21^{st} *Computer Vision Winter Workshop*
Luka Čehovin, Rok Mandeljc, Vitomir Štruc (eds.)
Rimske Toplice, Slovenia, February 3–5, 2016

Fast L_1-based RANSAC for homography estimation

Jonáš Šerých, Jiří Matas, Ondřej Drbohlav
Czech Technical University in Prague, Faculty of Electrical Engineering,
Department of Cybernetics, Center for Machine Perception,
Technická 2, 166 27 Praha 6, Czech Republic
{serycjon,matas,drbohlav}@fel.cvut.cz

Abstract. *We revisit the problem of local optimization (LO) in RANSAC for homography estimation. The standard state-of-the-art LO-RANSAC improves the plain version's accuracy and stability, but it may be computationally demanding, it is complex to implement and requires setting multiple parameters. We show that employing L_1 minimization instead of the standard LO step of LO-RANSAC leads to results with similar precision. At the same time, the proposed L_1 minimization is significantly faster than the standard LO step of [8], it is easy to implement and it has only a few of parameters which all have intuitive interpretation. On the negative side, the L_1 minimization does not achieve the robustness of the standard LO step, its probability of failure is higher.*

1. Introduction

RANSAC [3] is a robust model fitting algorithm that is the standard method used for two-view geometry estimation [5]. The plain version of RANSAC proceeds as follows: (i) randomly sample the minimum number of points required to calculate model parameters, (ii) compute the cardinality of the set consistent with that model, i.e. the number of inliers, and (iii) terminate if the probability that a better model than the one best so far will be found falls under a predefined threshold. The precision of the model returned by the algorithm is typically improved by least square fitting of the inliers of the best mode.

It has been observed [11] that the termination criterion (iii) stops the process later than expected given the recovered percentage of inliners. The discrepancy is due to a generally incorrect (overoptimistic) assumption that every minimal sample of inliers generates a "good" model, i.e. a model that will be consistent with all correct correspondences.

The problem was first addressed in a paper by Chum et al [2] who proposed an additional RANSAC step, the so called local optimization (LO). The LO step is employed whenever a new candidate model M is the best one so far found in the RANSAC loop, i.e. it has more inliers than any of the models estimated from the random minimal samples evaluated so far. Chum et al [2] proved that with the strategy, the LO step is run only $\log(k)$ times, where k is the number of random models tested.

The local optimization step[2] performs various heuristic procedures with the objective of increasing the accuracy of M, such as generating hypotheses by resampling the inliers of M and performing iterative least square estimation combined with scheduled inlier threshold changes. The standard implementation of RANSAC with the local optimization step, found in the commonly used publicly available code [10], is a combination of the above-mentioned heuristic procedures.

The choice and parameter settings of local optimization methods influence the speed and accuracy of the algorithm. In the state-of-the-art version [8], the LO step executes a complex procedure which involves repeated sampling from inliers of M and repeated iterative least squares minimisation. As the sampling is involved, it is stochastic[1]. Due to both repeated sampling and iterative least squares, it is so computationally demanding, in comparison with RANSAC steps (i) and (ii), that despite being executed only rarely, the LO step significantly influences the overall running time.

In this paper, we propose to replace the complex LO procedure of Lebeda et al. by minimization of the

[1]Since the outer loop of RANSAC is stochastic, the inner sampling does not change the character of the algorithm.

procedure STANDARD LO
  Input: M (model estimated by LSq), I (inliers)
  for $r = 1 \rightarrow reps$ **do**
    sample S drawn from inliers
    model M is estimated from S
    iterative least squares on M
  end for
  return best model
end procedure

procedure L_1-BASED LO
  Input: M (minimum sample model), I (inliers)
  while stopping condition not met **do**
    $M \leftarrow$ model estimated from inliers by IRLS optimization
    $I \leftarrow$ inliers to M
  end while
  return M
end procedure

Table 1: Comparison of the standard and proposed local optimization procedures in RANSAC – left and right columns, respectively. IRLS stands for interated re-weighted least squares. Note that the standard LO method includes several rounds of IRLS s which are themselves computationally demanding (for details, see [8]).The iteration stops if either the change in the cost function is below 10^{-3} or the maximum number of iterations is reached (set to 5).

sum of L_1 norms of the residuals, ie. the algebraic errors of the model on individual points. The minimizer of the L_1 norm, also known as geometric median, is robust to a modest contamination by outliers. This means that RANSAC becomes less sensitive to the inlier-outlier threshold. The threshold, a critical parameter of RANSAC, can be set more loosely and thus cover a wide range of problems. Moreover, the L_1-based procedure need not include least square estimation with multiple thresholds, thus saving time.

