

Guided Sampling via Weak Motion Models and Outlier Sample Generation for Epipolar Geometry Estimation

Liran Goshen

Ilan Shimshoni*

Faculty of Industrial Engineering & Management
Technion - Israel Institute of Technology
Haifa, Israel 32000

Abstract

The problem of automatic robust estimation of the epipolar geometry in cases where the correspondences are contaminated with a high percentage of outliers is addressed. This situation often occurs when the images have undergone a significant deformation, either due to large rotation or wide baseline of the cameras. An accelerated algorithm for the identification of the false matches between the views is presented. The algorithm generates a set of weak motion models (WMMs). Each WMM roughly approximates the motion of correspondences from one image to the other. The algorithm represents the distribution of the median of the geometric distances of a correspondence to the WMMs as a mixture model of outlier correspondences and inlier correspondences. The algorithm generates an outlier correspondence sample from the data. This sample is used to estimate the outlier rate and to estimate the outlier pdf. Using these two pdfs the probability that each correspondence is an inlier is estimated. These probabilities enable to guide the sampling. In the RANSAC process this guided sampling accelerates the search process. The resulting algorithm when tested on real images achieves a speedup of between one or two orders of magnitude!

1. Introduction

Recovery of epipolar geometry is a fundamental problem in computer vision. The RANdom SAMple Consensus algorithm (RANSAC) [4] has been widely used in computer vision in particular for recovering the epipolar geometry.

The RANSAC algorithm is simple but powerful. Repeatedly, random subsets are selected from the input data and the model parameters fitting the subset are calculated. The size of a random sample is the smallest sufficient for determining the model parameters. In each iteration the quality of the model is evaluated on the full data set. At the end, the

*This work was supported partly by grant 01-99-08430 of the Israeli Space Agency through the Ministry of Science Culture and Sports of Israel.

$s \setminus \epsilon$	25%	50%	60%	70%	80%	85%
3	8	34	70	168	573	1,362
7	33	588	2,808	21,055	3.5E05	2.6E06

Table 1: The number of samples I required to ensure, with probability $p = 0.99$, that at least one sample has no outliers for a given size of sample s and proportion of outliers ϵ .

model with the largest consensus set is returned. Two issues immediately arise:

1. How to check the quality of the model? Different cost functions may be used, the standard being the number of data points consistent with the model. In [9] a method of maximum likelihood estimation by sampling consensus (MLESC) is described. MLESC evaluates the likelihood of the hypothesis, representing the error distribution as a mixture model in which the inlier error is Gaussian and the outlier error is uniform. In [2] the need for user supplied threshold is eliminated by reformulating another robust method, the M-estimator, as a projection pursuit optimization problem. The projection based pbM-estimator automatically derives the threshold from univariate kernel density estimates.

2. When to stop the algorithm? The number of iterations is chosen sufficiently high [4, 7] to ensure with probability p that at least one of the random samples of s points is free from outliers. Usually p is chosen to be 0.99. Suppose ϵ is the probability that any selected feature is an outlier. Then at least I selections are required, where $(1 - (1 - \epsilon)^s)^I = 1 - p$, thus

$$I = \log(1 - p) / \log(1 - (1 - \epsilon)^s). \quad (1)$$

Table 1 gives the values of I for different values of s and ϵ .

Although ϵ is not generally known in advance, a lower bound can be estimated from the largest consistent set observed so far. It is widely appreciated that this stopping criterion is often wildly optimistic [6, 3] because with noisy data it is not enough to have a sample composed only of inliers, they must be inliers that span the manifold. This

significantly reduces the number of sample sets that will accurately span the manifold. Several approaches have been suggested to speed-up RANSAC. ROR [1] can speed-up the sampling process by reducing the contamination level ϵ of the data points when the camera’s internal parameters are known. ROR exploits the possibility of rotating one of the images to achieve some common behavior of the inliers. ROR may be run as a postprocessing step on output from any point matching algorithm. In R-RANSAC [5] the computational savings are achieved by typically evaluating only a fraction of the data points for each of the hypothesized models. Hypothesized models which pass the initial test are evaluated with all of the data points. LO-RANSAC [3] exploits the fact that the model hypothesis from an uncontaminated minimal sample is often sufficiently near the optimal solution and a local optimization step is applied to selected models. The number of samples which the LO-RANSAC performs achieves a good agreement with the theoretical predictions of Eq. (1).

