

Randomized RANSAC with Sequential Probability Ratio Test

958

Abstract

A randomized model verification strategy for RANSAC is presented. The proposed method finds, like RANSAC, a solution that is optimal with user-controllable probability η . A provably optimal model verification strategy is designed for the situation when the contamination of data by outliers is known, i.e. the algorithm is the fastest possible (on average) of all randomized RANSAC algorithms guaranteeing $1 - \eta$ confidence in the solution. The derivation of the optimality property is based on Wald's theory of sequential decision making. The R-RANSAC with SPRT, which does not require the a priori knowledge of the fraction of outliers and has results close to the optimal strategy, is introduced. We show experimentally that on standard test data the method is 2 to 10 times faster than the standard RANSAC and up to 4 times faster than previously published methods.

1 Introduction

The RANSAC (**R**ANd**o**m **S**AMple **C**onsensus) algorithm introduced by Fishler and Bolles in 1981 [1] is a widely used robust estimator that has become a de facto standard in the field of computer vision [2]. RANSAC and related hypothesize-and-verify methods [9, 5] have been applied in many vision problems: short baseline stereo [8, 10], wide baseline stereo matching, motion segmentation [8], mosaicing, detection of geometric primitives, robust eigenimage matching, structure and motion estimation [6, 7], object recognition and elsewhere.

The RANSAC algorithm proceeds as follows. Repeatedly, subsets of the input data (e.g. a set of tentative correspondences) are randomly selected and model parameters fitting the sample are computed. In a second step, the quality of the model is evaluated on the input data. Different cost functions have been proposed [9], the standard being the number of inliers, i.e. the number of data points consistent with the model. The process is terminated when the probability of finding a better model becomes lower than a user-controlled probability η . The $1 - \eta$ confidence in the solution holds for all levels of contamination of the input data, i.e. for any number of outliers within the input data.

The speed of standard RANSAC depends on two factors.

The percentage of outliers determines the number of random samples needed to guarantee the $1 - \eta$ confidence in the solution. The time needed to assess the quality of a hypothesized model parameters is proportional to the number N of the input data points. Almost all models whose quality is verified are incorrect with arbitrary parameters originating from contaminated samples. Such models are consistent with only a small number of the data points. In [4], Matas and Chum showed how this property can be exploited to increase the speed of RANSAC. The algorithm, called R-RANSAC, speeds up the model evaluation step by introducing a two-stage procedure. First, a statistical test is performed on d randomly selected data points ($d \ll N$). Evaluation of the remaining $N - d$ data points is carried out only if the first d data points are inliers. The speed up of the R-RANSAC depends on the probabilities of the two types of errors committed in the pre-test, the rejection of an uncontaminated model and the acceptance of a contaminated model. The idea was modified by Nistér to include competitive verification of models and exploited in an impressive real-time structure from motion system [6]. The main limitation of Nistér's pre-emptive RANSAC is that a fixed number of models is evaluated, which is equivalent to an a priori assumption that the fraction of inliers is known. This limits the applicability of pre-emptive RANSAC as in many problems, e.g. wide baseline stereo, the fraction of inliers ranges widely.

It is noted in [4] that the two-stage procedure of R-RANSAC is not optimal. *As a main contribution of this paper, we define an optimal hypothesis evaluation procedure, i.e. a method for randomized model quality evaluation that returns, in the fastest average time possible, a solution with the confidence $1 - \eta$. The derivation of the optimality property is based on Wald's theory of sequential decision making [11].*

From the statistical point of view, the evaluation of the model quality can be formulated as a sequential test as follows. The hypothesis generation step proposes a model. It is either 'good', i.e. it leads to the optimal solution (the solution with maximal support), or it is 'bad', i.e. one of the data points in the sample is an outlier. The property 'good' is a hidden state that is not directly observable but is statistically linked to observable events. The observable events are "data point (correspondence) is/is-not consistent with the model".

