

Fixing the Locally Optimized RANSAC

Karel Lebeda
karel@lebeda.sk

Jiri Matas
matas@cmp.felk.cvut.cz

Ondrej Chum
chum@cmp.felk.cvut.cz

Center for Machine Perception,
Czech Technical University,
Faculty of Electrical Engineering,
Department of Cybernetics,
Karlovo namesti 13,
121 35 Prague, Czech Republic

Abstract

The paper revisits the problem of local optimization for RANSAC. Improvements of the LO-RANSAC procedure are proposed: a use of truncated quadratic cost function, an introduction of a limit on the number of inliers used for the least squares computation and several implementation issues are addressed. The implementation is made publicly available.

Extensive experiments demonstrate that the novel algorithm called LO⁺-RANSAC is (1) very stable (almost non-random in nature), (2) very precise in a broad range of conditions, (3) less sensitive to the choice of inlier-outlier threshold and (4) it offers a significantly better starting point for bundle adjustment than the Gold Standard method advocated in the Hartley-Zisserman book.

1 Introduction

One of the attractive properties of RANSAC [5] is, at least with the top-hat (inlier 1, outlier 0) cost function, that it returns an optimal solution with a predefined, user-controllable probability. The theoretical guarantee is based on the assumption that all all-inlier (minimal) samples lead to the optimal solution. It has been observed [6, 7] that the assumption is not valid in practice and that often a significant data-dependent fraction of all-inlier samples does not lead to an acceptable solution.

To address the “not all all-inlier samples are good” problem, Chum *et al.* [8] introduced the LO-RANSAC which applies a local optimization (LO) step to promising hypotheses generated from random minimal samples. Experiments in [8] show that LO-RANSAC is superior to plain RANSAC in terms of accuracy and its probability of obtaining a correct solution is close to the theoretical value derived from the stopping criterion. The LO-RANSAC method is popular, highly cited and has been used in a number of applications [9, 10].

Chum *et al.* [8] stated that the improvements of the accuracy and the probability of obtaining a correct solution may even speed the algorithm up since the increased number of found inliers triggers the stopping criterion earlier. The LO is run only rarely, the number of runs being close to the logarithm of the number of samples.

As the first contribution of the paper we show that the “no extra time” statement is true only for estimation problems with low inlier ratios. For image pairs with a high fraction of inliers, where a small number of random samples is sufficient for finding the solution, the original LO procedure significantly effects the running time, sometimes becoming a dominating factor that may increase the running time by an order of magnitude. To alleviate the problem and reduce the overhead, we modify the iterative least squares by introducing a limit on the number of inliers used in the least squares computation. Nevertheless, the modified LO-RANSAC is slower than plain RANSAC, fortunately mainly for easy datasets where the procedure is very fast anyway. Essentially the result shows that the local optimization is not always a free lunch and that there is a trade-off between estimation quality (accuracy and repeatability) and the computational time. As a second contribution, we introduce a fast version – LO⁺ that has an execution time close to the standard RANSAC and perform close to LO-RANSAC in almost all cases.

The LO procedure is relatively complex, with a high number of parameters. As a third contribution of the paper, we are making public an ultimate description of the method: a C/C++ implementation of the improved LO⁺. The implementation has been extensively experimentally tested and performed well on dozens of geometry estimation problems with the same parameter settings¹. The proposed method is very stable - for many tested geometric problems it returned the identical set of inliers in 10000 out of 10000 test runs. We also show that the proposed algorithm is insensitive to the choice of the *error scale* which defines the inlier-outlier separation. In this context we confirm the slight advantage of the MSAC-like truncated quadratic [22] over the the top-hat, 0-1 loss function. The precision of the LO procedure for both methods is almost identical, but the MSAC-like kernel increases tolerance to the choice of the inlier threshold. Therefore, the proposed LO⁺ differs from the standard LO [6], besides a number of implementation details, by using an inlier limit and the truncated quadratic cost function.

Finally, the accuracy of the proposed LO⁺ method is tested within a standard Bundle adjustment method [23]. Perhaps surprisingly the bundler is rather sensitive to initialization. The LO initialized non-linear optimization is always superior in terms of residual errors to the Gold Standard method advocated by Hartley and Zissermann [10].

The structure of this paper is as follows. We describe LO-RANSAC algorithm [6] in section 2 and lists its parameters. In section 3, procedures to speed the LO up are presented. The speed and precision of the proposed methods are experimentally evaluated in section 4. Finally, section 5 draws the conclusions.

