

VSAC: Efficient and Accurate Estimator for H and F

Maksym Ivashchkin¹, Daniel Barath², and Jiri Matas¹

¹ Centre for Machine Perception, Czech Technical University in Prague, Czech Republic

² Computer Vision and Geometry Group, Department of Computer Science, ETH Zürich

{ivashmak, matas}@cmp.felk.cvut.cz dbarath@ethz.ch

Abstract

We present VSAC, a RANSAC-type robust estimator with a number of novelties. It benefits from the introduction of the concept of independent inliers that improves significantly the efficacy of the dominant plane handling and, also, allows near error-free rejection of incorrect models, without false positives. The local optimization process and its application is improved so that it is run on average only once. Further technical improvements include adaptive sequential hypothesis verification and efficient model estimation via Gaussian elimination. Experiments on four standard datasets show that VSAC is significantly faster than all its predecessors and runs on average in 1-2 ms, on a CPU. It is two orders of magnitude faster and yet as precise as MAGSAC++, the currently most accurate estimator of two-view geometry. In the repeated runs on EVD, HPatches, PhotoTourism, and Kusvod2 datasets, it never failed.

1. Introduction

The Random Sample Consensus (RANSAC) algorithm introduced by Fischler and Bolles [14] is one of the most popular robust estimators in computer science. The method is widely used in computer vision, its applications include stereo matching [33, 35], image mosaicing [15], motion segmentation [33], 3D reconstruction, detection of geometric primitives, and structure and motion estimation [28].

The textbook version of RANSAC proceeds as follows: random samples of minimal size sufficient to estimate the model parameters are drawn repeatedly. Model consistency with input data is evaluated, *e.g.*, by counting the points closer than a manually set inlier-outlier threshold. If the current model is better than the *so-far-the-best*, it gets stored. The procedure terminates when the probability of finding a better model falls below a user-defined level. Finally, the estimate is polished by least-squares fitting of inliers.

Many modifications of the original algorithm have been proposed. Regarding sampling, PROSAC [8] exploits an a priori predicted inlier probability rank. NAPSAC [27] sam-

ples in the neighborhood of the first, randomly selected, point. Progressive NAPSAC [2] combines both and adds gradual convergence to uniform spatial sampling.

In textbook RANSAC, the model quality is measured by its support, *i.e.*, the number of inliers, points consistent with the model. MLESAC [34] introduced a quality measure that makes it the maximum likelihood procedure. To avoid the need for a user-defined noise level, MINPRAN [32] and A-contrario RANSAC [13] select the inlier-outlier threshold so that the inliers are the least likely to occur at random. Reflecting the inherent uncertainty of the threshold estimate, MAGSAC [5] marginalizes the quality function over a range of noise levels. MAGSAC++ [4] proposes an iterative re-weighted least-squares optimization of the *so-far-the-best* model with weights calculated from the inlier probability of points. The Locally Optimized RANSAC [9] refines the *so-far-the-best* model using a non-minimal number of points, *e.g.*, by iterated least-squares fitting. Graph-Cut RANSAC [3], in its local optimization, exploits the fact that real-world data tend to form spatial structures. The model evaluation is usually the most time-consuming part as it depends both on the number of models generated and the number of input data points. A quasi-optimal speed-up was achieved by the Sequential Probability Ratio Test (SPRT) [25] that randomizes the verification process itself.

In many cases, points in degenerate configuration affect the estimation severely. For example, correspondences lying on a single plane is a degenerate case for F estimation. DEGENSAC [11] detects such cases and applies the plane-and-parallax algorithm. USAC [29] was the first framework integrating many of the mentioned techniques, including PROSAC, SPRT, DEGENSAC, and LO-RANSAC.