In sequential testing, as applied e.g. in industrial inspection, the problem is to decide whether the model (or the batch of products) is ‘good’ or ‘bad’ in the shortest possible time (i.e. making the smallest number of observations) and yet satisfying the predefined bounds on the probabilities of the two possible errors – accepting a ‘bad’ model as ‘good’ and vice versa. Wald proposed the sequential probability ratio test (SPRT) and showed [11] that, given errors bound on the errors of the first and second kind, it minimizes the the number of observations (time to decision) ¹.

Wald’s SPRT test is a solution of a *constrained optimization* problem. The user supplies the acceptable probabilities of the errors of the first and the second kind and the resulting optimal test is a trade-off between time to decision (or cost of observations) and the errors committed. However, when evaluating RANSAC, the situation is different. First of all, a ‘good’ model is always evaluated for all data points (correspondences) since the number of inliers is one of the outputs of the algorithms. So the only error that can be committed is an early rejection of a ‘good’ model (error of the first kind). But this only means that more samples have to be drawn to achieve the required confidence $1 - \eta$ of finding the optimal solution. So unlike in the classical setting, we are solving a *global optimization* problem, minimizing a single real number – the time to decision, since the consequence of an error is also a loss of time.

The structure of this paper is as follows. First, in Section 2, we introduce the relevant parts of Wald’s decision theory and show how its results can be brought to bear on the problem of minimizing RANSAC run time. The RANSAC with SPRT algorithm is described in detail in Section 3. In Section 4, the theoretical results are experimentally verified on standard stereo matching problems. The paper is concluded in Section 5.

2. The Optimal Sequential Test

The model evaluation step of the optimal R-RANSAC proceeds as Wald’s sequential probability ratio test (SPRT) with the probability α of rejecting a ‘good’ sample set to achieve maximum speed of the whole RANSAC process. To understand the operation of R-RANSAC with SPRT, some familiarity with Wald’s decision theory is required. We therefore introduce its relevant parts. Some of the results are presented in a form that is not fully general, but sufficient for the derivation of the R-RANSAC with SPRT algorithm. Some of Wald’s terminology is modified in order to make the exposition more accessible.

In the model evaluation step, our objective is to decide between the hypothesis H_g that model is ‘good’ and the

¹Precisely speaking, the SPRT is only approximately optimal. However, the approximation has been shown by Wald to be so close to the optimum that, for practical purposes, it is considered the optimal test.

alternative hypothesis H_b that the model is ‘bad’. A ‘good’ model is computed from an all-inlier sample. The Wald’s SPRT is based on the likelihood ratio [11]

$$\lambda_j = \prod_{r=1}^j \frac{p(x_r|H_b)}{p(x_r|H_g)} = \lambda_{j-1} \cdot \frac{p(x_j|H_b)}{p(x_j|H_g)}, \quad (1)$$

a ratio of two conditional probabilities of an observation x_r under the assumptions of H_g and H_b respectively. In RANSAC, x_r is equal to 1 if the r -th data point is consistent with a model with parameters θ and 0 otherwise. For example, a correspondence is consistent with (i.e. supporting) an epipolar geometry represented by a fundamental matrix F if its Sampson’s error is smaller than some predefined threshold [2]. The probability $p(1|H_g)$ that any randomly chosen data point is consistent with a ‘good’ model is approximated by the fraction of inliers ε among the data points². The probability of a data point being consistent with a ‘bad’ model is modeled as a probability of a random event with Bernoulli distribution with parameter δ : $p(1|H_b) = \delta$. The process of estimation of δ and ε is discussed in Section 3.

Output: model accepted/rejected, number of tested data points, a fraction of data points consistent with the model

Set $j = 1$

- 1 Check whether j -th data point is consistent with the model
 - 2 Compute the likelihood ratio λ_j eq. (1)
 - 3 If $\lambda_j > A$, decide the model is ‘bad’ (model “rejected”), else increment j or continue testing
 - 4 If $j = N$ the number of correspondences decide model “accepted”
-

Algorithm 1: The adapted sequential probability ratio test (Adapted SPRT).