As can be seen in Table 1, the number of needed iterations increases when the percentage of outliers is over 50% and it dramatically increases after 70%. The reason this is that for ϵ close to 1 the expression in Eq. (1) can be approximated by $I \approx \log(1 - p)/(1 - \epsilon)^s$, which is a s degree polynomial in $1/(1 - \epsilon)$. Note that the iteration numbers in the table are the theoretic ones and in practice they can be much higher. Situations with high percentages of outliers often occur when the images have undergone a significant deformation, either due to large rotation or wide baseline of the cameras [8].

In this paper we propose novel improvements to the robust estimation of epipolar geometry. The main goal is to assign probabilities to the putative correspondences and to use them in the RANSAC step. We use *weak motion models* (WMMs) to estimate these probabilities which are more informative than the correlation scores used for this purpose in [6].

The algorithm generates a set of WMMs which are similar in spirit to weak classifiers in supervised learning cases. Each WMM roughly approximates the motion of correspondences from one image to the other. The algorithm represents the distribution of the median of the geometric distances of a correspondence to the WMMs as a mixture model of outlier correspondences and inlier correspondences. Typically the inlier correspondences are closer to the WMMs than the outlier correspondences. The algorithm then generates an outlier correspondence sample. This sample is used to estimate the outlier rate and to estimate the outlier pdf. When generating the sample we take into consideration the corner distribution on the images, the similarity between the corners and the corner matching technique. Using these two distributions, the probability that each correspondence is an inlier is estimated. These probabilities

enable to guide the sampling, i.e. in each iteration a random subset of seven correspondences is chosen where each correspondence in this subset is chosen according to its probability.

The two main contributions of this paper are: i) the use of WMMs to probe the correspondence set and ii) the generation and the use of the outlier correspondence sample. Our approach makes it possible to break the iteration number theoretical boundary given by Eq. (1). The acceleration is especially significant for outlier percentages above 70%. In such outlier percentages our algorithm achieves a speedup of between one to two orders of magnitude!!

The paper is organized as follows. Section 2 describes the WMMs. Section 3 presents the mixture density model and the generation of the outlier correspondence sample. The generation of the WMMs is described in Section 4. Section 5 presents the details of the algorithm. The results are shown and discussed in Section 6. The paper is concluded in Section 7.

2 Weak motion model

Given a set of N putative correspondences $\{p_i \leftrightarrow p'_i\}$, we seek to find the set of inlier correspondences and we wish to compute the fundamental matrix between the two images. The algorithm that we describe in the next sections uses WMMs to solve these problems. A WMM is a motion model that roughly approximates the motion that a point underwent from the first image to the second image. In general each point in the first image moves to its corresponding epipolar line. In [11] a probabilistic motion model was presented limiting the motion to a segment of the epipolar line. A WMM should approximate the motion of the inlier correspondences. Thus the agreement of the putative correspondences to the WMMs should yield a probability that a correspondence is an inlier. These probabilities enable guided sampling from the putative correspondences. We use an affine transformation as a WMM, i.e.

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} & t_x \\ a_{21} & a_{22} & t_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix},$$

where the matrix $A = [a_{ij}]$ is an invertible 2×2 matrix and (t_x, t_y) is a $2D$ translation. Note that three points in the joint image space [10] define a unique affine transformation and that an affine transformation is a linear transformation that defines a two dimensional manifold in the joint image space. We have found that very often an affine transformation that has been formed from three inlier correspondences can be used as a WMM. Let $\{w_j\}$ be a set of N_w WMMs. Let d_{ij} be the geometric distance from correspondence $p_i \leftrightarrow p'_i$ to the manifold defined by the WMM

w_j , i.e.

$$d_{ij} = \min_{\tilde{p}} \|p_i - \tilde{p}\|^2 + \|p'_i - H_{a_j}\tilde{p}\|^2,$$

where H_{a_j} is the affine transformation representing w_j . This geometric distance can be easily analytically calculated because the affine transformation is a linear transformation.