In this paper, we present VSAC¹, a RANSAC-type estimator that exploits a number of novelties. It is significantly faster than all its predecessors, and yet as precise as MAGSAC++, the currently most accurate method both in our experiments and according to a recent survey [23]. The accuracy reaches, or is very near, the geometric error of the

¹VSAC has multiple novelties and we found no natural abbreviation reflecting them. We chose "V" as the letter following "U", as in USAC.

ground truth, estimated by cross-validation. For homography \mathbf{H} and epipolar geometry \mathbf{F} estimation, VSAC runs on average in 1-2 ms (on a CPU) on all datasets, two orders of magnitude faster than MAGSAC++. In the repeated runs on datasets EVD [26], HPatches [1], PhotoTourism [31] and Kusvod2 [19], it never failed.

Moreover, VSAC is able to reject non-matching image pairs, with a zero false positive rate on hundreds of random image pairs and a zero false negative rate on pairs from the above-mentioned datasets. The ability is underpinned by a novel concept of *independent random inliers* in the contrario context. We show that if dependent random inliers, e.g. spatially co-located points, are not counted, the support of random models follows very closely a Poisson distribution with a single parameter λ that is easy to estimate reliably² for the given pair. The easily calculated CDF of Poisson raised to the power of the number of evaluated models provides the probability that a certain model quality was reached by chance. VSAC thus provides two confidence measures together with its result. The first is the classical one – the probability that RANSAC returned the model with the highest support. The second is the confidence that the returned solution was not obtained by chance.

The concept of independent random inliers plays a critical role in VSAC’s accuracy and robustness. Experiments show that most failures of USAC-like methods for \mathbf{F} estimation occur in the presence of a dominant plane, despite the DEGENSAC algorithm. In such cases, there are only few out-of-plane inliers, and due to structures in the outliers, incorrect models with high support exist. Removing the dependent structures addresses the problem. Further improvements of dominant plane handling include a heuristic guess of the calibration matrix allowing to deal with fully planar scenes and detects pure rotation. If the guess is wrong, the support reveals it and nothing but a microsecond is lost.

The speed of VSAC is achieved with several technical improvements. Most significantly, we attack the problem of expensive local optimization. In the LO-RANSAC paper [9], the authors prove that the local optimization is run at most $\log(k)$ times, where k is the number of iteration. Nevertheless, despite $\log(k) \ll k$, the complex local optimization may end up being the efficiency bottleneck. We show that a fast local optimization combined with a *single complex final optimization* leads to a faster, yet equally precise algorithm. Moreover, by detecting the intersection over union of *so-far-the-best* and the current set of inliers and by not optimizing similar models, an algorithm is obtained that runs the local optimization on average about once and almost always fewer than two times.

²For geometric problems, the Poisson distribution is a tight approximation of the binomial. Moreover, only the mean λ of independent random inlier counts needs to be estimated, instead of T (number of trials) and δ (success probability) for the binomial.

Further speed up is gained by adaptive SPRT. We use the estimated expected number of random inliers to tune SPRT [25] to the outlier density of the processed pair. We also measure, on the fly, the actual time of model estimation and model verification on the given hardware at the given moment, which is needed for calculating the quasi-optimal thresholds of the SPRT.

To broaden its application potential, VSAC provides novel outputs. Employing the highly efficient Lindstrom method for triangulation [20], it obtains the point pair exactly fitting the returned \mathbf{F} that minimizes the geometric error. VSAC can be thus employed for noise filtering. VSAC returns all input points sorted by the residual, allowing the user to set his own ex-post trade-off between the density, and possibly spread, of points on the one hand, and their accuracy on the other.

2. Detecting Random Models

One issue of RANSAC-like robust estimators is the inability to recognize failures. The estimator always returns a model maximizing some quality function, e.g. the inlier count, even if the tentative inliers stem from outlier structures – sets of neighboring data points that do not originate from the sought model manifold and significantly affect the quality function when considered inliers. In such cases, the returned model might have a reasonably large number of inliers while being inconsistent with the underlying scene geometry. See Fig. 1 for examples.

