Locally Optimized RANSAC †

Ondřej Chum[1,2], Jiří Matas[1,2], and Josef Kittler[2]

[1] Center for Machine Perception, Czech Technical University, Faculty of Electrical Engineering
Dept. of Cybernetics, Karlovo nám. 13, 121 35 Prague, Czech Republic

[2] CVSSP, University of Surrey, Guildford GU2 7XH, United Kingdom

Abstract

A new enhancement of RANSAC*, the locally optimized* RANSAC *(*LO-RANSAC*), is introduced. It has been observed that, to find an optimal solution (with a given probability), the number of samples drawn in* RANSAC *is significantly higher than predicted from the mathematical model. This is due to the incorrect assumption, that a model with parameters computed from an outlier-free sample is consistent with all inliers. The assumption rarely holds in practice. The locally optimized* RANSAC *makes no new assumptions about the data, on the contrary - it makes the above-mentioned assumption valid by applying local optimization to the solution estimated from the random sample.*

The performance of the improved RANSAC *is evaluated in a number of epipolar geometry and homography estimation experiments. Compared with standard* RANSAC*, the speed-up achieved is two to three fold and the quality of the solution (measured by the number of inliers) is increased by 10-20%. The number of samples drawn is in good agreement with theoretical predictions.*

1 Introduction

Many computer vision algorithms include a robust estimation step where model parameters are computed from a data set containing a significant proportion of outliers. The RANSAC algorithm introduced by Fishler and Bolles in 1981 [3] is possibly the most widely used robust estimator in the field of computer vision. RANSAC has been applied in the context of short baseline stereo [13, 12], wide baseline stereo matching [9, 15, 10, 6], motion segmentation [13], mosaicing [7], detection of geometric primitives [2], robust eigenimage matching [5] and elsewhere.

In a classical formulation of RANSAC, the problem is to find all inliers in a set of data points. The number of inliers I is typically not known a priori. Inliers are data points consistent with the 'best' model, e.g. epipolar geometry or homography in a two view correspondence problem or line or ellipse parameters in the case of detection of geometric primitives. The RANSAC procedure finds, with a certain probability, all inliers and the corresponding model by repeatedly drawing random samples from the input set of data points.

† The authors were supported by the European Union under projects IST-2001-32184, ICA 1-CT-2000-70002 and by the Czech Ministry of Education under project LN00B096 and by the Czech Technical University under project CTU0306013. The images for experiments **A**,**B**, and **E** were kindly provided by T. Tuytelaars (VISICS, K.U.Leuven), **C** by M. Pollefeys (VISICS, K.U.Leuven), and **E** by K. Mikolajczyk (INRIA Rhône-Alpes).

RANSAC is popular because it is simple and it works well in practice. The reason is that almost no assumptions are made about the data and no (unrealistic) conditions have to be satisfied for RANSAC to succeed. However, it has been observed experimentally that RANSAC runs much longer (even by an order of magnitude) than theoretically predicted [11]. The discrepancy is due to one assumption of RANSAC that is rarely true in practice: it is assumed that a model with parameters computed from *an uncontaminated sample* is consistent with *all* inliers.

In this paper we propose a novel improvement of RANSAC exploiting the fact that the model hypothesis from an uncontaminated minimal sample is almost always sufficiently near the optimal solution and a local optimization step applied to selected models produces an algorithm with near perfect agreement with theoretical (i.e. optimal) performance. This approach not only increases the number of inliers found and consequently speeds up the RANSAC procedure by allowing its earlier termination, but also returns models of higher quality. The increase of average time spent in a single RANSAC verification step is minimal. The proposed optimization strategy guarantees that the number of samples to which the optimization is applied is insignificant.

The main contributions of this paper are (a) modification of the RANSAC that simultaneously improve the speed of the algorithm and and the quality of the solution (which is near to optimal) (b) introduction of two local optimization methods (c) a rule for application of the local optimization and a theoretical analysis showing the local optimization is applied at most $\log k$ times, where k is the number of samples drawn. In experiments on two image geometry estimation (epipolar geometry and homography) the speed-up achieved is two to three fold.

The improvement proposed in this paper requires no extra input information or prior knowledge, and it does not interfere with other modifications of the algorithm, the MLESAC [14], R-RANSAC [1] and NAPSAC [8]. MLESAC, proposed by Torr and Zisserman, defines a cost function in the maximal likelihood framework.

The structure of this paper is as follows. First, in Section 2, the motivation of this paper is discussed in detail and the general algorithm of locally optimized RANSAC is described. Four different methods of local optimization are proposed in Section 3. All methods are experimentally tested and evaluated through epipolar geometry and homography estimation. The results are shown and discussed in Section 4. The paper is concluded in Section 5.