In practice we replace the L_1 norm by the Huber robust kernel response to the inlier algebraic errors. The Huber cost function is defined in Eq. 6. The Huber kernel is differentiable and convex and the global minimum of the cost function can be found by gradient descent. The gradient minimization alternates with the inlier-outlier selection process. The alternating minimisation can be seen as *local optimization* of the truncated Huber kernel. The procedure has only a small number of parameters that have intuitive meaning, it is simple, and deterministic.

We show that the minimization produces errors which are comparable to the standard LO-RANSAC, while being computationally much less expensive – of an order of magnitude in our experiments compared to the standard local optimization.

2. Method

The difference of the standard and proposed LO method is presented in Table 1. The L_1 minimization is carried out by iterated reweighted least squares (IRLS). The particular instantiation of IRLS is knows as the generalized Weiszfeld algorithm [1]. Weiszfeld proved that the geometric mean minimiation by IRLS requires solving repeated least squares problems where each data point is weighted by the reciprocal of its residual to the current estimate of the model. The algorithm has to be modified to avoid singularities when some point is exactly consistent with the model, i.e. it has a zero residual. To avoid the problem, we replace L_1 minimization with Huber kernel minimization. In the implementation, it only means that points with small residuals are not scaled.

First, the necessary notation is introduced. The L_2^2 norm (for a vector $\mathbf{r} \in \mathbb{R}^D$) is defined as:

$$\|\mathbf{r}\|_2^2 = \sum_{j=1}^{D} |r_j|^2, \tag{1}$$

the L_2^1 norm (for $\mathbf{r} \in \mathbb{R}^D$) as

$$\|\mathbf{r}\|_2^1 = \sqrt{\sum_{j=1}^{D} |r_j|^2} \tag{2}$$

2.1. Homography estimation by algebraic error minimization in L_2^2 and L_2^1 norms

Let the number of correspondences be M. The data matrix $\mathbf{Z}$ is computed from correspondences by a standard procedure ([5]): Let (x, y) and (x', y') be the correspondence pair. It generates two rows into the data matrix $\mathbf{Z}$:

$$\begin{bmatrix} x & y & 1 & 0 & 0 & 0 & -x'x & -x'y & -x' \\ 0 & 0 & 0 & x & y & 1 & -y'x & -y'y & -y' \end{bmatrix}. \tag{3}$$

Let $\mathbf{z}^{(i)}$ denote the two rows generated by i-th correspondence. The homography $\mathbf{h}$ is estimated from $\mathbf{Z}$ by one of the following optimizations:

The L_2^2 optimization

$$\mathbf{h} = \operatorname*{argmin}_{\hat{\mathbf{h}}} \sum_{i=1}^{M} \|\mathbf{z}^{(i)}\hat{\mathbf{h}}\|_2^2 \,, \quad (\text{subj. to } \hat{\mathbf{h}}^\top\hat{\mathbf{h}} = 1) \tag{4}$$

The minimization is solved by computing the spectral decomposition of $\mathbf{Z}^\top\mathbf{Z}$ and taking the eigenvector corresponding to the smallest eigenvalue. The algorithm has the following properties: it is fast, but not robust with a breakdown point of zero [7] – in general a single outlier can make $\mathbf{h}$ arbitrarily wrong[2].

The L^1 **optimization**, defined as

$$\mathbf{h} = \operatorname*{argmin}_{\hat{\mathbf{h}}} \sum_{i=1}^{M} \|\mathbf{z}^{(i)}\hat{\mathbf{h}}\|^1 \,, \quad (\text{subj. to } \hat{\mathbf{h}}^\top\hat{\mathbf{h}} = 1) \tag{5}$$

is robust and can be solved by the generalized Weiszfeld algorithm, an instance of IRLS. Instead of modifying the algorithm to take of the technical problems associated with the Weiszfeld algorithm arising if one of the residuals vanishes, we instead optimize the response to the Huber kernel.