Let d_i be the median distance of $p_i \leftrightarrow p'_i$ to the manifolds, i.e.

$$d_i = \text{med}\{d_{ij}\}_{j=1}^{N_w}.$$

This median distance can be thought of as a random variable and is modeled as a mixture model:

$$f_d(d_i) = f_{in}(d_i)(1 - \epsilon) + f_{out}(d_i)\epsilon,$$

where $f_{in}(d_i) = f(d_i|p_i \leftrightarrow p'_i \text{ inlier})$, $f_{out}(d_i) = f(d_i|p_i \leftrightarrow p'_i \text{ outlier})$ and ϵ is the mixing parameter which is the probability that any selected correspondence is an outlier.

3 Probability estimation using outlier sample

The probability, $P_{in}(i)$, that correspondence $p_i \leftrightarrow p'_i$ is an inlier can be calculated by

$$P_{in}(i) = \frac{f_{in}(d_i)(1 - \epsilon)}{f_d(d_i)}.$$

We estimate this probability in a non-parametric manner. All we assume is that the median distances of the inlier correspondences, d_i , are bounded by an unknown parameter D and that the outlier correspondences are not. Thus

$$f_d(d_i) = \begin{cases} f_{in}(d_i)(1 - \epsilon) + f_{out}(d_i)\epsilon, & d_i \leq D; \\ f_{out}(d_i)\epsilon, & \text{otherwise.} \end{cases}$$

We obtain

$$P_{in}(i) = \begin{cases} \frac{f_{in}(d_i) - f_{out}(d_i)\epsilon}{f_d(d_i)}, & d_i \leq D; \\ 0, & \text{otherwise.} \end{cases}$$

The algorithm estimates $f_d()$ using a kernel density estimator. The kernel estimator with given kernel K is defined by

$$p\widehat{df}(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{x - X_i}{h}\right),$$

where h is the kernel width, n is the number of sampled data points and X_i is the i^{th} observation of the random variable

X . We use a Gaussian kernel function and the L-Stage Plug In method [12] to estimate the bandwidth, i.e.

$$h = \left(\frac{4\hat{\sigma}^5}{3n}\right)^{1/5},$$

where $\hat{\sigma}$ is the sample standard deviation.

The estimation of $f_{out}()$ is more problematic. We usually do not have any prior knowledge about this pdf. Therefore we turn to generate a sample of outlier correspondences. Given such a sample $\{o_i \leftrightarrow o'_i\}, i = 1, \dots, N_o$ we can estimate $f_{out}()$ like we estimated $f_d()$.

We have tried three methods to generate this outlier sample.

1. **Uniformly.** Each point of each correspondence is sampled uniformly from the images.
2. **Corner based.** Each point of each correspondence is sampled uniformly from the corner sets of the images.
3. **Algorithm guided.** Using the same algorithm that generates the putative correspondences $\{p_i \leftrightarrow p'_i\}$, only this time it excludes from the input of the algorithm the correspondences that already have been used to generate $\{p_i \leftrightarrow p'_i\}$. In our experiments only mutually best candidates were selected as putative correspondences. So the entries that were selected for the putative correspondences were removed and the outlier sample was generated using the same method.

We found that the algorithm guided method gives the best estimation to the outlier distribution. To demonstrate this, a set of ten WMMs was randomly generated for the image pair in Fig. 1. Fig. 2 shows the distribution of the actual outliers and of the three aforementioned methods. The distributions in this figure and in rest of the figures in this paper are shown with Gaussian kernel smoothing with the L-Stage Plug In method for bandwidth estimation. In Fig. 2 the algorithm guided method gives a much better estimate of the $f_{out}()$ than the other two methods.

The reason for this is that the outliers have a distribution that depends on several factors. It depends on the distribution of the corners in the images. Usually there are regions in the image that have a larger density of corners than others. For example in Fig. 1 the region of the white wall has much less corners than the slide projector region. The distribution also depends on the texture around the corners. For example a corner on the desk of Fig. 1 has a better chance to be matched to another corner on the desk, because the texture around the corners on the desk is more similar to each other than to corners from other regions. The second method takes into consideration the distribution of the corners in the images. Fig. 1 shows that the second method gives a slightly better estimation relative to the first method. The algorithm guided method gives the best estimation. It

takes into consideration the corner distribution, the similarity between the corners and the matching technique.