2 Algorithm

The structure of the RANSAC algorithm is simple but powerful. Repeatedly, subsets are randomly selected from the input data and model parameters fitting the sample are computed. The size of the random samples is the smallest sufficient for determining model parameters. In a second step, the quality of the model parameters is evaluated on the full data set. Different cost functions may be used [14] for the evaluation, the standard being the number of inliers, i.e. the number of data points consistent with the model. The process is terminated [3, 13] when the likelihood of finding a better model becomes low, i.e. the probability η of missing a set of inliers of size I within k samples falls under predefined threshold

$$\eta = (1 - P_I)^k. \tag{1}$$

Repeat until the probability of finding better solution falls under predefined threshold, as in (1):

1. Select a random sample of the minimum number of data points S_m.
2. Estimate the model parameters consistent with this minimal set.
3. Calculate the number of inliers I_k, i.e. the data points their error is smaller than predefined threshold θ.
4. If new maximum has occurred ($I_k > I_j$ for all $j < k$), run **local optimization**. Store the best model.

Algorithm 1: A brief summary of the LO-RANSAC

Symbol P_I stands for the probability, that an uncontaminated sample of size m is randomly selected from N data points

$$P_I = \frac{\binom{I}{m}}{\binom{N}{m}} = \prod_{j=0}^{m-1} \frac{I-j}{N-j} \approx \varepsilon^m, \tag{2}$$

where ε is the fraction of inliers $\varepsilon = I/N$. The number of samples that has to be drawn to ensure given η is

$$k = \log(\eta)/\log(1 - P_I).$$

From equations (1) and (2), it can be seen, that termination criterion based on probability η expects that a selection of a single random sample not contaminated by outliers is followed by a discovery of whole set of I inliers. However, this assumption is often not valid since inliers are perturbed by noise. Since RANSAC generates hypotheses from minimal sets, the influence of noise is not negligible, and the set of correspondences the size of which is smaller than I is found. The consequence is an increase in the number of samples before the termination of the algorithm. The effect is clearly visible in the histograms of the number of inliers found by standard RANSAC. The first column of Figure 2 shows the histogram for five matching experiments. The number of inliers varies by about 20-30%.

We propose a modification that increases the number of inliers found near to the optimum I. This is achieved via a local optimization of so-far-the-best samples. For the summary of the locally optimized RANSAC see Algorithm 1. The local optimization step is carried out only if a new maximum in the number of inliers from the current sample has occurred, i.e. when standard RANSAC stores its best result. The number of consistent data points with a model from a randomly selected sample can be thought of as a random variable with unknown (or very complicated) density function. This density function is the same for all samples, so the probability that k-th sample will be the best so far is $1/k$. Then, the average number of reaching the so-far-the-best sample within k samples is

$$\sum_{1}^{k} \frac{1}{x} \leq \int_{1}^{k} \frac{1}{x}\, dx + 1 = \log k + 1.$$

Note, that this is the upper bound as the number of correspondences is finite and discrete and so the same number of inliers will occur often. This theoretical bound was confirmed experimentally, the average numbers of local optimization over an execution of (locally optimized) RANSAC can be found in Table 3. For more details about experiments see Section 4.

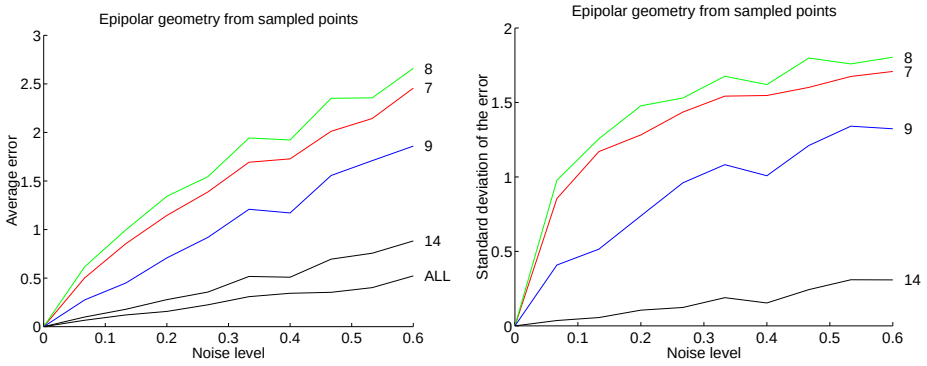

Fig. 1. The average error (left) and the standard deviation of the error for samples of 7,8,9, 14 and all 100 points respectively with respect to the noise level.