After each observation the standard Wald’s SPRT makes one of three decisions: accept a ‘good’ model, reject a ‘bad’ model, or continue testing. Since in RANSAC the total number of inliers is needed to decide on termination, nothing is gained by an early decision in favor of a ‘good’ model. Therefore the option of an early acceptance of the model has been removed in the adapted SPRT (Alg. 1). The full SPRT is described e.g. in Wald [11] and, in a more accessible form, in Lee [3].

2.1. The optimal value of the decision threshold

The decision threshold A is the only parameter of the Adapted SPRT. We show how to set it to achieve optimal

²The probability ε would be exact if the data points were selected with replacement. Since the objective of the verification is to count the size of the support of the model, the correspondences are drawn without replacement. However, the approximation is close.

performance, i.e. minimal average RANSAC running time given the probabilities δ and ε . We use the following theorems (for proofs, see [11]).

Theorem 1 *The probability α of rejecting a ‘good’ model in SPRT $\alpha \leq 1/A$.*

Proof: Wald’s theorem [11, p. 41] states $\alpha \leq (1 - \beta)/A$, where β stands for the probability that a ‘bad’ model is incorrectly accepted as ‘good’. In the adapted SPRT, since the only decision of the test can be “reject”, $\beta = 0$ and thus $\alpha \leq 1/A$. \square

The approximation $\alpha \approx 1/A$ is close and is often used.

Theorem 2 (Wald’s lemma) *The average number of observations (checked data points) carried out while testing a ‘bad’ model is $C^{-1} \log A$, where*

$$C = p(0|H_b) \log \frac{p(0|H_b)}{p(0|H_g)} + p(1|H_b) \log \frac{p(1|H_b)}{p(1|H_g)}. \quad (2)$$

Proof: According to [11, p. 53]

$$C = E \left(\log \frac{p(x|H_b)}{p(x|H_g)} \right).$$

The value of x is from $\{0, 1\}$. The expectation E is a sum of two terms weighted by probability $p(x|H_b)$. Equation (2) follows. \square

In the particular case of RANSAC, $p(1|H_b) = \delta$, $p(0|H_b) = 1 - \delta$, $p(0|H_g) = 1 - \varepsilon$, and $p(1|H_g) = \varepsilon$. Therefore the average number of verified correspondences per model is:

$$\frac{\log A}{C} = \left((1 - \delta) \log \frac{1 - \delta}{1 - \varepsilon} + \delta \log \frac{\delta}{\varepsilon} \right)^{-1} \log A. \quad (3)$$

The value of A influences the total running time in two opposing ways. The larger the value of A , the smaller the probability of rejection of a ‘good’ model. On the other hand, the number of correspondences verified per model increases with $\log A$ (eq (3)). We wish to set A to achieve minimal average time needed to find the solution.

The average time-to-solution in R-RANSAC is $t = \bar{k} \bar{t}_s$, where \bar{k} is the average number of samples drawn until a ‘good’ model and \bar{t}_s is the average testing time per sample. In the following, the time unit will be the time needed to check one data point. The probability P_g of drawing a ‘good’ model is $P_g = \varepsilon^m$, where m is the number of data points in the RANSAC sample. The number of tested samples before a ‘good’ one is drawn and not rejected is a random variable with geometric distribution and mean $\bar{k} = 1/(P_g(1 - \alpha)) \approx 1/(P_g(1 - 1/A))$. The average time

\bar{t}_s of processing a sample consists of two components: time t_M needed to instantiate a model hypotheses given a sample³, and the average time of testing each hypothesis. Let \bar{m}_S be the number of models that are verified per sample⁴ and $C^{-1} \log A$ be the average length of the SPRT (Theorem 2). The average time to the solution expressed as a function of A is

$$t(A) = \frac{1}{P_g(1 - 1/A)} \left(t_M + \bar{m}_S \frac{\log A}{C} \right). \quad (4)$$

The formula (4) can be simplified to

$$t(A) = \frac{K_1 + K_2 \log A}{1 - 1/A},$$

where $K_1 = t_M/P_g$ and $K_2 = \bar{m}_S/(P_g C)$. We are interested in the optimal value of A , i.e.