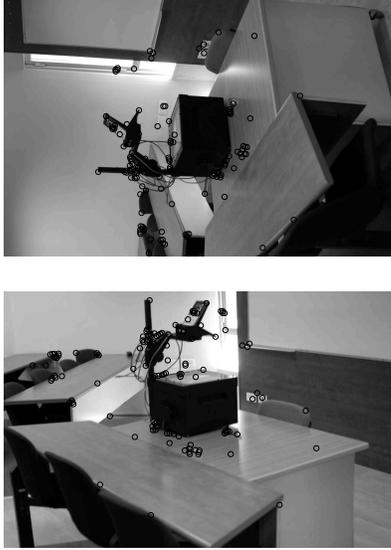


Figure 1: Slide projector image pair, with correct matching points marked.

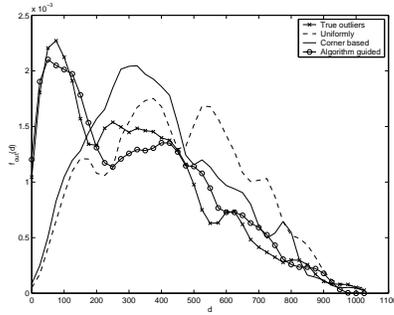


Figure 2: The distributions for the slide projector image pair of the actual outliers, uniformly method, corner based method and the algorithm guided method. The algorithm guided method gives the best estimation.

The algorithm estimates D as the smallest distance at which the number of putative correspondences less the number of expected outliers according to the outlier sample is equal to the number of expected inliers, i.e.

$$D = \min\{d \mid \left[\sum_{i=1}^N I(d_i < d) - \frac{N}{N_o} \sum_{i=1}^{N_o} I(d_i^o < d) \right] = \lceil (1 - \epsilon)N \rceil\}, \quad (2)$$

where d_i^o is the median of the distances of $o_i \leftrightarrow o'_i$ to the WMM manifolds and $I()$ is the index function, i.e.

$I(True) = 1$ and $I(False) = 0$. The estimation of ϵ is described in Section 5.

The following figures show results for the slide projector image pair. Fig. 3 shows $f_d()$, $f_{out}()$ and $f_{in}()$ density functions and Fig. 4 shows histograms of $P_{in}()$ for inlier and outlier correspondences. Note that a large number of outlier correspondences get probability zero to be an inlier and a large number of the inlier correspondences get probability 0.6 to be an inlier while the original inlier rate is only 0.22. These probabilities are used to guide the RANSAC step of the algorithm enabling to accelerate the search process. For illustration Fig. 5 shows some of the inlier correspondences with solid white lines, three outlier correspondences with a probability greater than zero to be an inlier with black solid lines and three outlier correspondences with probability zero to be an inlier with white dashed lines. Note that outlier correspondences with probabilities greater than zero have motions similar to the inliers and can not be discarded using the WMMs. However outlier correspondences with probability zero have a different motion relative to the inliers and have been successfully detected using the WMMs.

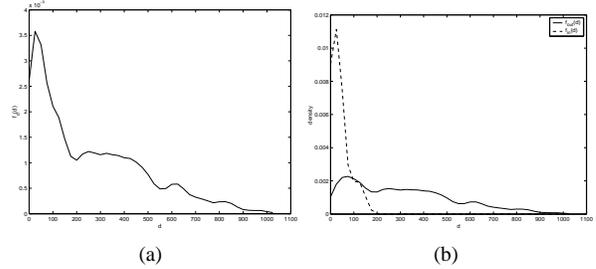


Figure 3: Density function for the slide projector image pair (a) $f_d(d)$ density function (b) $f_{out}(d)$ and $f_{in}(d)$ density functions.

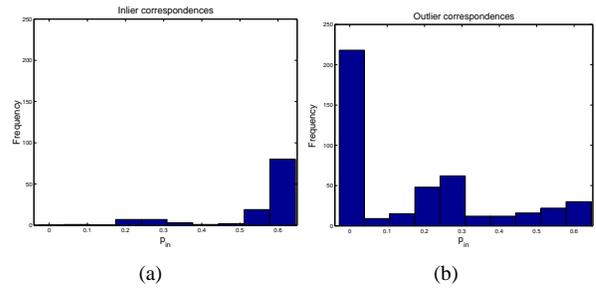


Figure 4: Histograms for $P_{in}()$ of (a) inlier correspondences (b) outlier correspondences for the slide projector image pair.