3 Local Optimization Methods

The following methods of local optimization have been tested. The choice is motivated by the two observations that are given later in this section.

1. Standard. The standard implementation of RANSAC without any local optimization.
2. Simple. Take all data points with error smaller than θ and use a linear algorithm to hypothesize new model parameters.
3. Iterative. Take all data points with error smaller that $K \cdot \theta$ and use linear algorithm to compute new model parameters. Reduce the threshold and iterate until the threshold is θ.
4. Inner RANSAC. A new sampling procedure is executed. Samples are selected only form I_k data points consistent with the hypothesised model of k-th step of RANSAC. New models are verified against whole set of data points. As the sampling is running on inlier data, there is no need for the size of sample to be minimal. On the contrary, the size of the sample is selected to minimize the error of the model parameter estimation. In our experiments the size of samples are set to $\min(I_k/2, 14)$ for epipolar geometry (see results in Section 3) and to $\min(I_k/2, 12)$ for the case of homography estimation. The number of repetitions is set to ten in the experiments presented.
5. Inner RANSAC with iteration. This method is similar to the previous one, the difference being that each sample of the inner RANSAC is processed by method 3.

The local optimization methods are based on the two following observations.

Observation 1: The Size of Sample
The less information (data points) is used to estimate the model parameters in the presence of noise, the less accurate the model is. The reason for RANSAC to draw minimal samples is that every extra point exponentially decreases the probability of selecting an outlier-free sample, which is approximately ε^m where m is the size of the sample (i.e. the number of data points included in the sample).

It has been shown in [13], that the fundamental matrix estimated from a seven point sample is more precise than the one estimated form eight points using a linear algorithm [4]. This is due to the singularity enforcement in the eight point algorithm. However,

the following experiment shows, that this holds only for eight point samples and taking nine or more points gives more stable results than those obtained when the fundamental matrix is computed from seven points only.

Experiment: This experiment shows, how the quality of a hypothesis depends on the number of correspondences used to calculate the fundamental matrix. For seven points, the seven point algorithm was used [13] and for eight and more points the linear algorithm [4] was used. The course of experiment was as follows. Noise of different levels was added to the noise-free image points correspondences divided into two sets of hundred correspondences. Samples of different sizes were drawn from the first set and the average error over the second was computed. This was repeated 1000 times for each noise level. Results are displayed in Figure 1.

This experiment demonstrates, that the more points are used to estimate the model (in this case fundamental matrix) the more precise solution is obtained (with the exception of eight points). The experiment also shows that the minimal sample gives hypotheses of rather poor quality. One can use different cost functions that are more complicated than simply the number of inliers, but evaluating this function only at parameters arising from the minimal sample will get results at best equal to the proposed method of local optimization.

Observation 2: Iterative Scheme

It is well known from the robust statistic literature, that pseudo-robust algorithms that first estimate model parameters from all data by least squares minimization, then remove the data points with the biggest error (or residual) and iteratively repeat this procedure do not lead to correct estimates. It can be easily shown, that a single far–outlying data point, i.e. leverage point, will cause a total destruction of the estimated model parameters. That is because such a leverage point overweights even the majority of inliers in least-squares minimization. This algorithm works only well, when the outliers are not overbearing, so the majority of inliers have bigger influence on the least squares.

In local optimization method 3 there are no leverage points, as each data point has error below $K \cdot \theta$ subject to the sampled model.

4 Experimental Results

The proposed algorithm was extensively tested on the problem of estimation of the two view relations (epipolar geometry and homography) from image point correspondences. Five experiments are presented in this section, all of them on publicly available data, depicted in Figures 3 and 4. In experiments A and B, the epipolar geometry is estimated in a wide-baseline setting. In experiment C, the epipolar geometry was estimated too, this time from short-baseline stereo images. From the point of view of RANSAC use, the narrow and wide baseline problems differ by the number of correspondences and inliers (see Table 1), and also by the distribution of errors of outliers. Experiments D and E try to recover homography. The scene in experiment E is the same as in experiment A and this experiment could be seen as a plane segmentation. All tentative correspondences were detected and matched automatically.

