

Matching with PROSAC – Progressive Sample Consensus

Ondřej Chum

Jiří Matas

Center for Machine Perception
 Department of Cybernetics
 Czech Technical University in Prague, Czech Republic

Abstract

A new robust matching method is proposed. The Progressive Sample Consensus (PROSAC) algorithm exploits the linear ordering defined on the set of correspondences by a similarity function used in establishing tentative correspondences. Unlike RANSAC, which treats all correspondences equally and draws random samples uniformly from the full set, PROSAC samples are drawn from progressively larger sets of top-ranked correspondences.

Under the mild assumption that the similarity measure predicts correctness of a match better than random guessing, we show that PROSAC achieves large computational savings. Experiments demonstrate it is often significantly faster (up to more than hundred times) than RANSAC.

For the derived size of the sampled set of correspondences as a function of the number of samples already drawn, PROSAC converges towards RANSAC in the worst case. The power of the method is demonstrated on wide-baseline matching problems.

1. Introduction

Finding¹ reliable correspondences in two or more images is a difficult and critical step in many computer vision problems such as narrow and wide-baseline stereo matching [10, 17, 11, 7, 6] structure and motion estimation [8, 14], image retrieval, and object recognition [12].

It is generally accepted that incorrect matches cannot be avoided in the first stage of the matching process where only local image descriptors are compared. The mismatches, due to phenomena like occlusions, depth discontinuities and repetitive patterns, are detected and removed by robust methods that search for sets of matches consistent with some global constraint. Random sample consensus RANSAC [3] and similar robust hypothesize-and-verify methods [15, 14] have become the methods of choice for outlier removal.

¹The authors were supported by the Czech Science Foundation under project GACR 102/03/0440 and by the European Commission under project IST-004176.




Figure 1: The Great Wall image pair with an occlusion. Given 250 tentative correspondences as input, both PROSAC and RANSAC found 57 correct correspondences (inliers). To estimate the epipolar geometry, RANSAC tested 106,534 seven-tuples of correspondences in 10.76 seconds while PROSAC tested only 9 seven-tuples in 0.06 sec (on average, over hundred runs). Inlier correspondences are marked by a line segment joining the corresponding points.

Standard RANSAC does not model the local matching process. It is viewed as a black box that generates N tentative correspondences, i.e. the error-prone matches established by comparing local descriptors. The set \mathcal{U} of tentative correspondences contains an *a priori* unknown number I of

correct matches (inliers). The inliers are consistent with a global geometric model that is found by fitting a model to a randomly selected subset of \mathcal{U} . The hypothesize-and-test loop is terminated when the probability of finding a superior solution falls below a pre-selected threshold. The time complexity of RANSAC depends on N , I , and the complexity m of the geometric model. The average number of samples drawn is proportional to $(N/I)^m$ [4].

In this paper, we introduce a new hypothesize-and-verify (sample-and-test) matching approach called PROSAC (PROgressive SAmple Consensus). The method achieves large computational savings (with speed-up factors of the order of 10^2 compared to RANSAC) by exploiting the linear ordering structure of \mathcal{U} . The ordering is defined at least implicitly in all commonly used local matching methods, because the set of tentative correspondences is obtained by first evaluating a real-valued similarity function (or “quality”) $q(\cdot)$ that is subsequently thresholded to obtain the N correspondences. Correlation of intensities around points of interest [18], Mahalanobis distance of invariant descriptors [17] or the ratio of distances in the SIFT space of the first to second nearest neighbor [5] are commonly used examples of $q(\cdot)$.

In PROSAC, samples are semi-randomly drawn from progressively larger sets of tentative correspondences. The improvement in efficiency rests on the mild assumption that tentative correspondences with high similarity are more likely to be inliers. More precisely, we assume that the ordering defined by the similarity used during the formation of tentative matches is not worse than random ordering. The assumption was found valid in our experiments, for all quality function and for all tested image pairs. Experiments presented in Section 3 demonstrate that the fraction of inliers among the top n sorted correspondences falls off fairly rapidly and consequently PROSAC is orders of magnitude faster than the worst-case prediction.