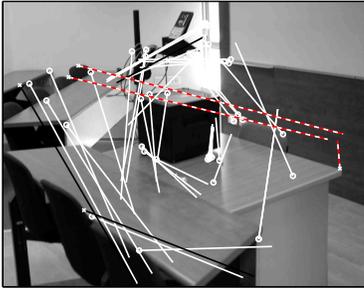


Figure 5: Several inlier correspondences are shown with solid white lines, three outlier correspondences with a probability greater than zero to be an inlier are shown with black solid lines and three outlier correspondences with probability zero to be an inlier are shown with dashed white lines.

4 Weak motion model generation

In order to get a set of approximately N_w WMMs. A RANSAC style algorithm is used. At each iteration a subset of three correspondences is randomly chosen from the putative correspondences. Each subset uniquely defines an affine transformation. After $N_w/(1-\epsilon)^3$ iterations we have on average N_w subsets containing only inlier correspondences. Each subset gets a score. The score is the geometric distance of the $(1-\epsilon)N$ th closest correspondence to the manifold of the affine transformation of the subset. A subset that fits an affine transformation that can be used as a WMM gets a low score because at least $(1-\epsilon)N$ inliers are relatively close to the manifold. An affine transformation based on a subset that includes outliers usually can not be used as a WMM and does not have a set of $(1-\epsilon)N$ correspondences that are relatively close to its manifold and because of that will get a relatively high score. We take the set of N_w models with the lowest scores to be the set of the WMMs.

5 The algorithm

In this section we describe the flow of the algorithm. The generation of the WMMs and the computation of the probability of each correspondence to be an inlier depend on ϵ . In general ϵ is not known in advance. The algorithm searches for the correct ϵ starting from the lowest rate level up to the highest rate level of a rate level set, $\{\epsilon_j\}$, given as input by the user. The algorithm is summarized in Algorithm 1. The details of the algorithm are as follows:

Generation of outlier sample

The algorithm generates an outlier sample. This step is done only once at the beginning of the algorithm using the algo-

Algorithm 1 Guided Sampling via WMMs and Outlier Sample Generation

- 1: Generate an outlier sample
 - 2: For each ϵ_j starting from the smallest up to the highest:
 - a) Generate random WMMs for the current rate level ϵ_j
 - b) Fine tune the estimate of the outlier rate
 - c) Estimate inlier probabilities
 - d) Estimate the number of iterations, \hat{N}_s , for the guided LO-RANSAC
 - e) If $\hat{N}_s < N_t$ then execute guided LO-RANSAC
 - f) Check the stopping criteria of the algorithm
-

rithm guided method to generate the outlier sample.

Generation of random weak motion models

The generation of the random WMMs is done in the same manner described in Section 4. However in Section 4 we assumed that ϵ is known in advance. As in practice this is not the case we work on the rate level set $\{\epsilon_j\}$ in parallel. Each rate level, ϵ_j , is associated with a number of iterations $N_{I_j} = \lceil N_w/(1-\epsilon_j)^3 \rceil$. Each rate level has a set, B_j , of N_w suspected WMMs. These sets are updated during the generation of the random affine transformations.

Each set B_j keeps the N_w best affine transformations found until now that are suspected to be WMMs. In each iteration a new subset of three correspondences is randomly chosen and the appropriate affine transformation is formed. The geometric distance of each of the correspondences to the manifold of the affine transformation is calculated. Each rate level ϵ_j gives a different score to this affine transformation. The score is the distance of the $(1-\epsilon_j)N$ th closest correspondence to the manifold of the affine transformation. Each of the rate levels that the algorithm did not pass until now, keeps the current affine transformation if it has one of the N_w lower scores obtained till now. In our implementation each of the WMM sets, B_j , is kept in a different heap. This enables very efficient updates of the sets. The random generation of the affine transformations continues until the number of iterations reaches N_{I_j} .