In this section, we propose a new approach for detecting failures. Towards that end, we differentiate between *independent* and *dependent* inliers. This split is conceptual, helping exposition – in the a contrario calculation of the probability, we pick one point in a group of structured points and count it as an inlier arising by chance, it is an independent random inlier. The other inliers in the structure are ignored, since their inlier status is not a random event, but rather a consequence of their spatial dependence and the fact that the independent inlier is consistent. A *non-random* model must have a sufficient number of *independent* inliers.

We define data point \mathbf{p} a *dependent* inlier if its point-to-model residual is smaller than the inlier-outlier threshold and one of the following conditions hold.

1. Point \mathbf{p} is in the minimal sample used for estimating the model parameters. In such cases, the point will have zero residual by definition.
2. Point \mathbf{p} is close to an independent inlier \mathbf{q} , $\|\mathbf{p} - \mathbf{q}\| \rightarrow 0$. In such cases, points \mathbf{p} and \mathbf{q} form a spatial structure that affects the model quality significantly. Thus, only point \mathbf{q} is considered independent random. Other points from the structure, e.g., \mathbf{p} , are dependent inliers.

These conditions are valid for general data points and model to be estimated. In case of estimating epipolar geometry

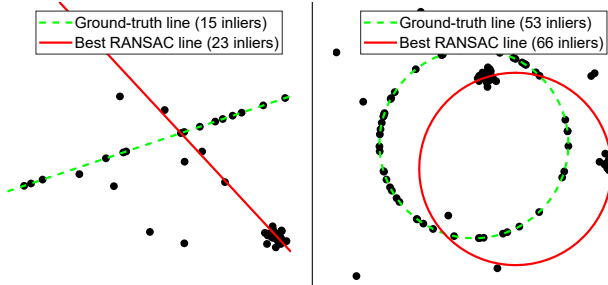


Figure 1: Examples of random models found by RANSAC with a large number of dependent inliers.

from point correspondences, we define the following additional conditions as well.

3. A correspondence $(\mathbf{p}, \mathbf{p}')$ where \mathbf{p} or \mathbf{p}' is close to the epipole in the corresponding image is considered dependent since $(\mathbf{p}, \mathbf{p}')$ always satisfies epipolar constraint $\mathbf{p}'^T \mathbf{F} \mathbf{p} = 0$. This stems from the fact that $\mathbf{F} \mathbf{p} \simeq \mathbf{F} \mathbf{e} = \mathbf{0}$ if $\|\mathbf{p} - \mathbf{e}\| \rightarrow 0$, where \mathbf{e} is the epipole in the first image. The same holds in the second one.
4. Correspondence $(\mathbf{p}, \mathbf{p}')$ is a dependent inlier if it does not pass the chirality check [10].
5. Let $(\mathbf{l}, \mathbf{l}')$ be the corresponding epipolar lines of an independent inlier correspondence. All correspondences that are closer to lines $(\mathbf{l}, \mathbf{l}')$ than the inlier-outlier threshold are considered dependent.

Data points that have a point-to-model residual smaller than the inlier-outlier threshold and do not satisfy any of the previous conditions are considered independent inliers.

To decide if a model estimated by RANSAC is random and inconsistent with the underlying scene geometry and, thus, should be considered failure, we use the univariate theory [30]. Suppose that we are given N models estimated inside RANSAC during its run and the corresponding numbers of random independent inliers $I_1, I_2, \dots, I_N \in \mathbb{N}$. The number of I points being consistent with a random model follows the binomial distribution. The sequence of inliers is an *i.i.d.* random variable with cumulative binomial distribution $C_B(T, \delta)$, where T is the number of points and δ is the probability that a point is an independent inlier to a bad (random) model. The distribution of $I_{\max} = \max\{I_1, I_2, \dots, I_N\}$ over N models is $C_B(T, \delta)^N$. In order to recognize a good (non-random model) with confidence $p \rightarrow 1$, the following condition must hold.

$$C_B(I_{\max}; T, \delta)^N \geq p. \quad (1)$$

In our experiments, we found that probability δ is fairly low. Therefore, even a few independent inliers are enough to decide whether a model is random or not. In this case, the binomial distribution can be approximated by Poisson distribution which is faster to compute. The only parameter of

the Poisson distribution is $\lambda = T \delta$ that represents the mean number of independent inliers.