2 LO-RANSAC algorithm

The structure of the RANSAC algorithm is simple. Repeatedly, minimal subsets are randomly selected from the input data (i.e. tentative correspondences in a two-view geometry estimation) and model parameters fitting the sample are computed. Subsequently, in a verification step all inliers to the model are found and the quality of the model parameters evaluated. The model maximising the cost function is returned. The locally optimized RANSAC adds an optimization step after the verification phase, if a *so-far-the-best* model is found. The LO-RANSAC is summarized in Algorithm 2. All symbols used in the algorithm are described in Table 1. The algorithm is written as a for-loop where K denotes the number of cycles. As

¹With the geometric error tolerance a fixed function of the image dimensions

in standard RANSAC, K is a function of the user-defined desired probability η of finding the optimal solution and the number of inliers of the best model.

Algorithm 1 LO-RANSAC [9].

```

1: for  $k = 1 \rightarrow K(|\mathcal{I}^*|, \eta)$  do
2:    $\mathcal{S}_k \leftarrow$  randomly drawn minimal sample
3:    $M_k \leftarrow$  model estimated from sample  $\mathcal{S}_k$ 
4:    $\mathcal{I}_k \leftarrow \text{find\_inliers}(M_k, \theta)$ 
5:   if  $|\mathcal{I}_k| > |\mathcal{I}_s^*|$  then
6:      $M_s^* \leftarrow M_k; \mathcal{I}_s^* \leftarrow \mathcal{I}_k$ 
7:      $M_{LO}, \mathcal{I}_{LO} \leftarrow$  run Local Optimization      (Alg. 2)
8:     if  $|\mathcal{I}_{LO}| > |\mathcal{I}^*|$  then
9:        $M^* \leftarrow M_{LO}; \mathcal{I}^* \leftarrow \mathcal{I}_{LO}$ 
10:    update  $K$ 
11:   end if
12: end if
13: end for
14: return  $M^*$ 

```

Table 1: Notation.

Symbol	Description	Value for EG	Value for HG
s_{is}	size of inner sample	$\min(14, \frac{ \mathcal{I}_{base} }{2})$	$\min(12, \frac{ \mathcal{I}_{base} }{2})$
m_θ	threshold multiplier	$\sqrt{2}$	$\sqrt{2}$
m'_θ	threshold multiplier – LO'	$4 \cdot m_\theta$	m_θ
$reps$	inner sam. repetitions	10	10
$iters$	iterations of LSq	4	4
$iters'$	iterations of LSq – LO'	10	4
Δ_θ	threshold decrement	$\frac{m_\theta \cdot \theta - \theta}{iters - 1}$	$\frac{m_\theta \cdot \theta - \theta}{iters - 1}$
θ	inlier-outlier error threshold (error scale)		
η	user-required probability of finding the optimal solution		
LSq	Least Squares solution, normalized and possibly weighted		
find_inliers	find points with error smaller than θ w.r.t. model M		
\mathcal{S}	random sample		
$M(M^*, M_s^*)$	model (the best found, the best from minimal sample)		
$\mathcal{I}(\mathcal{I}^*, \mathcal{I}_s^*)$	inlier set (the largest found, the largest from minimal sample)		

The LO procedure of Chum *et al.* includes least squares on inliers, iterative least squares with a narrowing threshold and random sampling of non-minimal samples (from inliers of the *so-far-the-best* model); the standard LO is outlined in Algorithms 2 and 3. The procedure is fairly complex with a number of parameters. It is a result of extensive experimentation that showed that all attempts to simplify the algorithm lead to a decrease in stability and accuracy of the algorithm.

It should be noted that the proposed local optimization does not handle degenerate scenes and it must be combined with approaches like DEGENSAC [9] or QDEGSAC [9] that cover this problem.

