	0.433	0.540	0.656	0.773	0.892
$\overline{\varepsilon_{ROR}}$	0.046	0.101	0.171	0.199	0.309	0.405	0.525	0.684	0.867
STD ε_{quad}	0.013	0.016	0.022	0.021	0.021	0.022	0.025	0.025	0.015
STD ε_{ROR}	0.055	0.093	0.122	0.155	0.175	0.188	0.181	0.140	0.075
M	8	20	54	163	588	2,809	$2.1\ 10^4$	$3.6\ 10^5$	$4.6\ 10^7$
$\overline{M_{quad}}$	6	14	32	78	254	1,125	9,769	$2.8\ 10^5$	$4.9\ 10^7$
$\overline{M_{ROR}}$	4	12	34	401	2,084	$1.7\ 10^6$	$1.3\ 10^7$	$6.6\ 10^6$	$1.4\ 10^9$

Table 1: The results for synthetic data containing 50 points. The quantities indicated are the outlier ratio ε, the average outlier ratios ε_{quad} and ε_{ROR} and their standard deviations, the number of iterations M, and the average number of iterations M_{quad} and M_{ROR}.

In general, the results for 200 points are better than for 50 points both in terms of the averages as well as the standard deviations. More points allow a better estimate of the side of the quadric that is largest. We can see that our method reduces the outlier ratio in both cases. An exception is the somewhat extreme case $\varepsilon = 0.9$ for 50 points, where only 5 inliers are present.

The synthetic data essentially violates the assumption of a dominant depth value in the ROR algorithm. Yet, in most runs the algorithm performs well, as can be seen by the

[1] The code is obtained from http://www.cs.technion.ac.il/Labs/Isl/Project/Projects_done/ror/.

ε	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
$\overline{\varepsilon_{quad}}$	0.075	0.155	0.237	0.330	0.425	0.528	0.637	0.756	0.883
$\overline{\varepsilon_{ROR}}$	0.049	0.098	0.148	0.210	0.293	0.387	0.505	0.657	0.860
STD ε_{quad}	0.006	0.011	0.012	0.016	0.017	0.016	0.015	0.015	0.011
STD ε_{ROR}	0.042	0.082	0.109	0.134	0.164	0.183	0.178	0.139	0.071
M	8	20	54	163	588	2,809	$2.1\ 10^4$	$3.6\ 10^5$	$4.6\ 10^7$
$\overline{M_{quad}}$	6	13	29	75	224	906	5,815	$1.0\ 10^5$	$2.3\ 10^7$
$\overline{M_{ROR}}$	5	11	26	69	452	7,430	$3.4\ 10^4$	$1.0\ 10^6$	$1.2\ 10^{10}$

Table 2: The results for synthetic data containing 200 points. The quantities indicated are the outlier ratio ε, the average outlier ratios ε_{quad} and ε_{ROR} and their standard deviations, the number of iterations M, and the average number of iterations M_{quad} and M_{ROR}.

lower average outlier ratio that it achieves when compared to our method. However, the standard deviation of ROR is larger, so there will be several runs where the outlier ratio is substantially higher. Since the number of iterations increases quickly with the outlier ratio, the average number of iterations for high outlier ratios is therefore much larger for ROR.

5.2 Real data

We have also applied the algorithms to a set of real stereo images. The correspondences were found by applying the SIFT keypoint detector[2] [5]. For every keypoint descriptor in the left image we found the nearest descriptor in the right image using Euclidean distance. As proposed in [5], the correspondence is retained if the distance is smaller than 0.8 times the distance to the second nearest neighbor. The inliers among the resulting correspondences are found by robustly estimating the fundamental matrix, for which we have used [8]. We manually identified any remaining incorrect correspondences among the inliers and labeled them as outliers. The images used are shown in Fig. 3. The left images contain the set of inliers, and the right images show the interest points belonging to the outliers. The results of applying the algorithms on the images are shown in Table 3.

Although the effective outlier ratio ε_{quad} is decreased compared to ε for all image pairs, the ROR algorithm shows impressive results here. Most image pairs apparently meet the assumptions that are needed for the algorithm. However, if we introduce some additional transformation like image rotation, the ROR algorithm may not be able to handle it very well. We rotated the right images of all image pairs and noticed that ε_{ROR} generally increases. In two cases, which are indicated in Table 3, the increase was such that ε_{ROR} became larger than the original ε.

The image pair "tea box", taken from the Amsterdam Library of Object Images [3], provides a challenge for the algorithms. The box is rotated between the views, and due to the similar text on both sides many outliers arise from the same regions. As a result, the outlier distribution is somewhat structured and confuses the ROR algorithm. Our algorithm also has difficulty with this image pair, but it does show a reduction in the outlier ratio.

The strength of the proposed method is therefore being able to perform consistently for various stereo pairs, rather than showing a superior performance for a constrained set of stereo images.