Huber optimization is defined as

$$\mathbf{h} = \operatorname*{argmin}_{\hat{\mathbf{h}}} \sum_{i=1}^{M} \begin{cases} \frac{1}{2}\|\mathbf{r}^{(i)}\|_2^2 & : \|\mathbf{r}^{(i)}\|_2^1 \le k \\ k(\|\mathbf{r}^{(i)}\|_2^1 - \frac{k}{2}) & : \|\mathbf{r}^{(i)}\|_2^1 \ge k \end{cases}$$
$$(\text{subj. to } \hat{\mathbf{h}}^\top\hat{\mathbf{h}} = 1 \text{ and } \mathbf{r}^{(i)} = \mathbf{z}^{(i)}\hat{\mathbf{h}}) \tag{6}$$

The minimization is carried out by a slightly modified Weiszfeld algorithm ([12]), an iterative reweighted least squares method:

procedure IRLS OPTIMIZATION
    Initialize $\mathbf{h}$ as the estimate obtained from the minimal sample $\mathbf{h}$
    while stopping condition not met **do**
        Compute the geometric error $\mathbf{r}^{(i)}$:

$$\mathbf{r}^{(i)} = \|\mathbf{z}^{(i)}\mathbf{h}\|_2^1 \quad (\forall i = 1, 2, ..., M) \tag{7}$$

        Reweight $\mathbf{Z}$:

$$\mathbf{z}^{(i)} \leftarrow \sqrt{\mathbf{w}^{(i)}}\mathbf{z}^{(i)} \quad (\forall i = 1, 2, ..., M) \tag{8}$$

        Recompute $\mathbf{h}$ using L_2^2 optimization (4)
    end while
end procedure

The iteration stops when

$$\sum_i \mathbf{r}_t^{(i)} - \sum_i \mathbf{r}_{(t+1)}^{(i)} \approx 0$$

, i.e. if the value of the cost function does not change between consecutive iterations or after 5 iterations are completed. The second condition reflects the empirical observation that most of the time, the IRLS algorithm converges after 3 iterations and it is used only as a guarantee against an infinite loop.

In the case of L_2^1 optimization, the weight $\mathbf{w}^{(i)}$ is set to $1/\left(\|\mathbf{r}^{(i)}\|_2^1 + \delta\right)$. A small constant δ is used to avoid the problem of dividing by zero when the residuals vanish.

The L^1 optimization proposed above introduces additional parameter δ in order to deal with the division by zero, but its interpretation is not clear. Using the Huber cost function instead of the L^1 norm avoids the numerical issue. The weight $\mathbf{w}^{(i)}$ is set as follows ([13]).

$$\mathbf{w}^{(i)} = \begin{cases} 1 & : \|\mathbf{r}^{(i)}\|_2^1 \le k \\ k/\|\mathbf{r}^{(i)}\|_2^1 & : \|\mathbf{r}^{(i)}\|_2^1 \ge k \end{cases}$$

The additional Huber parameter k can be intuitively seen as a smoothing factor between L_2^2 and L_2^1 norms or, alternatively, like a lower bound on the inlier threshold.

The motivation for using this optimization is its robustness. It is closely related to geometric median computation and the formulation is convex.It is a well known property of median that it is robust to outliers for up to 50% contamination of samples by the outliers. The property makes the procedure non sensitive to the choice of the inlier-outlier threshold of the "outer" RANSAC loop.

3. Experiments.

We compared the standard RANSAC, the state-of-the-art LO-RANSAC and the proposed L_1-based RANSAC on a dataset consisting of 42 image pairs, including selected images from the ZuBuD dataset [4], images from Lebeda's homog dataset [9] used for evaluation of the LO-RANSAC, and images from the symbench dataset [6]. The Hessian Affine feature detector with SIFT descriptor was used for obtaining the tentative correspondences.