Algorithm 2 Local Optimization step.**Input:** $M_s^*, m_\theta, reps$

- 1: $M_{m_\theta} \leftarrow$ model estimated by LSq on $find_inliers(M_s^*, m_\theta \cdot \theta)$
- 2: $\mathcal{I}_{base} \leftarrow find_inliers(M_{m_\theta}, \theta)$
- 3: **for** $r = 1 \rightarrow reps$ **do**
- 4: $S_{is} \leftarrow$ sample of size s_{is} randomly drawn from \mathcal{I}_{base}
- 5: $M_{is} \leftarrow$ model estimated from S_{is} by LSq
- 6: $M_r \leftarrow$ Iterative Least Squares ($M_{is}, m_\theta, iters$) (Alg. 3)
- 7: **end for**
- 8: **return** the best of M_s^* , all M_{is} , all M_r , with its inliers

Algorithm 3 Iterative Least Squares.**Input:** $M_{is}, m_\theta, iters$

- 1: $M' \leftarrow$ model estimated by LSq on $find_inliers(M_{is}, \theta)$
- 2: $\theta' \leftarrow m_\theta \cdot \theta$
- 3: **for** $i = 1 \rightarrow iters$ **do**
- 4: $\mathcal{I}' \leftarrow find_inliers(M', \theta')$
- 5: $w' \leftarrow$ computed weights of \mathcal{I}' (depend on model)
- 6: $M' \leftarrow$ model estimated by LSq on \mathcal{I}' weighted by w'
- 7: $\theta' \leftarrow \theta' - \Delta_\theta$
- 8: **end for**
- 9: **return** the best M'

3 Speeding up Local Optimization

The full local optimization step uses the (iterated) linear least squares to improve the quality of the estimated model. For a large number of inliers, this procedure can dominate the execution time, even if executed only once. Therefore, LO-RANSAC is considerably slower than plain RANSAC in scenarios with large number (both relative and absolute) of inliers. Formally, the total running time of LO-RANSAC is

$$t_{tot} = C_R \cdot K + C_{LO} \cdot \lceil \log(K) \rceil \quad (1)$$

where C_{LO} is the average time of LO procedure and C_R is the time of standard RANSAC single hypothesis generation and verification round. For small K , the significant difference of C_{LO} and C_R means that $C_{LO} \cdot \lceil \log(K) \rceil \gg C_R \cdot K$. E.g. for the head pair with 74 inliers (86%) where $K \approx 22$, the standard RANSAC running time $C_R \cdot K$ is only 0.4 ms while a single local optimization (C_{LO}) takes 5.6 ms on average.

We propose to reduce the time consumption by introducing a limit on the number of correspondences that participate in estimation of model M' parameters (Alg. 3, line 6). If the number of inliers exceeds this limit, a subset (limit-sized) is randomly chosen and used for the estimation. In our experiments, the limit was empirically set to $7 \times$ minimal sample size. This approach reduces the cost of the least squares estimation to a constant independent of the number of inliers. The experiments show that the precision of the estimated model is not affected and even improved for a non-negligible number of epipolar geometry estimation tasks.

For fast applications, where no challenging camera motion or illumination changes are



Figure 1: The dataset for epipolar geometry estimation




Figure 2: The dataset for homography estimation

expected, a lightweight version (LO') is proposed. Instead of estimating models from non-minimal samples followed by iterative least squares, only a single iterative least squares are applied on each *so-far-the-best* model. The limit on the number of inliers used in each step of the least squares is applied. The experimental evaluation shows significant reduction in the execution time while sacrificing only little accuracy, especially for “easy” image pairs.

4 Experimental results

The performance of different RANSAC variants was evaluated on a collection of 16 image pairs (Figure 1) for epipolar geometry and 16 pairs (Figure 2) for homography estimation. These image pairs were previously used for evaluation in a number of publications ([4, 6, 7, 12, 15, 17, 18, 20, 24, 25]). The datasets are available at <http://cmp.felk.cvut.cz/data/geometry2view/>. In the paper, results on a subset (Figure 3) of those image pairs are reported. For full evaluation see the technical report ([1]).

What is measured. In standard RANSAC, the number of data points consistent with the estimated model (inliers) is optimized. Such a measure is a good indicator how well the image pair is matching ([9]). For some applications such as 3D reconstruction, the precision of the estimated model is of high importance. To compare the precision, we record the error of the estimated model on manually annotated ground truth correspondences. These correspondences are *not included* in the estimation process. Note that such a measure includes the error in the manual selection as well as error induced by deviation of the imaging process from a pin-hole camera model (such as radial distortion). Statistically, smaller error indi-



Figure 3: Scenes with results presented in the paper (for more scenes and further experiments see [10])

cates better precision, however, small differences in small values are insignificant. Finally, the time complexity was evaluated by both number of samples drawn by RANSAC, and the execution duration in seconds (all the tests were performed on 2 GHz dual core machine with 2 GB RAM at 667 MHz and 2 MB L2 cache).