$$A^* = \arg \min_A t(A).$$

The minimum is found by solving

$$\frac{dt}{dA} = - \frac{K_1 + K_2 - K_2 A + K_2 \log A}{(A - 1)^2} = 0.$$

After rearrangements, we have

$$A^* = \frac{K_1}{K_2} + 1 + \log A^* = \frac{t_M C}{\bar{m}_S} + 1 + \log A^*. \quad (5)$$

Equation (5) has two real solutions for positive K_1/K_2 , $A_1^* < 1 < A_2^*$. Since $\delta < \varepsilon$, the contribution to the likelihood ratio (eq. (1)) of a correspondence that is not consistent with the model is greater than 1, therefore the solution of interest is $A^* > 1$. This solution can be obtained as $A^* = \lim_{n \rightarrow \infty} A_n$, where $A_0 = K_1/K_2 + 1$ and $A_{n+1} = K_1/K_2 + 1 + \log(A_n)$. The series converges fast, typically within four iterations.

3. R-RANSAC with SPRT

The R-RANSAC with SPRT algorithm is outlined in Alg. 2. To fully specify the details of the algorithm, two issues have to be addressed. First, the estimation of parameters δ and ε ; second, the termination criterion guaranteeing $1 - \eta$ confidence in the solution has to be derived.

The algorithm proceeds like standard RANSAC [1, 2], only instead of checking all data points in the model verification step, the data points are evaluated sequentially and hypotheses with low support are rejected early. After a hypothesis is rejected, δ is re-estimated (Alg. 2, step 2a). Accepted hypotheses are candidates for the RANSAC outcome

³Computing model parameters from a sample takes the same time as verification of t_M data points.

⁴In the 7-pt algorithm for epipolar geometry estimation, 1 to 3 models have to be verified.

Initialize ε_0, δ_0 , calculate A_0 and set $i = 0$.

Repeat until the probability η (eq. (8)) of finding a model with support larger than $\hat{\varepsilon}$ falls under a user defined confidence value η_0 :

1. Hypothesis generation

- Select a random sample of minimum size m from the set of data points.
- Estimate model parameters θ fitting the sample.

2. Verification

Execute the SPRT (Alg. 1) and update the estimates if

- Model rejected: re-estimate δ . If the estimate $\hat{\delta}$ differs from δ_i by more than 5% design $(i+1)$ -th test ($\varepsilon_{i+1} = \varepsilon_i, \delta_{i+1} = \hat{\delta}, i = i + 1$)
 - Model accepted and the largest support so far: design $(i+1)$ -th test ($\varepsilon_{i+1} = \hat{\varepsilon}, \delta_{i+1} = \hat{\delta}, i = i + 1$). Store the current model parameters θ .
-

Algorithm 2: The structure of R-RANSAC with SPRT.

(see below). The overhead of the evaluation of the likelihood ratio λ_j eq. (1) is negligible compared to the evaluation of the model versus data point error function.

The optimal test derived in Section 2 requires the knowledge of two parameters, ε and δ . These probabilities are different for different data sets and we assume they are unknown. The proposed algorithm uses values of ε and δ that are estimated during the sampling process and the test is adjusted to reflect the current estimates.

If the probabilities ε and δ are available a-priori, e.g. in some standard setting where the algorithm is run repeatedly, they can be used in the initialisation of the algorithm.

Estimation of δ . Since almost all tested models are ‘bad’⁵, the probability δ can be estimated as the average fraction of consistent data points in rejected models. When current estimate δ differs from the estimate used to design the SPRT (by more than 5%, for example), new $(i+1)$ -th test is designed. The initial estimate δ_0 is obtained by geometric considerations, i.e. as a fraction of the area that supports a hypothesised model (a strip around an epipolar line in case of epipolar geometry) to the area of possible appearance of outlier data (the area of the search window). Alternatively, a few models can be evaluated without applying SPRT in order to obtain an initial estimate of δ .