[2] In RANSAC, the error on a single point is bounded by the inlier threshold. In practice, points close to the the inlier-outlier boundary make the outcome of standard RANSAC unstable.

Image		Qty↓	RANSAC			LO-RANSAC			L_1-based RANSAC		
05	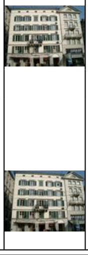	I	**953.2**	±0.9	(950-956)	**953.0**	±0.0	(953-953)	**953.0**	±0.1	(952-953)
		LO time $_{(\mu s)}$	**0.0**	±0.0	(0-0)	**29158.8**	±3383.8	(27499-42497)	**3934.6**	±1035.1	(1901-6479)
		I (%)	**76.9**	±0.1	(77-77)	**76.9**	±0.0	(77-77)	**76.9**	±0.0	(77-77)
		Samp	**11.8**	±5.7	(7-35)	**11.8**	±5.7	(7-35)	**7.5**	±1.9	(7-19)
		Time$_{(ms)}$	**6.1**	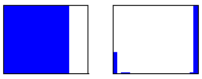		**35.5**	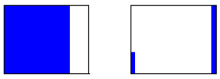		**13.7**	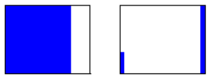	
		Error	**0.74**	±0.05	(0.6-0.9)	**0.72**	±0.00	(0.7-0.7)	**0.73**	±0.01	(0.7-0.8)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**2.2**	±0.9	(1-5)
adam	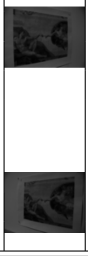	I	**250.8**	±1.1	(244-252)	**251.0**	±0.0	(251-251)	**251.0**	±0.0	(251-251)
		LO time $_{(\mu s)}$	**0.0**	±0.0	(0-0)	**10922.9**	±1797.6	(8737-15292)	**1318.5**	±214.1	(917-1917)
		I (%)	**97.6**	±0.4	(95-98)	**97.7**	±0.0	(98-98)	**97.7**	±0.0	(98-98)
		Samp	**5.0**	±2.6	(2-14)	**5.0**	±2.6	(2-14)	**2.0**	±0.3	(2-4)
		Time$_{(ms)}$	**2.3**	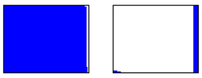		**14.4**	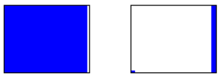		**4.4**	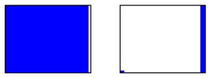	
		Error	**1.15**	±0.45	(0.4-2.8)	**0.77**	±0.05	(0.6-0.8)	**0.79**	±0.02	(0.6-0.8)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**1.4**	±0.5	(1-2)
boat	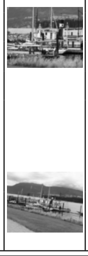	I	**328.4**	±0.5	(328-330)	**328.0**	±0.2	(328-329)	**328.0**	±0.0	(328-328)
		LO time $_{(\mu s)}$	**0.0**	±0.0	(0-0)	**13874.9**	±2006.0	(11071-16489)	**1738.3**	±323.7	(917-2428)
		I (%)	**86.2**	±0.1	(86-87)	**86.1**	±0.1	(86-86)	**86.1**	±0.0	(86-86)
		Samp	**6.2**	±2.5	(4-15)	**6.2**	±2.5	(4-15)	**4.1**	±0.4	(4-7)
		Time$_{(ms)}$	**2.6**	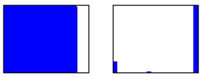		**17.8**	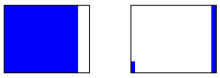		**5.9**	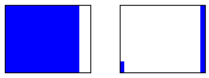	