The PROSAC process is in principle simple, but to fully specify it, two problems must be addressed. First, the growth function $n = g(t)$ that defines the set \mathcal{U}_n of n top-ranked correspondences that is sampled after t trials must be selected. Second, a stopping criterion giving guarantees similar to RANSAC about the optimality of the obtained solution must be found. We propose a growth function $g(t)$ guaranteeing that PROSAC is at least equally likely to find the optimal solution as RANSAC. However, we have not been able to prove analytically that PROSAC and RANSAC have the same performance for the worst-case situation, i.e. when the correspondences are ordered randomly. Nevertheless, the comparison of PROSAC and RANSAC on randomly ordered sets of correspondences showed that their performance was effectively identical.

The PROSAC algorithm has two other desirable features. The size N of the set of tentative correspondences has limited influence on its speed, since the solution is typically

found early, when samples are taken from a smaller set. One parameter of the matching process is thus effectively removed. Instead, the user controls the behavior of PROSAC by specifying the time when the sampling distribution of PROSAC and RANSAC become identical. For the growth function $g(t)$ selected according to the above-mentioned criteria, PROSAC can be interpreted as a process running RANSAC processes in parallel for all $\mathcal{U}_n, n \in \{m \dots N\}$. In experiments presented in Section 3, PROSAC speed was close to that of RANSAC that would operate on (the *a priori* unknown) set of correspondences with the highest inlier ratio.

Related work. Tordoff and Murray [14] combine the MLESAC [15] algorithm with non-uniform (guided) sampling of correspondences. This is the published work closest to PROSAC, that differs in two important aspects. First, guided sampling requires estimates of the probability of correctness of individual correspondences while here we only assume that some quantity monotonically related to the probability is available. Second, PROSAC dynamically adapts the sampling strategy to the information revealed by the sampling process itself. The hypothesize-and-verify loop is a series of incorrect guesses until the first success. Each failure decreases the likelihood that the correspondences used to estimate the model parameters are correct. Gradually, the observed evidence against *a priori* preferred correspondences should result in the reduction of their preference. PROSAC can be viewed as an instance of a process that starts by deterministically testing the most promising hypotheses and then converging to uniform sampling as the confidence in the “quality” of the *a priori* sorting declines after unsuccessful tests.

The objective of PROSAC is to find inliers in the set of all tentative correspondences \mathcal{U}_N in the shortest possible time and to guarantee, with a certain probability, that all inliers from \mathcal{U}_N are found. Issues related to the precision of the model that is computed from the set of inliers are not discussed since they are not directly related to the problem of efficient sampling. Bundle adjustment [16] can be performed *ex post*.

1.1. Notation

The set of N data points (tentative correspondences) is denoted as \mathcal{U}_N . The data points in \mathcal{U}_N are sorted in descending order with respect to the quality function q

$$\mathbf{u}_i, \mathbf{u}_j \in \mathcal{U}_N : i < j \Rightarrow q(\mathbf{u}_i) \geq q(\mathbf{u}_j).$$

A set of n data points with the highest quality is denoted \mathcal{U}_n . A sample \mathcal{M} is a subset of data points $\mathcal{M} \subset \mathcal{U}_N, |\mathcal{M}| = m$ where m is the size (cardinality) of the sample. The quality function on samples is defined as the lowest quality of a data

point included in the sample

$$q(\mathcal{M}) = \min_{\mathbf{u}_i \in \mathcal{M}} q(\mathbf{u}_i).$$

2. Algorithm

The structure of the PROSAC algorithm is similar to RANSAC. First, hypotheses are generated by random sampling. The samples, unlike in RANSAC, are not drawn from all data, but from a subset of the data with the highest quality. The size of the hypothesis generation set is gradually increased. The samples that are more likely to be uncontaminated are therefore examined early. In fact, PROSAC is designed to draw the same samples as RANSAC, only in a different order. The hypotheses are verified against all data. As in RANSAC, the algorithm is terminated when the probability of the existence of solution that would be better than the best so far becomes low (smaller than 5%). Two important issues, the choice of the size of the hypothesis generation set and the termination criterion of the sampling process, are discussed below.