[2]The code is obtained from http://www.cs.ubc.ca/∼lowe/keypoints/.

image pair	n	ε	ε_{quad}	ε_{ROR}	M	M_{quad}	M_{ROR}
books	740	0.74	0.679	0.165	$5.7\ 10^4$	$1.3\ 10^4$	14
pile of books	548	0.82	0.788	0.286	$7.5\ 10^5$	$2.4\ 10^5$	47
Valbonne	299	0.58	0.512	0.056	1,996	697	5
U. British Columbia	911	0.56	0.481	0.089	1,441	452	7
U. British Columbia (rotated 180°)	911	0.56	0.481	0.598	1,441	452	2,712
corridor	262	0.43	0.380	0.082	234	129	6
Wadham college	921	0.71	0.655	0.183	$2.7\ 10^4$	7,914	17
Wadham college (rotated 60°)	921	0.71	0.657	0.737	$2.7\ 10^4$	8,243	$5.3\ 10^4$
tea box	221	0.71	0.698	1	$2.7\ 10^4$	$2.0\ 10^4$	∞

Table 3: The results on the image pairs of Fig. 3. The quantities indicated are the total number of correspondences n, the outlier ratios ε, ε_{quad} and ε_{ROR}, and the number of iterations M, M_{quad} and M_{ROR}.

6 Conclusion

We have proposed an algorithm that compares the point correspondences of stereo images to a series of quadrics. Using the fact that outliers in the correspondences tend to have an arbitrary position with respect to a quadric, we can decrease the effective outlier ratio by gathering statistics over multiple quadrics. Subsequent application of a robust estimator will therefore be less complex. The quadric algorithm has a low computational complexity and requires no assumptions on the data. Experiments show that the algorithm reduces the original outlier ratio of both synthetic and real data sets. In the synthetic data experiment, the quadric algorithm shows better performance than a previously proposed method, which itself is much more complex. Although this method outperforms the quadric algorithm for many real data sets, the quadric algorithm shows the most consistent performance.

Future work includes investigating to what extent the camera positions influence the performance of the algorithm.

References

[1] A. Adam, E. Rivlin, and I. Shimshoni. Ror: Rejection of outliers by rotations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(1):78–84, 2001.

[2] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the Association for Computing Machinery*, 24(6):381–395, 1981.

[3] J. M. Geusebroek, G. J. Burghouts, and A. W. M. Smeulders. The Amsterdam library of object images. *Int. Journal of Computer Vision*, 61(1):103–112, January 2005.

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

[5] D.G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. Journal of Computer Vision*, 60(2):91–110, 2004.

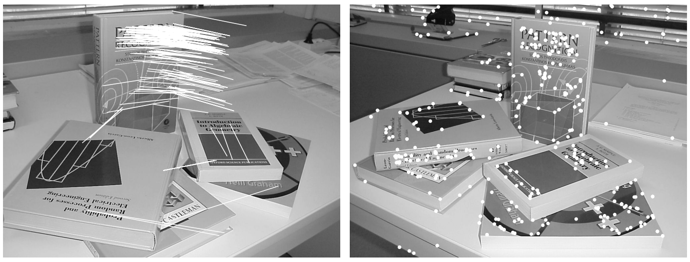

(a) books

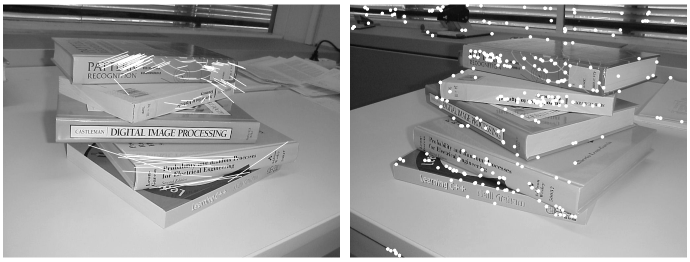

(b) pile of books

(c) Valbonne

Figure 3: Image pairs used in the experiments. The left images contain the inlying correspondences and the right images the outlying interest points.

[6] K. Mikolajczyk and C. Schmid. Scale and affine invariant interest point detectors. *Int. Journal of Computer Vision*, 60(1):63–86, 2004.

[7] B. Tordoff and D.W. Murray. Guided sampling and consensus for motion estimation. In *Proc. ECCV*, 2002.

[8] P.H.S. Torr. A structure and motion toolkit in matlab. Technical Report MSR-TR-2002-56, Microsoft Research, June 2002.

[9] P.H.S. Torr and D.W. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *Int. Journal of Computer Vision*, 24(3):271–300, 1997.

(d) U. British Columbia

(e) corridor

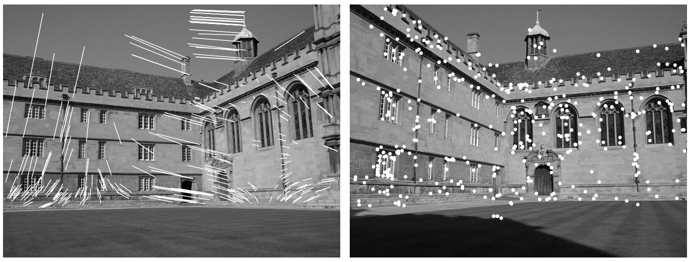

(f) Wadham college

(g) tea box

Figure 3: (continued)

[10] T. Tuytelaars and L. van Gool. Wide baseline stereo matching based on local, affinely invariant regions. In *Proc. BMVC*, pages 412–425, 2000.