Finding independent inliers is relatively expensive and, therefore, it is computationally prohibitive if done for every model generated inside RANSAC. Hence, we estimate parameter λ from the first $n \ll N$ generated models. Parameter λ is the mean number of independent inliers consistent with a bad model. In RANSAC, all models are considered bad that have fewer inliers than the *so-far-the-best* one. However, good models that are slightly worse than *so-far-the-best* can corrupt the estimation of λ . Parameter λ is found as follows: first, we discard the *so-far-the-best* model and the ones that have significantly overlapping inlier sets with it (measured via Jaccard similarity [18]) from the generated n models. Second, $\tilde{\lambda}$ is robustly found via taking the median of *independent* inlier counts of n models. Afterwards, we compute 95% percentile ($I_{95}^{\%}$) of the Poisson distribution with parameter $\tilde{\lambda}$ to filter out models with high supports. Eventually, $\hat{\lambda}$ is estimated as the average of inlier numbers lower than $I_{95}^{\%}$.

Finally, to decide if the model returned by RANSAC is non-random and, thus, should be accepted, the cumulative distribution of $(\max\{I_1, \dots, I_N\}; \hat{\lambda})^N$ is calculated, and the model is considered a good one if (1) holds.

3. DEGENSAC⁺

Chum *et al.* [11] proposed the DEGENSAC algorithm that handles estimation, in an uncalibrated setup, of epipolar geometry in scenes with a dominant plane. Such scenes often appear in real-world scenarios, especially, in a man-made environment containing, *e.g.*, building facades. In brief, after a minimal sample is selected to estimate the model parameters, DEGENSAC first checks whether correspondences in the sample are co-planar. If at least five of them are, the corresponding \mathbf{H} is estimated and, used in \mathbf{F} calculation by the plane-and-parallax algorithm [16].

When experimenting with DEGENSAC, we found the following issues that we address in this section. First, we observed that having five or six correspondences consistent with a homography does not necessarily mean that the estimated fundamental matrix is degenerate. This stems from the fact that the inlier-outlier threshold, used for selecting co-planar correspondences via thresholding the re-projection error, is usually wide enough to consider slightly curved surfaces to be planar. Second, the plane-and-parallax algorithm applied for recovering \mathbf{F} from a homography and two additional correspondences may fail, *e.g.*, when the scene is entirely planar and when a large number of random inliers supports an incorrect \mathbf{F} .

To solve these issues, we propose the DEGENSAC⁺ algorithm by incorporating the inlier randomness criteria in the procedure. Both of the above-mentioned problems are solved by checking if the estimated epipolar geometry has

a reasonable number of independent out-of-plane inliers. This builds on the observation that if a fundamental matrix is degenerate due to a dominant planar structure, it only has out-of-plane inliers by chance – otherwise, it would not be degenerate. Thus, it is enough to check the number of independent inliers that are inconsistent with the homography.

Calibrated DEGENSAC⁺. The plane-and-parallax algorithm that is used in DEGENSAC to recover from a degenerate sample assumes that a reasonable number of out-of-plane correspondences exists and the camera motion is not rotation-only. In case one of these conditions does not hold, the algorithm fails while often requiring many RANSAC iterations to recognize this failure.