Estimation of ε . In general, it is not possible to obtain an unbiased estimate of ε , since this would require the knowledge of the solution to the optimization problem we are solving. The tightest lower bound on ε is provided by the size of the largest support so far. It was shown in [4] that a sample with the largest support so far appears $\log k$

⁵RANSAC verifies, on average, $-\log(\eta_0)$ ‘good’ models, e.g. for the typical $\eta_0 = 0.05$ a ‘good’ model is hypothesised three times prior to termination of the algorithm.

times, where k is the number of samples drawn. When such a sample (with support of size I_{i+1}) appears, new test is designed for $\varepsilon_{i+1} = I_{i+1}/N$. Throughout the course of the algorithm, a series of different tests with

$$\varepsilon_0 < \dots < \varepsilon_i < \dots < \varepsilon$$

are performed. The initial value of ε_0 can be derived from the maximal time the user is willing to assign to the algorithm.

The termination criterion. The algorithm is terminated, when the probability η of missing a set of inliers larger than the largest support found so far falls under a predefined threshold η_0 . In standard RANSAC, where the probability of rejection of a ‘good’ model is zero, the probability is equal to

$$\eta_R = (1 - P_g)^k.$$

In R-RANSAC, the probability of hypothesising and not rejecting a ‘good’ model is $P_g(1 - \alpha)$ and the probability η becomes as

$$\eta = (1 - P_g(1 - \alpha))^k.$$

In R-RANSAC with SPRT, the SPRT is adjusted to current estimates of δ_i and ε_i , so α is no more constant. Theorem 1, which gives the probability α of rejecting a ‘good’ model for the test designed for optimal value of ε , does not cover this situation. The following theorem is needed:

Theorem 3 *The probability of rejecting a ‘good’ model with fraction of inliers ε in a SPRT designed for ε_i and δ_i with threshold A_i is*

$$\alpha_i = A_i^{-h_i}, \tag{6}$$

where h_i is given by

$$\varepsilon \left(\frac{\delta_i}{\varepsilon_i} \right)^{h_i} + (1 - \varepsilon) \left(\frac{1 - \delta_i}{1 - \varepsilon_i} \right)^{h_i} = 1. \tag{7}$$

Proof: For proof see [11, p. 50]. □

Equation (7) has two solutions, one being $h_i = 0$. Since $\varepsilon_i < \varepsilon$, $h_i > 1$ holds for other solution. This solution is found numerically.

Let for each of l tests the following values be stored: the expected fraction of inliers ε_i , the SPRT threshold A_i , the number of samples k_i processed by the test, and h_i satisfying (7). Then, the probability η is given by

$$\eta(l) = \prod_{i=0}^l \left(1 - P_g(1 - A_i^{-h_i}) \right)^{k_i}. \tag{8}$$

The number k_l of samples that are needed to be drawn with current (i.e. l -th) SPRT follows from (8) as

$$k_l = \frac{\log \eta_0 - \log(\eta(l-1))}{\log(1 - P_g A_l^{-1})} \tag{9}$$

Implementation note: since $\eta > \eta_R$ the equation (9) does not have to be evaluated before $\eta_R < \eta_0$ is satisfied.

4 Experiments

Several experiments were performed comparing the proposed R-RANSAC with SPRT with three other RANSAC algorithms: (1) standard RANSAC that verifies all correspondences for every model, (2) R-RANSAC with the $T_{d,d}$ test [4] that rejects the model when the first checked correspondence is not consistent with it ($d = 1$), and (3) R-RANSAC with the a priori SPRT, i.e. the R-RANSAC with SPRT designed for the true values of ε and δ (labelled SPRT*). The results achieved with a priori SPRT show the best achievable performance of RANSAC with a randomized verification step for a problem characterized by given δ and ε .

For epipolar geometry estimation, the time needed to compute model parameters $t_M = 200$ was set within the range observed in a large number of experiments (i.e. in our implementation, checking whether a correspondence is consistent with a fundamental matrix is 200 times faster than estimating the matrix). The exact value depends on many factors including the CPU speed and type. The constant $\bar{m}_S = 2.38$ was set to the experimentally observed average of the number of models generated by the 7-point algorithm per sample⁶. The initial values of δ and ε were set to $\delta_0 = 0.05$ and $\varepsilon = .2$ respectively.