		Error	**1.30**	±0.14	(1.1-2.1)	**1.23**	±0.01	(1.2-1.2)	**1.24**	±0.00	(1.2-1.2)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**1.7**	±0.7	(1-3)
Brussels	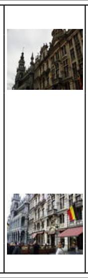	I	**450.0**	±3.5	(428-451)	**451.0**	±0.0	(451-451)	**451.0**	±0.0	(451-451)
		LO time $_{(\mu s)}$	**0.0**	±0.0	(0-0)	**16342.6**	±2310.4	(13755-19648)	**2094.9**	±347.1	(1090-3084)
		I (%)	**87.2**	±0.7	(83-87)	**87.4**	±0.0	(87-87)	**87.4**	±0.0	(87-87)
		Samp	**8.3**	±4.2	(4-22)	**8.3**	±4.2	(4-22)	**4.1**	±0.3	(4-6)
		Time$_{(ms)}$	**3.4**	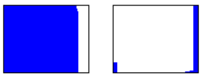		**20.6**	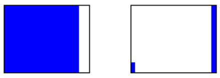		**6.6**	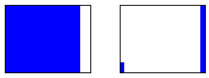	
		Error	**1.39**	±0.37	(1.1-3.3)	**1.24**	±0.00	(1.2-1.2)	**1.24**	±0.00	(1.2-1.3)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**1.8**	±0.7	(1-3)
graf	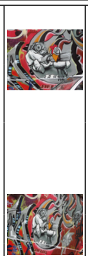	I	**840.1**	±9.8	(808-848)	**846.2**	±0.4	(846-847)	**846.0**	±0.0	(846-846)
		LO time $_{(\mu s)}$	**0.0**	±0.0	(0-0)	**24032.7**	±2219.6	(21845-29919)	**4274.4**	±834.5	(1794-6007)
		I (%)	**89.9**	±1.1	(87-91)	**90.6**	±0.0	(91-91)	**90.6**	±0.0	(91-91)
		Samp	**7.3**	±3.5	(3-20)	**7.3**	±3.5	(3-20)	**3.2**	±0.7	(3-8)
		Time$_{(ms)}$	**4.8**	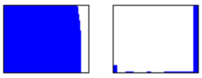		**29.5**	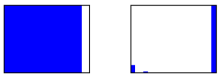		**11.9**	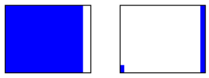	
		Error	**1.69**	±0.22	(1.4-2.7)	**1.45**	±0.00	(1.4-1.5)	**1.45**	±0.01	(1.4-1.6)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**1.7**	±0.7	(1-4)
sym_notredame13		I	**89.6**	±2.4	(77-93)	**91.0**	±0.2	(91-92)	**91.0**	±0.2	(90-92)
		LO time $_{(\mu s)}$	**0.0**	±0.0	(0-0)	**8090.3**	±1025.8	(7196-10973)	**707.8**	±131.0	(437-1009)
		I (%)	**48.4**	±1.3	(42-50)	**49.2**	±0.1	(49-50)	**49.2**	±0.1	(49-50)
		Samp	**110.6**	±53.7	(45-257)	**54.0**	±4.0	(45-67)	**46.1**	±14.8	(37-123)
		Time$_{(ms)}$	**4.2**	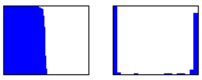		**11.7**	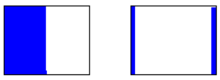		**5.9**	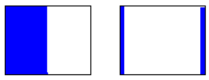	
		Error	**1.81**	±0.62	(1.1-4.7)	**1.13**	±0.02	(1.1-1.2)	**1.15**	±0.09	(1.1-1.7)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**2.9**	±1.2	(1-7)