Algorithms were implemented in C and the experiments were ran on AMD K7 1800+ MHz processor. The terminating criterion based on equation (1) was set

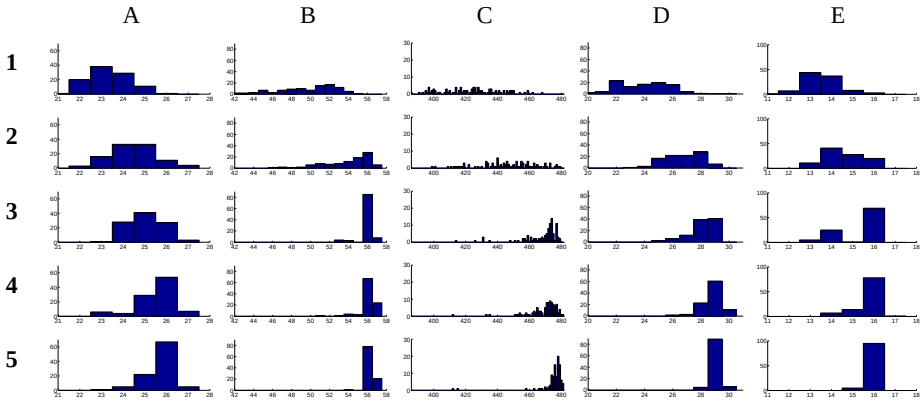

Fig. 2. Histograms of the number of inliers. The methods 1 to 5 (1 stands for standard RANSAC) are stored in rows and different dataset are shown in columns (A to E). On each graph, there is a number of inliers on the x-axis and how many times this number was reached within one hundred repetitions on the y-axis.

to $\eta < 0.05$. The threshold θ was set to $\theta = 3.84\sigma^2$ for the epipolar geometry and $\theta = 5.99\sigma^2$ for the homography. In both cases the expected σ was set to $\sigma = 0.3$.

The characterization of the matching problem, such as number of correspondences, the total number of inliers and expected number of samples, are summarized in Table 1. The total number of inliers was set to the maximal number of inliers obtained over all methods over all repetitions. The expected number of samples was calculated according to the termination criterion mentioned above.

Performance of local optimization methods 1 to 5 was evaluated on problems A to E. The results for 100 runs are summarized in Table 2. For each experiment, a table containing the average number of inliers, average number of samples drawn, average time spent in RANSAC (in seconds) and efficiency (the ratio of the number of samples drawn and expected) is shown. Table 3 shows both, how many times the local optimization has been applied and the theoretical upper bound derived in Section 2.

The method 5 achieved the best results in all experiments in the number of samples and differs slightly from the theoretically expected number. On the other hand standard RANSAC exceeds this limit 2.5 – 3.3 times. In Figure 2 the histograms of the sizes of the resulting inliers sets are shown. Each column shows results for one method, each row for one experiment. One can observe that the peaks are shifting to the higher values with the increasing identification number of method.

Method 5 reaches the best results in terms of sizes of inlier sets and consequently in number of samples before termination. This method should be used when the fraction of inliers is low. Resampling, on the other hand, might be quite costly in the case of high number of inliers, especially if accompanied by a small number of correspondences in total) as could be seen in experiment A (61 % of inliers out of 94 correspondences). In this case, method 3 was the fastest. Method 3 obtained significantly better results than the standard RANSAC in all experiments, the speed up was about 100%, and slightly worse than for method 5. We suggest to use method 5. Method 3 might be used in real-

Fig. 3. Image pairs and detected points used in epipolar geometry experiments (A - C). Inliers are marked as dots in left images and outliers as crosses in right images.

Fig. 4. Image pairs and detected points used in homography experiments (D and E). Inliers are marked as dots in left images and outliers as crosses in right images.

	A	B	C	D	E
# corr	94	94	1500	160	94
# inl	57	27	481	30	17
ε	61%	29%	32%	19%	18%
# sam	115	34529	8852	2873	3837

Table 1. Characteristics of experiments A-E. Total number of correspondences, maximal number of inliers found within all tests, fraction of inliers ε and theoretically expected number of samples.

		1	2	3	4	5
A	inl	49.7	53.9	55.9	56.0	**56.2**
	sam	383	205	129	117	**115**
	time	0.018	0.010	**0.007**	0.010	0.019
	eff	3.35	1.79	1.12	1.02	**1.01**
B	inl	23.3	24.4	25.0	25.5	**25.7**
	sam	90816	63391	49962	44016	**39886**
	time	3.911	2.729	2.154	1.901	**1.731**
	eff	2.63	1.84	1.45	1.27	**1.16**
C	inl	423.5	446.2	467.5	468.9	**474.9**
	sam	25205	16564	11932	10947	**9916**
	time	4.114	2.707	1.971	1.850	**1.850**
	eff	2.85	1.87	1.35	1.24	**1.12**
D	inl	23.9	26.7	28.1	28.8	**29.0**
	sam	8652	5092	3936	3509	**3316**
	time	0.922	0.543	0.423	**0.387**	0.391
	eff	3.01	1.77	1.37	1.22	**1.15**
E	inl	13.5	14.6	15.3	15.7	**15.9**
	sam	12042	8551	6846	5613	**5254**
	time	0.979	0.696	0.559	0.463	**0.444**
	eff	3.14	2.23	1.78	1.46	**1.37**