2.1. The growth function and sampling

The design of the growth function defining the \mathcal{U}_n must find a balance between the over-optimistic reliance on the pre-sorting by the quality and the over-pessimistic RANSAC approach that treats all correspondences equally. If the probabilities $P\{\mathbf{u}_i\} = P\{\text{correspondence } \mathbf{u}_i \text{ is correct}\}$ were known, it would be in principle possible to adopt a Bayesian approach. After each sample-and-test cycle, the posterior probability would be re-computed for all correspondences included in the sample. The correspondences would be sorted by their posterior probabilities and samples with the highest probability would be drawn. We pursued this line of research, but abandoned it for two reasons. Firstly, probabilities $P\{\mathbf{u}_i\}$ of correspondences tested are not independent after a test and it is not feasible to represent the joint probability for all but the simplest models. Secondly, errors in estimates of $P\{\mathbf{u}_i\}$ propagate through the Bayes formula and are accentuated. So if the initial estimate of $P\{\mathbf{u}_i\}$ based on the similarity of the correspondence is incorrect, the posterior probability becomes worthless soon.

The alternative, pursued here, is to make minimal assumptions about the link between $P\{\mathbf{u}_i\}$ and the similarity function $q(\mathbf{u}_j)$. In particular, we assume monotonicity, i.e.

$$q(\mathbf{u}_i) \geq q(\mathbf{u}_j) \Rightarrow P\{\mathbf{u}_i\} \geq P\{\mathbf{u}_j\}. \quad (1)$$

Sequences of correspondences satisfying

$$i < j \Rightarrow P\{\mathbf{u}_i\} \geq P\{\mathbf{u}_j\} \quad (2)$$

will be called not-worse-than-random.

Note that we are searching for a *single* growth function. It seems possible to adapt the growth function to reflect the result of previous sample-and-test cycles. However, all PROSAC (and RANSAC) runs are alike: a sequence of failures followed by a ‘hit’ due to an all-inlier sample. The history of the sampling process is thus fully captured by t , the number of tests carried so far.

The sampling strategy. Imagine standard RANSAC drawing T_N samples of size m out of N data points. Let $\{\mathcal{M}_i\}_{i=1}^{T_N}$ denote the sequence of samples $\mathcal{M}_i \subset \mathcal{U}_N$ that are uniformly drawn by RANSAC, and let $\{\mathcal{M}_{(i)}\}_{i=1}^{T_N}$ be sequence of the same samples sorted in descending order according to the sample quality

$$i < j \Rightarrow q(\mathcal{M}_{(i)}) \geq q(\mathcal{M}_{(j)}).$$

If the samples are taken in order $\mathcal{M}_{(i)}$, the samples that are more likely to be uncontaminated are drawn earlier. Progressively, samples containing data points with lower quality function are drawn. After T_N samples, exactly all RANSAC samples $\{\mathcal{M}_i\}_{i=1}^{T_N}$ were drawn.

Let T_n be an average number of samples from $\{\mathcal{M}_i\}_{i=1}^{T_N}$ that contain data points from \mathcal{U}_n only

$$T_n = T_N \frac{\binom{n}{m}}{\binom{N}{m}} = T_N \prod_{i=0}^{m-1} \frac{n-i}{N-i}, \text{ then}$$

$$\frac{T_{n+1}}{T_n} = \frac{T_N}{T_N} \prod_{i=0}^{m-1} \frac{n+1-i}{N-i} \prod_{i=0}^{m-1} \frac{N-i}{n-i} = \frac{n+1}{n+1-m}.$$