For homography estimation, the values were set as follows $t_M = 200$, $\bar{m}_S = 1$, $\delta_0 = 0.01$, $\varepsilon_0 = 0.1$.

The experimental image pairs are displayed in Fig. 1. The number N of correspondences as well as the true values of ε and δ estimated by evaluation 100,000 verifications of random models are summarized in Tab. 2. The results of compared algorithms are shown in Table 1.

As a consequence of the randomization of model verification that erroneously rejects some ‘good’ models, on average, the randomized algorithms must draw a larger number of samples than standard RANSAC. This is confirmed in the first column of Table 1. This small increase is more than compensated by the reduction in the number of data points (correspondences) checked on average per model. The running time of RANSAC is reduced by factors ranging from 2.8 to 10.9 In all experiments the SPRT outperforms the $T_{d,d}$ test.

⁶It is known that the 7-point algorithm produces 1 to 3 potential models. In experiments, the average number of models per sample – 2.38 – has been observed consistently in a number of scenes. No theoretical justification of the stability of this average is known to the authors.

The ROTUNDA A experiment					
	k	models	vpm	time	spd-up
RANSAC	195	470	886.0	4.3	1.0
$T_{d,d}$	345	834	49.0	1.4	3.1
SPRT*	201	486	29.1	0.8	5.4
SPRT	200	483	41.5	0.9	4.9
The GREAT WALL experiment					
	k	models	vpm	time	spd-up
RANSAC	48085	122223	507.0	667	1.0
$T_{d,d}$	224865	571059	6.9	642	1.0
SPRT*	56678	143993	14.4	178	3.8
SPRT	55796	141761	16.3	178	3.7
The CORRIDOR experiment					
	k	models	vpm	time	spd-up
RANSAC	145	344	600.0	2.3	1.0
$T_{d,d}$	243	576	86.6	1.2	2.0
SPRT*	154	366	32.9	0.6	3.7
SPRT	153	364	77.2	0.8	2.8
The LEUVEN experiment					
	k	models	vpm	time	spd-up
RANSAC	1855	4434	786.0	35.5	1.0
$T_{d,d}$	4490	10730	27.5	14.8	2.4
SPRT*	2025	4839	20.4	6.8	5.3
SPRT	1982	4736	23.2	6.8	5.2
The GRAFFITI experiment					
	k	models	vpm	time	spd-up
RANSAC	121	121	405.0	3.0	1.0
$T_{d,d}$	287	287	16.7	1.0	2.8
SPRT*	132	132	35.1	0.6	4.6
SPRT	129	129	47.1	0.7	4.1
The homography LEUVEN experiment					
	k	models	vpm	time	spd-up
RANSAC	1203	1203	789.0	53.1	1.0
$T_{d,d}$	5323	5323	3.4	15.0	3.5
SPRT*	1351	1351	14.5	4.7	11.2
SPRT	1266	1266	20.7	4.9	10.9

Table 1: The comparison of RANSAC, R-RANSAC with $T_{d,d}$ test, a priori SPRT* and SPRT: the number of samples (k), the number of models (models), the number of checked correspondences per model (vpm), time in ms (time), and relative speed-up with respect to standard RANSAC (spd-up). Averaged over 500 runs.

5 Conclusions

We have derived an optimal sequential strategy for randomised evaluation of model quality in RANSAC. A method for estimating the two probabilities characterising the problem and critically influencing the design of the optimal strategy was proposed. A termination criterion maintaining the $1 - \eta$ confidence in the solution was derived.