Table 2: Results on six pairs representing well the whole dataset with the exception of cases in Tab.4. The number of inliers found (*I*), the inlier ratio *I(%)*, the LO step time (LO time), the number of RANSAC samples (*Samp*), CPU time (*time*), the mean error on ground truth correspondences (*Error*) and the number of local optimizations (*LO*). Values in bold are means over 100 runs. The ± entries are standard deviations, minimum and maximum are shown in parentheses. The blue plots represent the stability of each algorithm over 100 runs. The left one represents a probability of a tentative correspondence to be an inlier (probability on the vertical axis, correspondence index on the horizontal axis). The correspondences were sorted so that the plot is non-increasing. In the ideal case, the plot should look like a rectangle. Any other shape indicates that some of the tentative correspondences were not classified as inliers/outliers consistently over the 100 runs. The second plot is a histogram of the first plot.

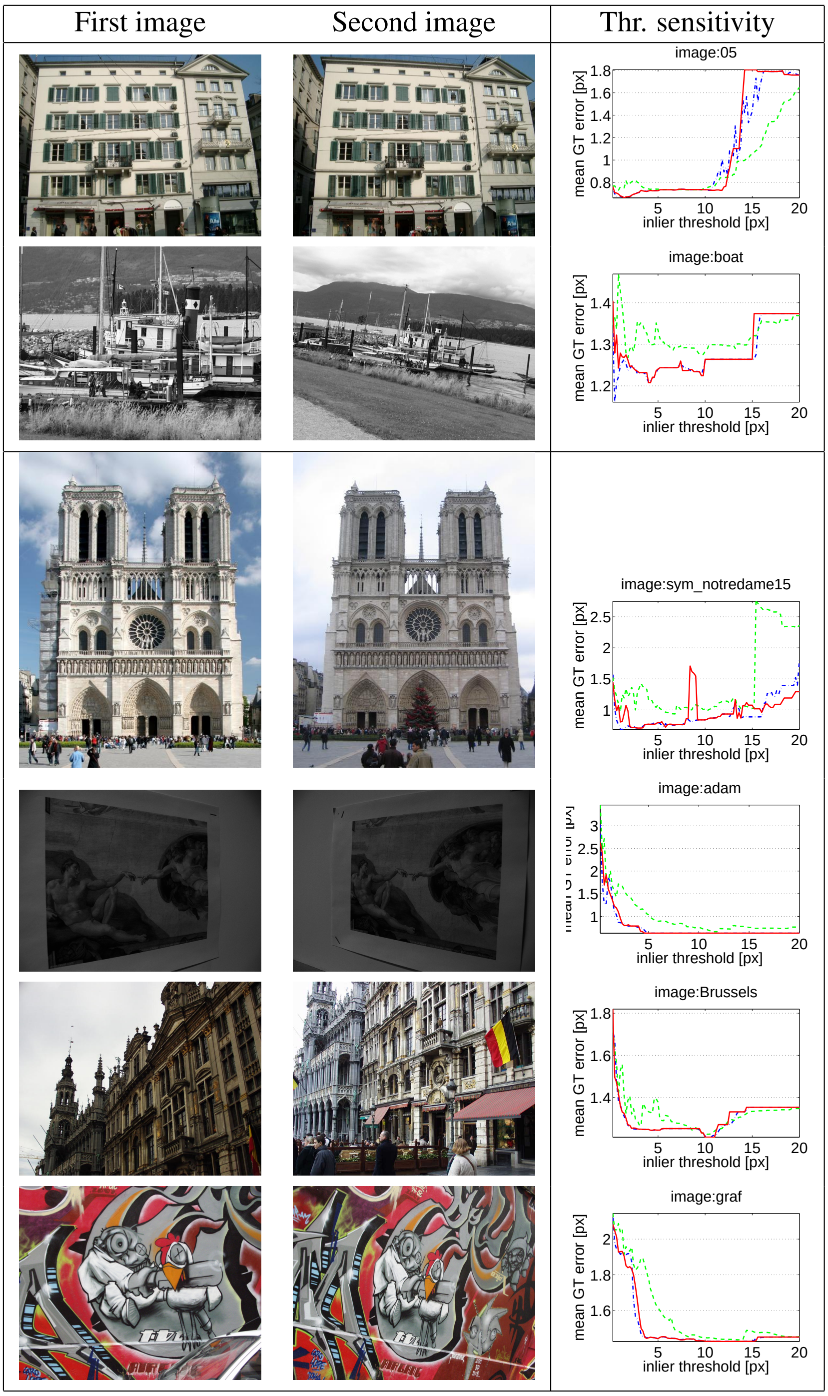

Table 3: The dependence of the ground truth error on the inlier threshold (RANSAC green, LO-RANSAC blue, L_1-based RANSAC red). Note that the proposed L_1 algorithm yields results very similar to LO-RANSAC. The ground truth error was averaged over 10 runs for each of the methods. Experimental results demonstrated on the same image pairs as in Table 2.