Finally, the recurrent relation for T_{n+1} is

$$T_{n+1} = \frac{n+1}{n+1-m} T_n. \quad (3)$$

There are T_n samples containing only data points from \mathcal{U}_n and T_{n+1} samples containing only data points from \mathcal{U}_{n+1} . Since $\mathcal{U}_{n+1} = \mathcal{U}_n \cup \{\mathbf{u}_{n+1}\}$, there are $T_{n+1} - T_n$ samples that contain a data point \mathbf{u}_{n+1} and $m-1$ data points drawn from \mathcal{U}_n . Therefore, the procedure that for $n = m \dots N$ draws $T_{n+1} - T_n$ samples consisting of a data point \mathbf{u}_{n+1} and $m-1$ data points drawn from \mathcal{U}_n at random efficiently generates samples $\mathcal{M}_{(i)}$.

As the values of T_n are not integer in general, we define $T'_m = 1$ and

$$T'_{n+1} = T'_n + \lceil T_{n+1} - T_n \rceil. \quad (4)$$

The growth function is then defined as

$$g(t) = \min\{n : T'_n \geq t\}. \quad (5)$$

In PROSAC, the t -th sample \mathcal{M}_t consists of

$$\mathcal{M}_t = \{\mathbf{u}_{g(t)}\} \cup \mathcal{M}'_t, \quad (6)$$

where $\mathcal{M}'_t \subset \mathcal{U}_{g(t)-1}$ is a set of $|\mathcal{M}'_t| = m - 1$ data points drawn from $\mathcal{U}_{g(t)-1}$ at random. The parameter T_N defines after how many samples the behavior of PROSAC and RANSAC becomes identical. In our experiments, the parameter was set to $T_N = 200000$.

$t := 0, n := m, n^* := N$

Repeat until a solution satisfying eqs. (12), (9) is found.

1. Choice of the hypothesis generation set

$t := t + 1$

if $(t = T'_n)$ & $(n < n^*)$ then $n := n + 1$ (see eqn. 4)

2. Semi-random sample \mathcal{M}_t of size m

if $T'_n < t$ then

The sample contains $m - 1$ points selected from \mathcal{U}_{n-1} at random and \mathbf{u}_n

else

Select m points form \mathcal{U}_n at random

3. Model parameter estimation

Compute model parameters p_t from the sample \mathcal{M}_t

4. Model verification

Find support (i.e. consistent data points) of the model with parameters p_t

Select termination length n^* if possible according to Section 2.2

Algorithm 1: Outline of the PROSAC algorithm.

2.2. Stopping criterion

The PROSAC algorithm terminates if the number of inliers I_{n^*} within the set \mathcal{U}_{n^*} satisfies the following conditions:

- *non-randomness* – the probability that I_{n^*} out of n^* data points are by chance inliers to an arbitrary incorrect model is smaller than Ψ (typically set to 5%)
- *maximality* – the probability that a solution with more than I_{n^*} inliers in \mathcal{U}_{n^*} exists and was not found after k samples is smaller than η_0 (typically set to 5%).

From all such solutions the one that causes the termination first is chosen.

The **non-randomness** requirement prevents PROSAC from selecting a solution supported by outliers that are by chance consistent with it. The constraint is typically checked ex-post in standard approaches [1]. The distribution of the cardinalities of sets of random ‘inliers’ is binomial

$$P_n^R(i) = \beta^{i-m}(1-\beta)^{n-i+m} \binom{n-m}{i-m}, \quad (7)$$

where β is the probability, that an incorrect model calculated from a random sample containing an outlier is supported by a correspondence not included in the sample.

We set β pessimistically based on geometric considerations. If needed, the estimate of β can be made more precise during the sampling of PROSAC.