Measurements were averaged over multiple runs with required confidence of success 95 %. For a fair comparison, the seeds of the pseudo-random generator were fixed so that different methods draw the same samples. Tentative correspondences were obtained by matching SIFT descriptors [13] of MSER’s [15]. In the supplementary material, experiments using Hessian Affine detector [16] are presented. The results on Hessian Affine features are even more favourable for the LO methods because of lower inlier ratios.

4.1 Sensitivity of cost functions to the choice of inlier error scale

The sensitivity of RANSAC with a hard decision on inliers and outliers (top-hat cost function) and MSAC (truncated parabolic cost function) on the expected noise level in inlier localization was measured. We model the inlier noise as a zero mean normal distribution with standard deviation dependent on the size of the image as $\sigma_{scaled} = \sigma \cdot sizeFactor$. Here, $sizeFactor$ is the ratio of the larger dimension of the image and constant 768. The inlier-outlier threshold θ for RANSAC was set as 95% percentile of the χ^2 distribution (with one DoF for epipolar geometry and two DoF’s in the case of homography) with standard deviation σ_{scaled} .

The error on the ground truth points as a function of σ is plotted in Figure 4 (every measured point is a mean from 1000 runs). The dashed lines show the error when using the MSAC-like cost function widened by factor $3/2$ to have an equivalent kernel as the top-hat function (the same area). The solid lines correspond to a top-hat cost function with the quadratic cost function used as a tie-breaker when comparing two models with an equal top-hat score (which happens often [27]). From the graph in Fig. 4, it can be seen that both cost functions in the vicinity of the optimal threshold output similar accuracy of the resulting geometries (the truncated quadratic scoring yields slightly better results). For larger thresholds this difference becomes more significant – the truncated quadratic cost function is more robust to the selection of the error scale. It provides a range of one order of magnitude of usable thresholds for most of the image pairs, making it easier to select a suitable one for a diverse set of data.




Figure 4: The dependence of the geometric accuracy on the cost function and the inlier-outlier threshold

4.2 Two-view geometry estimation evaluation

We used the truncated quadratic cost function in the following experiments as it outperforms the top-hat function. The parameter σ was empirically set to 0.3 (illustrated as vertical cyan lines in Fig. 4). Table 2 summarizes the evaluation results. Measurements were averaged over 10 000 runs with an exception of experiments with the bundle adjustment, where only 100 runs were performed. To assure maximal possible precision of time measurements, we measured overall duration of all RANSAC runs. Thus there is no information about the variance of consumed time. The first two columns give reference results of standard methods: plain MSAC [2], and MSAC followed by linear least squares on all inliers [10].

Then we show lightweight LO⁺ (with inlier-subset speed-up) and improved full version of local optimization. In the next column the effect of the proposed speed-up action is shown (LO⁺). The last two columns report results after non-linear bundle adjustment.

Stability of the results. RANSAC, as a randomized algorithm, returns different outputs each time it is executed, which is often considered as a drawback. The local optimization significantly reduces the variance in both the number of detected inliers and the accuracy. Here the stability of a model returned with usage of a truncated quadratic cost function should be emphasized. In particular, e.g. for the *Boston* image pair, LO⁺-RANSAC with MSAC-like gain function returned (with precision sufficient for all practical purposes) **the same resulting homography for all 10,000 runs**.

Iterative LSq on bounded number of inliers. The experiments show that LO⁺-RANSAC not only improves the speed, but often has also a positive effect on the accuracy. We explain such a behaviour as avoiding sticking in a local minimum by randomizing the set of points used in each iteration of the least squares. This technique is also used in the LO⁺ procedure. Compared to full LO-RANSAC, LO⁺ is up to six times faster, mostly reducing the accuracy only marginally.