The RANSAC parameters common to all three tested versions used in our experiments are summarized in table 6. The inlier threshold θ is set, following[8] given σ in the following way:

$$\theta = 5.99\,(\sigma S)^2$$

where $S = \max(w, h)/768$ is a scale factor dependent image dimensions. The 5.99 term is the 95% percentile of the χ^2 distribution with two degrees of freedom.

Additional parameters used for the standard LO-RANSAC are summarized in Table 7. The proposed

Image		Qty↓	RANSAC			LO-RANSAC			L_1-based RANSAC		
BruggeSquare	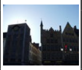	I	**201.0**	±12.9	(172-234)	**227.2**	±1.3	(224-232)	**214.9**	±11.4	(195-228)
		LO time (μs)	**0.0**	±0.0	(0-0)	**12567.0**	±1944.9	(10166-16251)	**993.8**	±187.2	(598-1650)
		I (%)	**60.0**	±3.9	(51-70)	**67.8**	±0.4	(67-69)	**64.1**	±3.4	(58-68)
		Samp	**52.6**	±24.8	(15-153)	**42.5**	±10.9	(15-59)	**17.7**	±5.7	(9-41)
		Time$_{(ms)}$	**5.2**	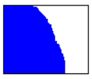	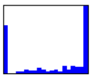	**18.0**		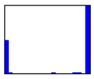	**5.8**	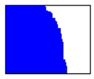	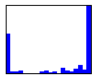
		Error	**3.50**	±1.25	(1.2-6.2)	**2.44**	±0.12	(2.0-2.7)	**2.93**	±0.91	(1.3-4.6)
		LO count	**0.0**	±0.0	(0-0)	**1.0**	±0.0	(1-1)	**2.7**	±1.2	(1-5)
dlazky	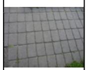	I	**11.0**	±0.2	(9-11)	**11.0**	±0.0	(11-11)	**10.9**	±0.6	(7-11)
		LO time (μs)	**0.0**	±0.0	(0-0)	**1531.3**	±746.7	(603-4434)	**249.7**	±68.9	(109-419)
		I (%)	**14.8**	±0.3	(12-15)	**14.9**	±0.0	(15-15)	**14.7**	±0.8	(9-15)
		Samp	**10745.0**	±5429.3	(6963-27356)	**8392.7**	±3432.8	(6963-25008)	**8220.6**	±3301.0	(4820-19947)
		Time$_{(ms)}$	**87.3**		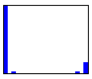	**76.0**		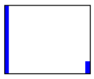	**71.6**	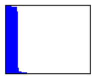	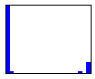
		Error	**2.99**	±0.95	(2.6-6.4)	**2.61**	±0.00	(2.6-2.6)	**5.43**	±20.63	(2.6-204.9)
		LO count	**0.0**	±0.0	(0-0)	**4.7**	±1.6	(1-9)	**7.0**	±3.1	(2-21)

Table 4: Results on two image pairs with unusual sensitivity to the inlier threshold. See caption of Tab. 2 for description of entries.

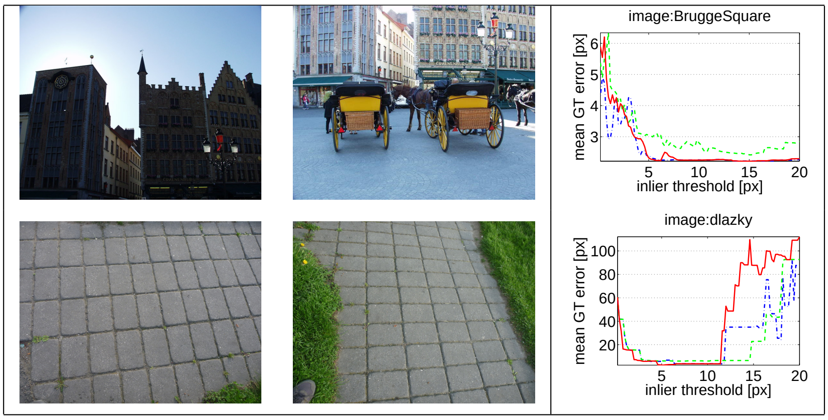

Table 5: The dependence of the ground truth error on the inlier threshold (RANSAC green, LO-RANSAC blue, L_1-based RANSAC red) for two failure cases.

confidence	0.95
σ	2.0
sample limit	500000

Table 6: RANSAC parameters

ILSQ iterations	4
ILSQ sample limit	28
threshold multiplier	4
inner RANSAC repetitions	10

Table 7: LO-RANSAC parameters

method does not introduce any new parameters.