Table 2. The summary of local optimization experiments: average number of inliers (inl) and samples taken (sam), average time in seconds and efficiency (eff). The best values for each row are highlighted in bold. For more details see the description in text in Section 4.

time procedures when a high number of inliers is expected. Methods 2 and 4 are inferior to methods with iteration (3 and 5 respectively) without any time saving advantage.

5 Conclusions

An inprovement of the RANSAC algorithm was introduced. The number of detected inliers increased, and consequently the number of samples drawn decreased. In all experiments, the running-time is reduced by a factor of at least two, which may be very

	1		2		3		4		5	
A	3.0	5.9	2.6	5.3	2.0	4.9	1.9	4.8	1.8	4.7
B	6.4	11.4	6.1	11.1	5.9	10.8	6.0	10.7	5.9	10.6
C	7.7	10.1	6.8	9.7	6.5	9.4	6.7	9.3	6.5	9.2

	1		2		3		4		5	
D	5.2	9.1	4.8	8.5	4.5	8.3	4.4	8.2	4.0	8.1
E	4.8	9.4	4.3	9.1	4.2	8.8	4.0	8.6	3.9	8.6

Table 3. The average number of local optimizations ran during one execution of RANSAC and logarithm of average number of samples for comparison.

important in real-time application incorporating a RANSAC step. It has been shown and experimentally verified that the number of local optimization steps is lower than logarithm of the number of samples drawn, and thus local optimization does not slow the procedure down. Four different methods of local optimization were tested and the efficiency of method 5 is almost 1. The proposed improvement allows to make precise quantitative statements about the number of samples drawn in RANSAC. The local optimization step applied to selected models produces an algorithm with near perfect agreement with theoretical (i.e. optimal) performance.

References

1. O. Chum and J. Matas. Randomized ransac with T(d,d) test. In *Proceedings of the British Machine Vision Conference*, volume 2, pages 448–457, 2002.
2. J. Clarke, S. Carlsson, and A. Zisserman. Detecting and tracking linear features efficiently. In *Proc. 7th BMVC*, pages 415–424, 1996.
3. M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *CACM*, 24(6):381–395, June 1981.
4. R. Hartley. In defence of the 8-point algorithm. In *ICCV95*, pages 1064–1070, 1995.
5. A. Leonardis and H. Bischof. Robust recognition using eigenimages. *Computer Vision and Image Understanding: CVIU*, 78(1):99–118, Apr. 2000.
6. J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide baseline stereo from maximally stable extremal regions. In *Proc. of the BMVC*, volume 1, pages 384–393, 2002.
7. P. McLauchlan and A. Jaenicke. Image mosaicing using sequential bundle adjustment. In *Proc. BMVC*, pages 616– 62, 2000.
8. D. Myatt, P. Torr, S. Nasuto, J. Bishop, and R. Craddock. Napsac: High noise, high dimensional robust estimation - it's in the bag. In *BMVC02*, volume 2, pages 458–467, 2002.
9. P. Pritchett and A. Zisserman. Wide baseline stereo matching. In *Proc. International Conference on Computer Vision*, pages 754–760, 1998.
10. F. Schaffalitzky and A. Zisserman. Viewpoint invariant texture matching and wide baseline stereo. In *Proc. 8th ICCVon, Vancouver, Canada*, July 2001.
11. B. Tordoff and D. Murray. Guided sampling and consensus for motion estimation. In *Proc. 7th ECCV, Copenhagen, Denmark*, volume 1, pages 82–96. Springer-Verlag, 2002.
12. P. Torr, A. Zisserman, and S. Maybank. Robust detection of degenerate configurations while estimating the fundamental matrix. *CVIU*, 71(3):312–333, September 1998.
13. P. H. S. Torr. *Outlier Detection and Motion Segmentation*. PhD thesis, Dept. of Engineering Science, University of Oxford, 1995.
14. P. H. S. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78:138–156, 2000.
15. T. Tuytelaars and L. Van Gool. Wide baseline stereo matching based on local, affinely invariant regions. In *Proc. 11th British Machine Vision Conference*, 2000.