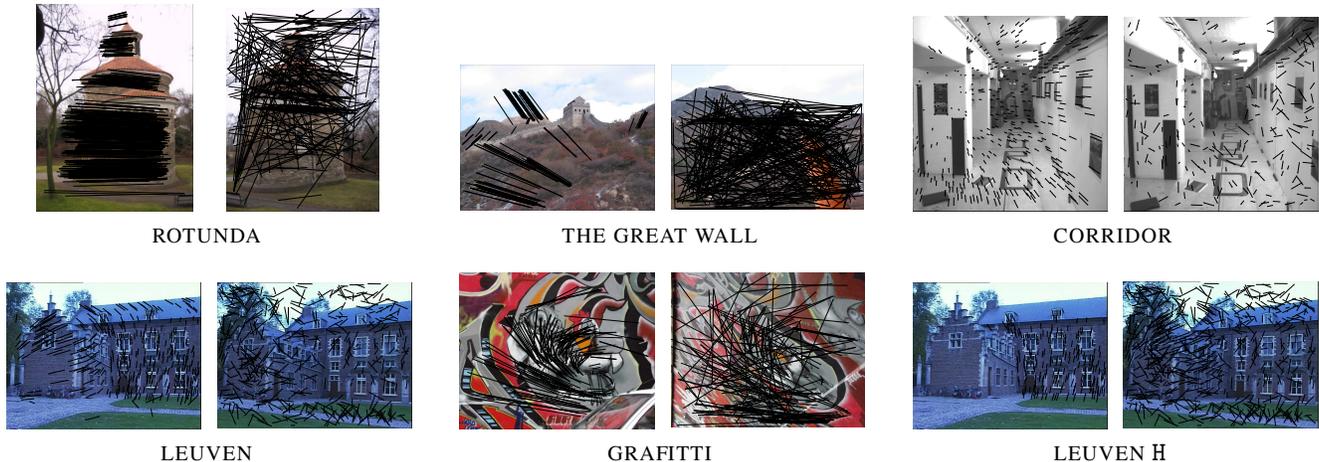


Figure 1: The experimental image pairs with inliers (left) and outliers (right) superimposed. Two wide-baseline epipolar geometry experiments ROTUNDA and the GREAT WALL; two narrow-baseline EG experiments LEUVEN and CORRIDOR; two homography experiments GRAFFITI and LEUVEN H.

	corr	ε	δ
ROTUNDA	893	0.60	0.052
THE GREAT WALL	514	0.28	0.014
CORRIDOR	607	0.56	0.142
LEUVEN	793	0.47	0.032
GRAFFITI	409	0.51	0.018
LEUVEN H	793	0.29	0.004

Table 2: Number of correspondences (corr), fraction of inliers (ε), the probability of a correspondence being consistent with bad model (δ).

Properties of R-RANSAC with SPRT were tested on wide range of standard data. Tests included epipolar geometry estimation in both wide baseline setting and narrow baseline settings and homography estimation.

The method was 2.8 to 10 times faster than the standard RANSAC and up to 4 times faster than R-RANSAC with $T_{d,d}$ test.

References

- [1] M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *CACM*, 24(6):381–395, June 1981.
- [2] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge University, Cambridge, 2nd edition, 2003.
- [3] Peter M. Lee. Sequential probability ratio test. www.york.ac.uk/depts/maths/teaching/pml/ais/sprt.ps.
- [4] Jiri Matas and Ondrej Chum. Randomized RANSAC with $T_{d,d}$ test. *Image and Vision Computing*, 22(10):837–842, September 2004.
- [5] D.R. Myatt, P.H.S. Torr, S.J. Nasuto, J.M. Bishop, and R. Craddock. NAPSAC: High noise, high dimensional robust estimation - it’s in the bag. In *BMVC02*, volume 2, pages 458–467, 2002.
- [6] D. Nister. Preemptive RANSAC for live structure and motion estimation. In *Proc. ICCV03*, volume 1, pages 199–206, October 2003.
- [7] B. Tordoff and D.W. Murray. Guided sampling and consensus for motion estimation. In *Proc. 7th ECCV*, volume 1, pages 82–96. Springer-Verlag, 2002.
- [8] P. H. S. Torr. *Outlier Detection and Motion Segmentation*. PhD thesis, Dept. of Engineering Science, University of Oxford, 1995.
- [9] P. H. S. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *CVIU*, 78:138–156, 2000.
- [10] P.H.S. Torr, A. Zisserman, and S.J. Maybank. Robust detection of degenerate configurations while estimating the fundamental matrix. *CVIU*, 71(3):312–333, Sep. 1998.
- [11] Abraham Wald. *Sequential analysis*. Dover, New York, 1947.