Table 2 shows a sample of six image pairs well representing the results on the whole dataset, with the exception of a few cases described later. Note that the proposed L_1 optimization is usually about 5 times faster than the standard LO step (see 'LO time').Table 4 summarizes the performance on the few exceptional cases.

The error (see 'Error' in the table) was computed by reprojecting hand-made ground truth correspondences (about 8 of them for each image pair) by the model found by the algorithm used.

Two observations summarize the results: i) the proposed procedure yields error which is comparable to the standard LO-RANSAC, and ii) it usually runs approximately 5 times faster (see 'LO time' in the table).

Table 3 shows the comparison of the dependence of the error on the inlier threshold for standard RANSAC, standard LO-RANSAC and the proposed method. The results shown on the same subset of six image pairs which are representative of the whole dataset. The experiment confirms that the proposed procedure is able to achieve results similar to the standard LO-RANSAC.

The results for two exceptional image pairs are shown in table 5. The standard LO-RANSAC achieves good results (high stability, low error), while our proposed algorithm fails to stabilize the plain RANSAC results (the 'dlazky' pair is one of the most challenging ones from our dataset, as there are only 11 inliers).

4. Conclusions

We have shown that replacing the standard LO step of LO-RANSAC with minimization of the sum of Huber kernel responses to residuals has the following properties: it is simple, deterministic and produces similar errors as the standard LO-RANSAC and is usually approximately 5 times faster. On the negative side, in the current implementation, it has higher probability of failure than the standard LO-RANSAC.

Acknowledgements

The research was supported by CTU student grant SGS15/155/OHK3/2T/13.

References

[1] A. Beck and S. Sabach. Weiszfeld's method: Old and new results. *Journal of Optimization Theory and Applications*, 164(1):1–40, 2014. 2

[2] O. Chum, J. Matas, and J. Kittler. Locally optimized ransac. In *Pattern Recognition*, pages 236–243. Springer, 2003. 1

[3] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981. 1

[4] L. V. G. H. Shao, T. Svoboda. Zubud - zurich buildings database for image based recognition. Technical report, 2003. 3

[5] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Cambridge, UK, 2000. Chapter 3 : Estimation - 2D Projective Transformations. 1, 2

[6] D. C. Hauagge and N. Snavely. Image matching using local symmetry features. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 206–213. IEEE, 2012. 3

[7] P. Huber. *Robust Statistics*. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section Series. Wiley, 2004. 3

[8] K. Lebeda, J. Matas, and O. Chum. Fixing the locally optimized ransac. In R. Bowden, J. Collomosse, and K. Mikolajczyk, editors, *Proceedings of the British Machine Vision Conference*, pages 1013–1023, London, UK, September 2012. BMVA. 1, 2, 5

[9] K. Lebeda, J. Matas, and O. Chum. Fixing the locally optimized RANSAC – Full experimental evaluation. Research Report CTU–CMP–2012–17, Center for Machine Perception, Czech Technical University, Prague, Czech Republic, September 2012. 3

[10] R. Raguram, O. Chum, M. Pollefeys, J. Matas, and J. Frahm. Usac: A universal framework for random sample consensus. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(8):2022–2038, Aug 2013. 1

[11] B. Tordoff and D. W. Murray. Guided sampling and consensus for motion estimation. In *Computer VisionECCV 2002*, pages 82–96. Springer, 2002. 1

[12] E. Weiszfeld and F. Plastria. On the point for which the sum of the distances to n given points is minimum. *Annals of Operations Research*, 167(1):7–41, 2009. 3

[13] Z. Zhang. Parameter estimation techniques: A tutorial with application to conic fitting. *Image and vision Computing*, 15(1):59–76, 1997. 3