For each n , the minimal number of inliers I_n^{\min} is calculated so that the probability of size of such support being random is smaller than Ψ

$$I_n^{\min} = \min\{j : \sum_{i=j}^n P_n^R(i) < \Psi\}. \quad (8)$$

A non-random solution found on \mathcal{U}_{n^*} must satisfy

$$I_{n^*} \geq I_{n^*}^{\min}. \quad (9)$$

A **maximality** constraint defines how many samples are needed to be drawn to ensure the confidence in the solution and is the (only) termination criterion of RANSAC [4].

For a hypothesis generation set \mathcal{U}_n , the probability P_{I_n} that an uncontaminated sample of size m is randomly selected from a set \mathcal{U}_n of n correspondences is

$$P_{I_n} = \frac{\binom{I_n}{m}}{\binom{n}{m}} = \prod_{j=0}^{m-1} \frac{I_n - j}{n - j} \approx \varepsilon_n^m, \quad (10)$$

where I_n is the number of inliers in \mathcal{U}_n and $\varepsilon_n = I_n/n$ is the fraction of inliers. The probability η of missing a set of inliers of the size I_n on the set \mathcal{U}_n after k samples of PROSAC, where $g(k) \leq n$, is

$$\eta = (1 - P_{I_n})^k. \quad (11)$$

The number of samples that have to be drawn to ensure the probability η falls under the predefined threshold η_0 is

$$k_{n^*}(\eta_0) \geq \log(\eta_0) / \log(1 - P_{I_{n^*}}). \quad (12)$$

The termination length n^* is chosen to minimize $k_{n^*}(\eta_0)$ subject to $I_{n^*} \geq I_{n^*}^{\min}$.

3. Experiments

The not-worse-than-random assumption about the ordering of correspondences was tested for two different similarity functions.

Matching based on SIFT descriptors [5] was used to obtain tentative correspondences in PLANT and MUG experiments². The similarity was defined as the ratio of the distances in the SIFT space of the best and second match. The threshold for the similarity is set to 0.8 as suggested in [5]. This setting has been shown experimentally [5, 13] to provide a high fraction of inliers in the tentative correspondences. However, this thresholding also leads to small absolute number of tentative correspondences.

²The code was kindly provided by David Lowe, UBC, Vancouver, Canada.




Figure 2: The fraction of inliers ε among top n correspondences sorted by quality for the GREAT WALL (left), MUG background (center) and MUG foreground (right) scenes. The circles mark the (average) size of the largest set of correspondences that PROSAC sampled, i.e. the size it sampled when it stopped. The circles are close to the optimal stopping size.



Figure 3: The PLANT scene. Depth discontinuities, self-occlusions and repetitive patterns reduce the probability that a correspondence with high quality (similarity) is indeed an inlier. RANSAC fails to estimate epipolar geometry on this image pair.




Figure 4: The dependence of the fraction of inliers ε on the ordering by the SIFT similarity on the PLANT scene. The first 30 correspondences have similarity below the threshold 0.8. The circle shows the optimal stopping length n^* chosen by PROSAC.

In the GREAT WALL experiment, the Euclidean distance of the first fifteen Discrete Cosine Transform (DCT) coefficients was used as a similarity function [9, 2]. The DCT was computed on normalized, affine invariantly detected, parallelograms. As tentative correspondences, points with mutually best similarity were selected.

Figures 2 and 4 show the dependence of the fraction of inliers ε on the order of a tentative correspondence induced by the similarity function. In all experiments, regardless of the similarity function used, the assumption of not-worse-than-random ordering held. The fraction of inliers ε decreased almost monotonically as a function of the number of tentative correspondences n .

Comparison of efficiency. The number of samples drawn by RANSAC and PROSAC as well as wall clock time of both algorithms were measured on three scenes.

For the **GREAT WALL scene** (Fig. 1) both PROSAC and RANSAC algorithms found 57 inliers among the 250 tentative correspondences. RANSAC needed on average (over hundred runs) 106,534 samples which took 10.76 sec. PROSAC estimated the 57-inlier epipolar geometry after 9 samples (0.06 sec) on average.