To solve this issue, we propose to use the intrinsic calibration matrices $\mathbf{K}_1, \mathbf{K}_2 \in \mathbb{R}^{3 \times 3}$ normalizing the coordinates of the co-planar correspondences in the minimal sample before estimating homography \mathbf{H} . In such cases, \mathbf{H} can be decomposed to rotation $\mathbf{R} \in \text{SO}(3)$ and translation $\mathbf{t} \in \mathbb{R}^3$ [24] that can be then used to find the sought non-degenerate fundamental matrix. When decomposing \mathbf{H} , we are given two candidate rotations \mathbf{R}_a and \mathbf{R}_b and translations $\pm \mathbf{t}_a, \pm \mathbf{t}_b$. Since the relative pose can only be estimated up-to-scale, $-\mathbf{t}_a$ and $-\mathbf{t}_b$ can be rejected, and the candidate solutions are $(\mathbf{R}_a, \mathbf{t}_a)$ and $(\mathbf{R}_b, \mathbf{t}_b)$. The solution is selected with the maximum inlier count. The final fundamental matrix is $\mathbf{F} = \mathbf{K}_2^{-\top} [\hat{\mathbf{t}}]_{\times} \hat{\mathbf{R}} \mathbf{K}_1^{-1}$, where $\hat{\mathbf{R}} \in \{\mathbf{R}_a, \mathbf{R}_b\}$ and $\hat{\mathbf{t}} \in \{\mathbf{t}_a, \mathbf{t}_b\}$. With this solution, the rotation-only cases are straightforwardly detected.

In many cases, the calibration is not known a priori but can be approximated. Setting the principal points to coincide with the image centers is a widely used approach. Estimating the focal length is, however, a more challenging problem. We found that testing a range of candidate focal lengths to find a reasonable approximation leads to a reliable \mathbf{F} estimate while still being faster than the original plane-and-parallax RANSAC applied inside DEGENSAC.

The proposed algorithm incorporating the non-randomness criteria proposed in the previous section and also the calibration matrix is shown in Alg. 1.

4. Adaptive SPRT

The Sequential Probability Ratio Test (SPRT) proposed by Matas *et al.* [25] aims at speeding up the robust estimation procedure by addressing the problem that, in RANSAC, a large number of models are verified, *e.g.* their support is calculated, even if they are unlikely to be better than the previous *so-far-the-best*. The time spent on these models is wasted. The SPRT test is based on Wald’s theory of sequential decision making. It interrupts the model verification when the probability of that particular model being a *good* one falls below a user-defined threshold.

SPRT has four user-defined parameters, *i.e.*, the initial probability of a correspondence being consistent with a

Algorithm 1: DEGENSAC⁺

Input: $\hat{\mathbf{F}}^*$ – *so-far-the-best* fundamental matrix; \mathcal{S} – minimal sample initializing $\hat{\mathbf{F}}^*$, I_F – initial guess for independent non-planar support.

- 1 $\hat{\mathbf{H}}^* := \text{estimateHomography}(\hat{\mathbf{F}}^*, \mathcal{S})$
- 2 **if** $\hat{\mathbf{H}}^* = \emptyset$ **then**
- 3 **return:** $\hat{\mathbf{F}}^*$ // no homography
- 4 **if** \mathbf{K} is given **then**
- 5 **assert** $(\mathbf{K}^{-1} \hat{\mathbf{H}}^* \mathbf{K}$ is not conj. to rotation)
- 6 $\hat{\mathbf{F}} := \text{findF}(\hat{\mathbf{H}}^*, \mathbf{K})$
- 7 **else**
- 8 $\hat{\mathbf{F}} := \text{findF}(\hat{\mathbf{H}}^*, \hat{\mathbf{K}})$ // use approx. $\hat{\mathbf{K}}$
- 9 **if** $\text{support}(\hat{\mathbf{F}}^*) \geq I_F$ **then**
- 10 **return:** $\hat{\mathbf{F}}^*$ // not degenerate
- 11 **if** \mathbf{K} is given **then**
- 12 **return:** $\hat{\mathbf{F}}$ // cannot be degenerate
- 13 **if** $\text{support}(\hat{\mathbf{F}}') \geq I_F$ **then**
- 14 **return:** $\hat{\mathbf{F}}'$
- 15 // re-estimate support I_F
- 16 $\hat{\mathbf{F}}'', I_F := \text{planeAndParallax}(\hat{\mathbf{H}}^*)$
- 17 **if** $\text{support}(\hat{\mathbf{F}}') \geq I_F$ or $\text{support}(\hat{\mathbf{F}}'') \geq I_F$ **then**
- 18 **return:** $\hat{\mathbf{F}}'$ or $\hat{\mathbf{F}}''$
- 19 **return:** \emptyset // reject $\hat{\mathbf{F}}^*$