Fine tune the estimate of the outlier rate

The outlier rate level set $\{\epsilon_j\}$ is given as an input from the user. So, no ϵ_j is expected to be the accurate outlier rate of the putative correspondences. In this step the algorithm fine tunes the estimation of the outlier rate. First the median distance sets of the putative correspondences $\{d_i\}$ and of the outlier sample $\{d_i^o\}$ are calculated according to the current set of WMMs, B_j , where j is the index of the current

outlier rate level. The search for the accurate outlier rate is carried out around ϵ_j in the region $A = [\frac{\epsilon_j + \epsilon_{j-1}}{2}, \frac{\epsilon_j + \epsilon_{j+1}}{2}]$. Note that for each $\epsilon \in A$ there is an appropriate D_ϵ according to Eq. (2). The search seeks to find $\hat{\epsilon}_j$ that gives the best fit between the outlier sample distribution and putative correspondence distribution for distances greater than D . Let $cdf_{out}()$ be the empirical distribution function of the outlier sample, i.e. $cdf_{out}(d) = \frac{\sum_{i=1}^{N_o} I(d_i^o < d)}{N_o}$ and let $cdf_{mix}()$ be the empirical distribution function of the outlier pdf given by the mixture model, i.e. $cdf_{mix}(d) = (\frac{\sum_{i=1}^N I(d_i < d)}{N} - (1 - \epsilon))/\epsilon$ for $d > D$. The quality of the fit between these cdfs is measured by the number of zero crossing of $cdf_{out}(d) - cdf_{mix}(d)$, where a large number represents a good fit. To ensure that this score will not be biased toward lower values of ϵ the parameter D is set to be $D_{(\epsilon_j + \epsilon_{j-1})/2}$ for all of the checked values of ϵ . Other methods for cdf comparison can also be used in this step, such as the Kolmogorov-Smirnov test.

Inlier probability estimation

The probability, $P_{in}(i)$, that correspondence $p_i \leftrightarrow p'_i$ is an inlier is calculated for all the correspondences by

$$P_{in}(i) = \begin{cases} \frac{f_d(d_i) - f_{out}(d_i)\hat{\epsilon}_j}{f_d(d_i)}, & d_i \leq D_{\hat{\epsilon}_j}; \\ 0, & \text{otherwise.} \end{cases}$$

The estimate of the density function at each point is done using the kernel density estimator described in Section 3.

Iteration number estimation

The inlier probability estimation is used to accelerate the search for the fundamental matrix using a guided RANSAC algorithm. The correspondences are sampled according to their inlier probability. In each guided RANSAC iteration a subset sample of seven correspondences is chosen. Correspondence i has probability $\frac{P_{in}(i)}{\sum_{i=1}^N P_{in}(i)}$ to be chosen. The number of samples N_s is chosen sufficiently high to ensure with probability p that at least one of the samples of seven correspondences is free from outliers. Thus Eq. (1) has to be generalized. Let S be the set of all possible subsets of seven correspondences. The probability of subset $s_i \in S$ to be chosen in the guided sampling step is

$$P_c(s_i) = \prod_{j=1}^7 \frac{P_{in}(c_{ij})}{\sum_{k=1}^N P_{in}(k)},$$

where c_{ij} is the j^{th} correspondence in subset s_i . The calculation of N_s has to take into consideration all the possible combination series of subsets s_i where each subset is chosen according the probability function $P_c()$. This kind of

calculation is infeasible, however we can estimate N_s using a Monte Carlo method. We generate a series M of length N_M of subsets s_i . Each member in the series is chosen in the same manner as in the guided sampling step, i.e. according to the probability function $P_c()$. The series M represents an ‘‘average’’ subseries of the guided sampling step. We calculate how many such subseries should be taken in order to ensure with probability p that at least one of the samples of seven correspondences is free from outliers. We have

$$\left(\prod_{i=1}^{N_M} (1 - P_c(m_i)) \right)^K = 1 - p,$$

where $m_i \in M$ is a subset of seven correspondences and K is the number of subseries needed to be taken in order to ensure the probability p . Now we can estimate N_s as

$$\hat{N}_s = K * N_M.$$

The length of series M should be sufficiently long to ensure that M represents an ‘‘average’’ subseries of the guided sampling step.

Guided LO-RANSAC

The guided LO-RANSAC step is not carried out when the WMMs did not provide enough probabilistic information to reduce the number of iterations \hat{N}_s below a user defined threshold N_t . In this step a slightly different version of LO-RANSAC algorithm is applied. The difference is in the sampling technique. Instead of random sampling, the samples are chosen according to the probabilities $P_{in}(i)$. In each iteration a subset of seven correspondences is sampled. The number of iterations is limited to \hat{N}_s . At the end of this step, the fundamental matrix F with the largest number of inliers N_{in_j} is kept.