The **PLANT scene** is challenging due to a large number of depth discontinuities and the presence of both repetitive (floor) and locally similar (leaves) patterns. Tentative

correspondences are obtained by SIFT matching [5] computed on MSERs [6]. For the 0.8 threshold, the tentative correspondences include only 12 inliers. The epipolar geometry could not be estimated reliably by either RANSAC or PROSAC. When the SIFT threshold was set to 0.95, there were $N = 559$ tentative correspondences and $I = 51$ inliers. In this case, RANSAC fails due to low fraction of inliers $\varepsilon = 9.2\%$; on average, RANSAC would need $8.43 \cdot 10^7$ samples to find the solution, which is not realistic. PROSAC, on the other hand, found all the inliers and estimated the correct epipolar geometry after 3,576 samples in 0.76 sec on average (over 100 execution).

The **MUG scene**, (Fig. 5) is non-rigid. The mug moved between the acquisition of the two images. Tentative correspondences were obtained matching MSER [6] and affine-invariant [7] regions using SIFT descriptors. First, epipolar geometry was estimated on all tentative correspondences. Then, inliers to the first epipolar geometry were removed and another EG was estimated to segment the motion in the scene. The results are summarized in Tab. 1.

In the Fig. 2, the dependency of ε on the ordering is shown for the background segmentation (center) and the foreground (the mug) segmentation (right). All the correct correspondences on the mug are outliers to the background



Figure 5: Motion segmentation. The motion of the background and the foreground (the mug) are denoted by light and black lines respectively.

Background		$N = 783, \varepsilon = 79\%$	
	I	k	time [sec]
PROSAC	617	1.0	0.33
RANSAC	617	15	1.10
Mug		$N = 166, \varepsilon = 31\%$	
	I	k	time [sec]
PROSAC	51.6	18	0.12
RANSAC	52.3	10,551	0.96

Table 1: The number of inliers (I) detected, samples (k) and the time needed in the motion estimation of the background (top) and the foreground (bottom) in the MUG experiment.

	k	min k	max k	time [sec]
RANSAC	106,534	97,702	126,069	10.76
PROSAC	9	5	29	0.06
PROSAC OR	61,263	1,465	110,727	6.28

Table 2: The comparison of the number of samples drawn by RANSAC, PROSAC, and PROSAC with random ordering on the GREAT WALL experiment. The values of k are average, minimum, and maximum over 100 runs respectively.

motion while having high similarity score. This can be observed in Fig. 2(center) as dips in the plot. This also shows that the probability of a correspondence being an inlier to a given geometric model depends on other factors besides the similarity of local descriptors.

Testing robustness to **the worst case situation**, i.e. to random ordering of tentative correspondences. To compare RANSAC and PROSAC in the least favorable situation for PROSAC, an experiment was carried out on sequences of randomly permuted correspondences. The 250 correspon-




Figure 6: Histogram of the stopping length n^* of PROSAC for 100 random orderings in the GREAT WALL scene.

dences from the GREAT WALL experiment were used. For PROSAC, the average number of samples taken over 100 different random permutations was 61, 263 (standard deviation $2.94 \cdot 10^4$). The number of samples ranged from 1, 465 to 110, 727. For RANSAC, the average and standard deviation were 130, 419 and $6.55 \cdot 10^3$ respectively. The results, together with results of PROSAC on tentative correspondences sorted by the similarity function are shown in Tab. 2.

PROSAC drew less samples and was faster than RANSAC in this experiment. The difference in the average number of samples can be attributed to the fact that even in a randomly permuted sequence there are sub-sequences with higher than average inlier fractions, allowing PROSAC to terminate earlier. The histogram of PROSAC termination lengths n^* is plotted in Fig. 6. Only a fraction of PROSAC executions were terminated on the full set of the tentative correspondences, where RANSAC is terminated. However, the maximal number of samples drawn are comparable for PROSAC and RANSAC, i.e. in the worst case PROSAC behaves as RANSAC.