good (ϵ_0) and a *bad* model (δ_0); avg. number of estimated models (\bar{m}_S); time to estimate the model parameters (t_M). The actual parameters that lead to the fastest procedure are challenging to find manually and require the user to acquire knowledge about the problem at hand. Even the architecture of the computer impacts the minimal solver and point verification times that should be considered when setting t_M . To avoid the manual setting, we propose the Adaptive SPRT (A-SPRT) algorithm that finds the optimal SPRT parameters in a data- and architecture-dependent manner.

The model estimation time depends not only on the computer architecture but, also, on the actual solver and error metric being used. Parameter t_M is calculated for free by measuring the model estimation and point verification run-times in the first n RANSAC iterations.

The average number of models \bar{m}_S is found as the average number of *valid* models per sample in the first n iterations.

Inlier probabilities ϵ_0 and δ_0 are estimated from the average number $\hat{\lambda}$ of inliers consistent with a bad model that is estimated from in the first n RANSAC iterations as $\delta_0 = \hat{\lambda} / T$, see Section 2. Parameter δ_0 is the probability of a point being an inlier and $\hat{\lambda}$ is the mean of the corresponding binomial distribution $\mathcal{B}(T, \delta_0)$. From δ_0 , we approximate the maximum number of inliers I_δ of a *bad* model as a high quantile (*e.g.*, 0.99) of the normal distribution with the same mean and standard deviation as \mathcal{B} . The approximation

is as follows:

$$I_\delta = \hat{\lambda} + 3.719 \sqrt{\hat{\lambda}(1 - \delta_0)} \quad (2)$$

The initial probability ϵ_0 of a correspondence being inlier of a *good* model is found as $\epsilon_0 = \max(I_\delta, \hat{I}^*) / T$, where \hat{I}^* is the inlier number of the *so-far-the-best* model.

Failure case. If probabilities δ_0 and ϵ_0 are similar, the original SPRT often rejects *good* models leading to increased run-time or, in extreme situations, total failure. To solve this issue, we propose to apply A-SPRT only if

$$\frac{1}{1 - \alpha} t_v^w E^w(T) < t_v T, \quad (3)$$

where t_v^w and t_v are the times for verifying a single correspondence, respectively, with and without SPRT and α is the probability of a false rejection [25], and $E^w(T)$ is the average number of points verified.

5. Model Accuracy

Simple and Fast Local Optimization. The key idea of the local optimization (LO) proposed in [9] is to address the fact that not all all-inlier samples lead to accurate models due to the noise in the data. We however observed that the LO step has a slightly different role in practice – finding a model that is good enough to trigger the termination criterion of RANSAC early. The final model accuracy depends mostly on the optimization procedure, *e.g.* least-squares fitting or numerical optimization, applied once, after the RANSAC main loop finished. Therefore, the LO can be made light-weight without compromising the final model accuracy.

In our formulation, the primary objective is to find a light-weight LO procedure that runs swiftly and is applied only when it likely leads to termination. To do so, we introduce the following conditions that control when the LO is applied. They are as follows:

1. *Non-random model.* The *so-far-the-best* model has the required number of independent inliers defined by (2).
2. *Low Jaccard similarity.* The local optimization is applied only if the inlier sets of the new best and previous best models has a lower than 0.95 Jaccard index, *i.e.*, the intersection over union. This condition is motivated by the tendency that if a model is just slightly different from the previous *so-far-the-best*, the LO step likely does not refine it significantly, but the final optimization does the main improvement.