Note that the algorithm has two RANSAC style steps, the WMM generation step and the guided LO-RANSAC step. Each iteration of the second type is more costly because in each iteration a fundamental matrix has to be computed based on seven correspondences. This involves computing the roots of a polynomial of degree three which may have up to three solutions for the fundamental matrix and the distances of the putative correspondences to each of the non linear manifolds of the fundamental matrices have to be calculated. On the other hand, in each iteration of the random WMM generation step a unique affine transformation is formed using a linear algorithm and the distances of the putative correspondences to the linear manifold of the affine transformation have to be calculated. In our implementation the run time of an average iteration of the LO-RANSAC is 4.5 times more costly than an average iteration of the random WMM generation step.

	Characteristics			LO-RANSAC			Proposed algorithm			Speedup rate	
	N	In.	ϵ	75%	End	In.	75%	End	In.	75%	End
Slide pr.	565	122	0.78	70,878	384,171	114.8	326 (427)	1,388 (1,274)	118.9	18.2	59.2
Desk	303	63	0.79	80,841	341,958	61.7	168 (411)	495 (908)	62.1	42.1	78.3
Lab	730	131	0.82	>232,596	>860,590	124.3	109 (2,216)	1,214 (2534)	129.0	>90	>114
Corridor	401	70	0.83	>228,675	1,000,000	59.7	1,038 (1,327)	1,645 (2,227)	66.4	>137	>178
Building	363	44	0.88	>821170	1,000,000	26.8	1,042 (4,635)	2,202 (7,832)	41.9	>208	>251
Road	401	50	0.88	>547,547	1,000,000	37.6	1,532 (1,937)	5,168 (10,000)	49.2	>270	>181
Yard	508	53	0.90	>508,519	1,000,000	43.5	91 (10,000)	2,843 (10,000)	50.9	>265	>185

Table 2: Experiment characteristics and results. N is the number of putative correspondences, “In.” is the number of inlier correspondences and ϵ is the outlier rate level. For each algorithm “75%” is the average number of samples until at least 75% of the true inliers are found, “End” is the average number of iterations until the termination of the algorithm and “In.” is the average number of inliers that were found until the termination of the algorithm. The numbers in parentheses are the average numbers of the WMMs that have been generated. The last two columns show the average time speedup rate until at least 75% of the possible inliers are found and until the termination of the algorithm. Several numbers in the table are shown with the “greater than” symbol that indicates that there have been cases in which the LO-RANSAC algorithm terminated when it exceeded one million iterations.

Checking the stopping criteria of the algorithm

In this step the algorithm checks the stopping criteria. The algorithm has two estimates of the number of inliers: $N(1-\hat{\epsilon}_j)$ from the mixture model and N_{in_j} from the guided LO-RANSAC. If the two estimations are approximately the same then the algorithm terminates and returns the fundamental matrix F and the set of N_{in_j} inlier correspondences. The algorithm also terminates when that the current estimation of the number of inliers is lower than the largest inlier set found until this stage. This situation is checked after the estimation of $\hat{\epsilon}_j$ in the fine tuning step of the algorithm. In this case the algorithm returns the largest found inlier set and its fundamental matrix.

6 Experiments

The proposed algorithm was tested on many image pairs of indoor and outdoor scenes several of which are presented here. All the putative correspondences were detected and matched automatically using a very simple method. Corners were detected by the Harris corner detector. Cross correlation scores were generated between all corners with each patch undergoing 36 evenly spaced rotations, the strongest matches over these rotations are stored. Only mutually best candidates were selected as putative correspondences. We have compared our algorithm with the LO-RANSAC algorithm which gives near perfect agreement with the theoretical performance that is given by Eq. (1). The termination criterion based on Eq. (1) was set to $p = 0.99$. In cases where the number of iterations exceeded one million the algorithm also terminated. The parameters for the proposed algorithm were as follows. The outlier rate level set, $\{\epsilon_j\}$, was set to $\{0.1, 0.25, 0.5, 0.6, 0.7, 0.75, 0.8, 0.85, 0.9, 0.925, 0.95\}$