Table 2: The geometric accuracy, stability and speed of RANSAC variants

		MSAC \square	MSAC+LSq \square	LO ⁺	LO	LO ⁺	MSAC+LSq+BA	LO+BA
corr	Inliers	62.7± 4.4	66.0± 4.2	69.8± 2.8	73.1± 1.6	73.3± 1.8	67.4± 4.2	73.2± 1.5
	Error (px)	0.48± 0.33	0.37± 0.33	0.31±0.12	0.18±0.11	0.18±0.10	0.34± 0.25	0.16±0.04
	Time (ms)	1.1	1.3	2.1	6.5	6.3	2459.5	2046.8
	Samples	61.0± 25.1	61.0± 25.1	49.7±16.1	49.5±15.9	49.5±15.9	63.7± 27.0	52.3±20.7
head	Inliers	66.9± 4.1	71.9± 2.7	73.7± 0.9	73.9± 0.6	74.0± 0.6	72.9± 2.0	74.0± 0.2
	Error (px)	0.78± 0.52	0.40± 0.19	0.30±0.03	0.31±0.03	0.31±0.03	0.38± 0.15	0.35±0.02
	Time (ms)	0.4	0.6	1.6	6.0	5.8	812.4	685.8
	Samples	21.8± 10.1	21.8± 10.1	21.7± 9.8	21.7± 9.8	21.7± 9.8	21.6± 9.9	21.6± 9.9
Kyoto	Inliers	295.2± 16.5	311.4± 15.3	325.1± 9.2	333.5± 6.7	330.7± 5.7	313.7± 16.7	332.1± 8.0
	Error (px)	2.25± 1.28	1.64± 1.14	1.07±0.54	0.81±0.32	0.78±0.23	1.47± 0.97	0.78±0.33
	Time (ms)	2.4	2.7	3.6	12.2	9.8	18499.7	12006.1
	Samples	65.4± 26.0	65.4± 26.0	49.6±12.6	49.2±12.1	49.1±12.1	66.8± 27.5	51.3±14.9
wash	Inliers	45.7± 3.5	50.1± 1.7	51.7± 0.5	51.3± 0.4	51.4± 0.5	50.6± 1.0	51.0± 0.2
	Error (px)	1.04± 0.61	0.39± 0.17	0.28±0.02	0.27±0.04	0.27±0.03	0.32± 0.13	0.26±0.03
	Time (ms)	0.3	0.4	1.4	5.4	5.4	132.2	107.4
	Samples	16.7± 9.8	16.7± 9.8	16.7± 9.7	16.7± 9.7	16.7± 9.7	15.8± 8.9	15.8± 8.9
Boston	Inliers	277.3± 21.5	303.0± 5.4	305.0± 0.1	305.0± 0.0	305.0± 0.0	305.0± 0.2	305.0± 0.0
	Error (px)	1.78± 1.01	0.72± 0.20	0.60±0.08	0.66±0.00	0.66±0.00	0.67± 0.03	0.66±0.00
	Time (ms)	1.1	1.3	1.9	16.0	11.0	82.6	26.0
	Samples	12.8± 5.8	12.8± 5.8	12.8± 5.8	12.8± 5.8	12.8± 5.8	12.3± 5.7	12.3± 5.7
Brussels	Inliers	328.7± 32.4	371.4± 18.2	387.9± 4.4	390.6± 1.3	390.5± 2.1	379.0± 8.6	387.1± 0.6
	Error (px)	3.65± 0.92	2.59± 0.50	2.25±0.20	2.88±0.05	2.86±0.08	3.15± 0.21	2.97±0.01
	Time (ms)	2.3	2.6	3.3	20.7	14.1	116.6	112.4
	Samples	21.0± 9.4	21.0± 9.4	20.9± 9.2	20.9± 9.2	20.9± 9.2	21.8± 9.6	21.8± 9.5
Eiffel	Inliers	60.9± 4.1	64.4± 3.2	66.0± 1.7	66.8± 1.1	66.7± 1.1	65.3± 2.6	66.5± 0.8
	Error (px)	1.23± 0.57	0.92± 0.44	0.82±0.28	0.88±0.16	0.88±0.15	0.91± 0.35	0.78±0.17
	Time (ms)	6.8	6.8	5.5	19.6	18.6	28.5	45.1
	Samples	438.9±155.3	438.9±155.3	273.2±40.7	254.5±18.6	254.4±17.2	444.8± 168.3	255.7±22.6
WhiteB.	Inliers	161.1± 13.2	171.6± 7.7	173.7± 1.8	174.0± 0.0	174.0± 0.0	172.7± 6.5	174.0± 0.0
	Error (px)	1.48± 0.49	1.09± 0.19	1.02±0.06	1.08±0.00	1.06±0.01	1.08± 0.12	1.05±0.00
	Time (ms)	0.7	0.8	1.3	9.7	7.6	61.2	53.1
	Samples	11.7± 5.8	11.7± 5.8	11.7± 5.8	11.7± 5.8	11.7± 5.8	10.2± 4.3	10.2± 4.3