The stopping criterion. In all experiments, the optimal stopping length n^* calculated ex-post was identical with the

length automatically selected by PROSAC. The values of n^* are depicted in Figs 2 and 4 as circles.

4. Conclusions

PROSAC – a novel robust estimator of the hypothesize-and-verify type was introduced. The PROSAC algorithm exploits the ordering structure of the set of tentative correspondences, assuming that the ordering by similarity computed on local descriptors is better than random. The assumption held in all our experiment for both quality measures that were tested.

The sampling on progressively larger subsets consisting of top-ranked correspondences brings very significant computational savings. In comparison to RANSAC, PROSAC was more than hundred time faster on non-trivial problems. Using synthetically generated sequences of correspondences, we showed that the worst-case behavior of PROSAC and RANSAC are effectively identical.

PROSAC removes one parameter of the matching process – the threshold on the similarity function for selection of tentative correspondences. Thus robustness against either having too few correspondences or a large number of tentative correspondences with high outlier percentage is achieved. In one tested problem, PROSAC found a solution of the matching problem that cannot be solved by RANSAC.

References

- [1] M. Brown and D. Lowe. Recognising panoramas. In *Proc. ICCV03*, volume I, pages 1218–1225, October 2003.
- [2] O. Chum, J. Matas, and Š. Obdržálek. Enhancing RANSAC by generalized model optimization. In *Proc. of the ACCV*, volume 2, pages 812–817, January 2004.
- [3] M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *CACM*, 24(6):381–395, June 1981.
- [4] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge University, Cambridge, 2nd edition, 2003.
- [5] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [6] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust wide-baseline stereo from maximally stable extremal regions. *Image and Vision Computing*, 22(10):761–767, Sep 2004.
- [7] K. Mikolajczyk and C. Schmid. An affine invariant interest point detector. In *Proc. ECCV*, volume 1, pages 128–142, 2002.
- [8] D. Nister. Preemptive RANSAC for live structure and motion estimation. In *Proc. ICCV03*, volume I, pages 199–206, October 2003.
- [9] Š. Obdržálek and J. Matas. Image retrieval using local compact DCT-based representation. In *Proc. of DAGM'03*, volume 1 of *LNCS*, pages 490–497. Springer-Verlag, 9 2003.
- [10] P. Pritchett and A. Zisserman. Wide baseline stereo matching. In *Proc. ICCV*, pages 754–760, 1998.
- [11] F. Schaffalitzky and A. Zisserman. Viewpoint invariant texture matching and wide baseline stereo. In *Proc. 8th ICCV, Vancouver, Canada*, July 2001.
- [12] H. Shao, T. Svoboda, T. Tuytelaars, and L. V. Gool. Hpat indexing for fast object/scene recognition based on local appearance. In *computer lecture notes on Image and video retrieval*, LNCS 2728, pages 71–80. Springer, July 2003.
- [13] J. Sivic and A. Zisserman. Video Google: A text retrieval approach to object matching in videos. In *Proceedings of the International Conference on Computer Vision*, pages 1470 – 1477, Oct. 2003.
- [14] B. Tordoff and D. Murray. Guided sampling and consensus for motion estimation. In *Proc. 7th ECCV*, volume 1, pages 82–96. Springer-Verlag, 2002.
- [15] P. H. S. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *CVIU*, 78:138–156, 2000.
- [16] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. A comprehensive survey of bundle adjustment in computer vision. In *Proc. Vision Algorithms: Theory and Practice. International Workshop on Vision Algorithms*, number 1883 in LNCS, pages 298–372. Springer Verlag, 1999.
- [17] T. Tuytelaars and L. Van Gool. Wide baseline stereo matching based on local, affinely invariant regions. In *Proc. 11th BMVC*, 2000.
- [18] Z. Zhang, R. Deriche, O. Faugeras, and Q.-T. Luong. A robust technique for matching two uncalibrated images through the recovery of the unknown epipolar geometry. *Artificial Intelligence, December 1995*, 78:87–119, 1995.