, the number of WMMs generated for each of the outlier rate levels was set to $N_w = 10$, the length of the series M was set to $N_M = 1000$ and the threshold for the guided LO-RANSAC step was $N_t = 3000$. Seven experiments are presented in this section. Some of the image pairs are shown in Fig. 6. The cases that are presented here are difficult cases in which the outlier rate is between 0.78-0.9. The characterization of the scenes, the number of putative correspondences, the number of inlier correspondences and the outlier rate are summarized in Table 2. Each algorithm has been applied to each image pair twenty times. For each algorithm the following statistics are presented: The number of samples until at least 75% of the true inliers are found (as was suggested in [6]), the number of iterations until the termination of the algorithm and the number of inliers that have been found until the termination of the algorithm. The number of generated WMMs is also reported. The results of the experiments are summarized in Table 2. The average speedup rate time achieved in the experiments until at least 75% of the possible inliers have been found is between 18.2-270 and the average speedup rate time until the termination of the algorithm is between 59.2-251. Note that several numbers in the table are shown with the “greater than” symbol that indicates that at least in one case the LO-RANSAC algorithm terminated when it exceeded one million iterations. Removal of this early termination criterion yields much higher results. In addition, the number of inlier correspondences found by the proposed algorithm is larger than in the ones achieved by LO-RANSAC and the results are near perfect.

7 Summary and Conclusions

In this work, we presented a novel algorithm for automatic robust estimation of the epipolar geometry in cases where



Figure 6: Some of the real image pairs on which the algorithms were tested.

the correspondences are contaminated with a high percentage of outliers. The algorithm uses WMMs that roughly approximate the motion of correspondences from one image to the other. The algorithm generates a set of WMMs and represents the distribution of the median of geometric distances of a correspondence to the WMMs as a mixture model of outlier correspondences and inlier correspondences. An outlier correspondence sample is then generated using the algorithm guided method. This method takes into consideration the corner distribution in the images, the similarity between the corner neighborhoods and the matching technique. The outlier sample is used to estimate the outlier rate and to estimate the outlier pdf. These pdfs are then used to estimate the probability that each correspondence is an inlier. These probabilities guide the sampling which accelerates the search process. The performance of the algorithm was compared to the performance of the LO-RANSAC algorithm that gives near perfect agreement with the number of iterations given by Eq. (1). The resulting algorithm when tested on real image pairs achieves speedup rates of 59-251!

References

- [1] A. Adam, E. Rivlin, and I. Shimshoni. ROR: Rejection of outliers by rotations. *IEEE Trans. Patt. Anal. Mach. Intell.*, 23(1):78–84, January 2001.
- [2] H. Chen and P. Meer. Robust regression with projection based m-estimators. In *International Conference on Computer Vision*, pages 878–885, 2003.
- [3] O. Chum, J. Matas, and J.V. Kittler. Locally optimized RANSAC. In *German Pattern Recognition Symposium*, pages 236–243, 2003.
- [4] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM*, 24(6):381–395, June 1981.
- [5] J. Matas and O. Chum. Randomized RANSAC with td,d test. *Image and Vision Computing*, 22(10):837–842, September 2004.
- [6] B. Tordoff and D.W. Murray. Guided sampling and consensus for motion estimation. In *European Conference on Computer Vision*, pages I: 82–96, 2002.
- [7] P.H.S. Torr. Motion segmentation and outlier detection. In *PhD thesis, Dept. of Engineering Science, University of Oxford*, 1995.
- [8] P.H.S. Torr and C. Davidson. IMPSAC: Synthesis of importance sampling and random sample consensus. *IEEE Trans. Patt. Anal. Mach. Intell.*, 25(3):354–364, March 2003.
- [9] P.H.S. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Comp. Vis. Im. Understanding*, 78(1):138–156, April 2000.
- [10] B. Triggs. Matching constraints and the joint image. In *International Conference on Computer Vision*, pages 338–343, 1995.
- [11] B. Triggs. Joint feature distributions for image correspondence. In *International Conference on Computer Vision*, pages II: 201–208, 2001.
- [12] M.P. Wand and M.C. Jones. *Kernel Smoothing*. Chapman & Hall, 1995.