Bundle adjustment (BA). To further minimize the reprojection error, non-linear iterative optimization is employed \square , using the output of the robust estimator as an initialization. The last two columns of Table 2 report results after application of the non-linear optimization using publicly available bundle adjustment of Lourakis \square . Two initializations are compared: the Gold Standard method \square (output of MSAC refined by one linear least squares) and the output of the fully locally optimized MSAC. The results consistently show, that LO-MSAC provides better starting point for the BA: lower reprojection error is acquired in shorter time. Note that the time saved in the BA is orders of magnitude higher than the overhead related to the local optimization.

4.3 LO overhead and the inlier ratio

Figure 5 shows the behaviour of RANSAC using truncated quadratic cost function with and without local optimization on the `booksh` (\square) image pair (the percentage of inliers is controlled by the threshold of the second closest match during the matching process). As expected, LO⁺ outperforms MSAC in both the number of inliers and the geometric error. In the cases of low inlier ratios, where the numbers of samples drawn by every RANSAC are high, LO⁺ is faster than MSAC. However, for high inlier ratios the curves cross.

4.4 Implementations details and issues

Since for some pairs each call of the LO has impact on the overall speed, it is efficient not to call the LO during first K_{start} iterations (set to 50 in our implementation). In such a case it is

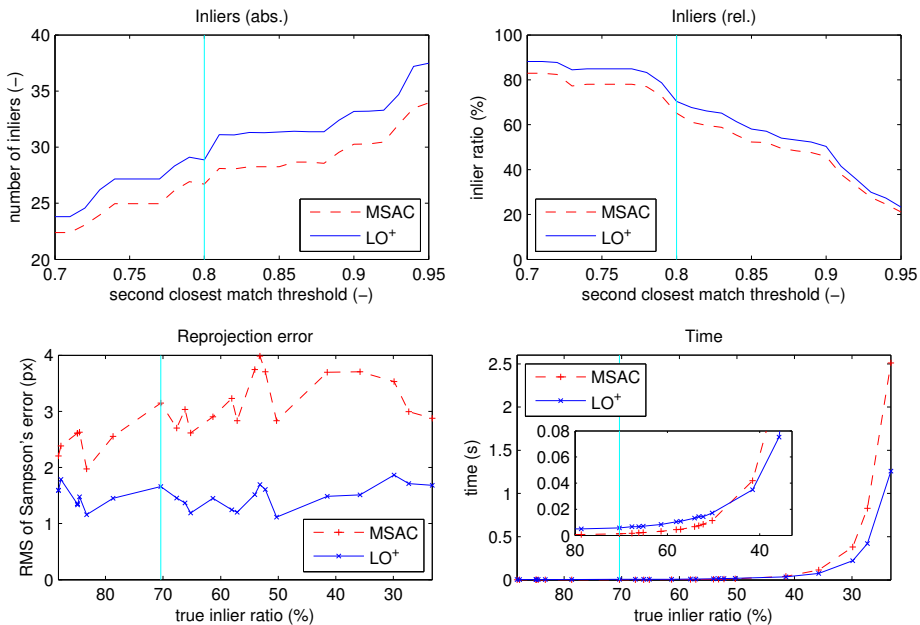


Figure 5: The dependence of the precision and speed on the percentage of inliers which is controlled by the threshold of the second closest match.

necessary to ensure that LO is executed at least once, after K_{start} iterations, or at the end of RANSAC (as some scenes may require less than K_{start} samples). In this case, the LO is run for the best model found in the first K_{start} iteration.

We observed that the choice of numerical methods is critical for getting fast and precise model estimates. While SVD is a convenient and stable way to compute a least square solution of a system of linear equations, it is significantly faster to use eigen-decomposition of the covariance matrix, especially for large systems of equations.

In our experiments, we have encountered a significant drop in performance for epipolar geometry estimation caused by the instability of SVD decomposition when using CCMATH library [2]. The instability was observed during fundamental matrix singularization. The final implementation uses the LAPACK library [1] that does not suffer by such an instability.

