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

One of the attractive properties of RANSAC [2], at least with the top-hat (inlier 1, outlier 0) cost function, is 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 [1, 9] 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.* [1] introduced the LO-RANSAC which applies a local optimization (LO) step to promising hypotheses generated from random minimal samples. Experiments in [1] 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.

Chum *et al.* [1] 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 a high fraction on 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 for 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 (see Figure 1 for an illustration of the dependence). 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.

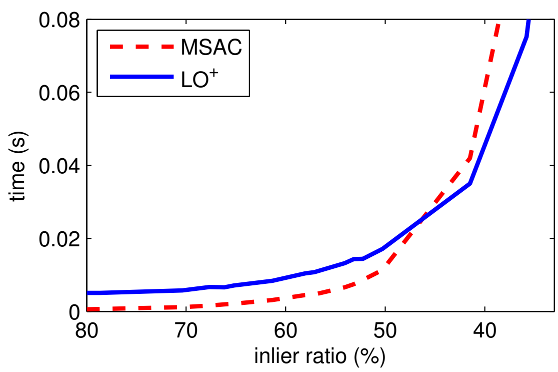

Figure 1: Dependence of the time complexity on the inlier ratio for a selected pair of images.

As a second contribution, we introduce a fast version – LO' that has execution time close to the standard RANSAC and perform close to LO-RANSAC in almost all cases. 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 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 settings. 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 [10] 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 by using the inlier limit and the truncated quadratic cost function.

The accuracy of the proposed LO method is tested within a standard Bundle adjustment method [5]. 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 [3].

In our experiments, tentative correspondences were obtained by matching SIFT descriptors [6] of MSER's [7]. In the supplementary material [4], also experiments using Hessian Affine detector [8] are presented. The results on Hessian Affine features are even more favourable for the LO methods because of lower inlier ratios. Basically, they show our conclusions are independent of the selection of detectors.

The experimental evaluation shows that: (1) the LO$^+$-RANSAC with MSAC cost function offers a stable robust estimation despite its randomized nature, (2) limiting the number of inliers included in the (iterative) least squares significantly reduces execution time and often even improves the precision, (3) the speed of the minimalistic version LO' is comparable to plain RANSAC even for easy problems with very high inlier ratios, and that (4) LO-RANSAC offers significantly better starting point for bundle adjustment than the Gold Standard [3].

[1] O. Chum, J. Matas, and J. Kittler. Locally optimized RANSAC. In *DAGM-Symposium*, pages 236–243, 2003.

[2] 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 ACM*, 24(6):381–395, 1981.

[3] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004.

[4] 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, 2012.

[5] 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.

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

[7] 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.

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

[9] B. Tordoff and D. W. Murray. Guided sampling and consensus for motion estimation. In *Proc. of ECCV*, pages 82–98, 2002.

[10] P. H. S. Torr and A. Zisserman. Robust computation and parametrization of multiple view relations. In *Proc. of ICCV*, pages 727 –732, 1998.