5 Conclusion

Several technical improvements of the LO-RANSAC were proposed. Our extensive evaluation shows that: (1) the LO^+ -RANSAC offers a stable robust estimation despite its randomized nature, (2) limiting the number of inliers included in the (iterative) least squares significantly reduces the execution time and often even improves the precision, (3) the speed of the lightweight LO^+ is comparable to plain RANSAC even for easy problems with very high inlier ratios, and that (4) LO^+ -RANSAC offers a significantly better starting point for bundle adjustment than the Gold Standard [1]. We addressed a number of implementation issues and extensively tested the method. An implementation of the proposed LO methods is available at <http://cmp.felk.cvut.cz/software/LO-RANSAC/>.

References

- [1] LAPACK: Linear algebra package.
<http://www.netlib.org/lapack/>.
- [2] D. A. Atkinson. CCMATH: Mathematics software library.
<http://freecode.com/projects/ccmath>.
- [3] J. Cech, J. Matas, and M. Perdoch. Efficient sequential correspondence selection by cosegmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1568–1581, 2009.
- [4] S. Choi, T. Kim, and W. Yu. Performance evaluation of RANSAC family. In *Proc. of BMVC*, pages 81.1–81.12, 2009.
- [5] O. Chum and J. Matas. Matching with PROSAC – progressive sample consensus. In *Proc. of the Conf. on CVPR*, pages 220–226, 2005.
- [6] O. Chum, J. Matas, and J. Kittler. Locally optimized RANSAC. In *DAGM-Symposium*, pages 236–243, 2003.
- [7] O. Chum, T. Werner, and J. Matas. Two-view geometry estimation unaffected by a dominant plane. In *Proc. of the Conf. on CVPR*, pages 772–779, 2005.
- [8] 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 ACM*, 24(6):381–395, 1981.
- [9] J.-M. Frahm and M. Pollefeys. RANSAC for (Quasi-)degenerate data (QDEGSAC). In *Proc. of the Conf. on CVPR*, pages 453–460, 2006.
- [10] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.
- [11] K. Lebeda, J. Matas, and O. Chum. Fixing the Locally Optimized RANSAC. Research Report CTU–CMP–2012–17, Center for Machine Perception, Czech Technical University, Prague, Czech Republic, 2012.
http://cmp.felk.cvut.cz/software/LO-RANSAC/Lebeda-2012-Fixing_LORANSAC-tr.pdf.
- [12] M. I. A. Lourakis and A. A. Argyros. SBA: A Software Package for Generic Sparse Bundle Adjustment. *ACM Trans. Math. Software*, 36(1):1–30, 2009.
- [13] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [14] D. Martinec and T. Pajdla. 3D reconstruction by fitting low-rank matrices with missing data. In *Proc. of the Conf. on CVPR*, pages 198–205, 2005.
- [15] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide baseline stereo from maximally stable extremal regions. In *Proc. of BMVC*, pages 384–396, 2002.
- [16] K. Mikolajczyk and C. Schmid. Scale and affine invariant interest point detectors. *International Journal of Computer Vision*, 60(1):63–86, 2004.

- [17] J.-M. Morel and G. Yu. ASIFT: A new framework for fully affine invariant image comparison. *SIAM Journal on Imaging Sciences*, 2(2):438–469, 2009.
- [18] M. Perdoch, J. Matas, and O. Chum. Epipolar geometry from two correspondences. In *Proc. of the ICPR*, pages 215–220, 2006.
- [19] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *Proc. of the Conf. on CVPR*, 2007.
- [20] M. Pollefeys, R. Koch, M. Vergauwen, and L. Van Gool. Automated reconstruction of 3D scenes from sequences of images. *ISPRS Journal Of Photogrammetry And Remote Sensing*, 55(4):251–267, 2000.
- [21] B. Tordoff and D. W. Murray. Guided sampling and consensus for motion estimation. In *Proc. of the ECCV*, pages 82–98, 2002.
- [22] P. H. S. Torr and A. Zisserman. Robust computation and parametrization of multiple view relations. In *Proc. of the ICCV*, pages 727–732, 1998.
- [23] P. Turcot and D.G. Lowe. Better matching with fewer features: The selection of useful features in large database recognition problems. In *International Conference on Computer Vision Workshops*, pages 2109–2116, 2009.
- [24] T. Tuytelaars and L. Van Gool. Wide baseline stereo matching based on local, affinely invariant regions. In *Proc. of BMVC*, pages 412–422, 2000.
- [25] G. Yang, C.V. Stewart, M. Sofka, and C.-L. Tsai. Registration of challenging image pairs: Initialization, estimation, and decision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):1973–1989, 2007.