The proposed LO procedure is shown in Alg. 2. It is important to note that in this procedure, larger-than-minimal samples are selected that is typically avoided in RANSAC due to increasing the problem complexity and, thus, the number of iterations required to provide probabilistic guarantees of finding the sought model parameters. In Alg. 2,

Algorithm 2: Local Optimization

Input: \mathcal{I}^* – inliers of *so-far-the-best* model,
 τ^* – score of *so-far-the-best* model, MAX_ITERs – maximum iterations of LO, s – sample size.

```

1  $\theta_{LO}^* := \emptyset$ ;  $I_{max} := |\mathcal{I}^*|$ 
2 for  $t := 0$ ;  $t < \text{MAX\_ITERs}$ ;  $t++$  do
3    $\theta_{LO} := \text{estimate}(\text{subset}(\mathcal{I}^*, \min\{|\mathcal{I}^*|, s\}))$ 
4    $\tau_{LO}, \mathcal{I}_{LO} := \text{evaluate}(\theta_{LO})$ 
5   if  $\tau^* < \tau_{LO}$  & non-degenerate( $\theta_{LO}$ ) then
6      $\mathcal{I}^*, \tau^*, \theta_{LO}^* := \mathcal{I}_{LO}, \tau_{LO}, \theta_{LO}$ 
7     if  $I_{max} < |\mathcal{I}_{LO}|$  then
8        $I_{max} := |\mathcal{I}_{LO}|$ 
9 return:  $\theta_{LO}^*, \tau^*, \text{MaximumIterations}(I_{max})$ 

```

the sample is selected from a set of points that likely are inliers. Therefore, the increased sample size does not affect the accuracy and processing time negatively.

We found experimentally the parameters that suit for two-view geometric problems, minimizing the total run-time while maintaining the accuracy. For **F** estimation the optimal sample size is 21 and the number of iterations is 20. For **H**, the optimal size is 32 and number of iterations is 10.

Final Optimization. Since the proposed local optimization does not intend to make the *so-far-the-best* model as accurate as possible, the final model polishing step should return the best possible model. In textbook RANSAC [14], the final model parameters are obtained by running a single least-squares fitting on the set of inliers returned by RANSAC.

In our experiments, we test two types of final optimizations. The first is an iterative LSQ fitting that re-selects the inliers. This improves the model parameters extremely efficiently. The second one is the iteratively re-weighted LSQ approach proposed in [4] that does not require a single inlier-outlier threshold, only its loose upper bound. This is, in practice, slightly slower than traditional iterative LSQ. Due to applying it only once, the deterioration in the run-time is at most 0.2-1.0 milliseconds in our experiments.

Minimal Model Estimation often requires finding the null-space of a linear system, typically, by SVD. The traditionally used solvers for **H/F/E** (homography, fundamental, essential) matrix estimation use null space parameterization either to directly solve the problem (**H** from 4 point correspondences) or to find the coefficients of some polynomials that are then solved (**F** from 7 matches; **E** from 5 matches) [16]. It is however slow when having large matrices. Given that it is applied in every RANSAC iteration, it affects the total run-time severely. Instead of SVD, we suggest to use Gaussian Elimination (GE) as in [6]. While GE is less stable numerically in theory, it has many advantages, *e.g.*, an order-of-magnitude speed-up, easy-to-implement, and lower memory complexity than SVD. The marginally

decreased stability of the minimal solver does not affect the final accuracy.

6. Point Correction and Ranking

RANSAC outputs the model with the highest support and the corresponding set of inliers. Some applications, *e.g.*, 3D reconstruction or bundle adjustment, rely heavily on the obtained set of inliers, and it is thus important to rank them according to the quality. A simple way is to sort the inliers by their residuals in an increasing order.

Correcting the inliers by making them consistent with the estimated model is an extremely important problem, *e.g.*, for improving data points provided as the ground truth. For instance, such points are obtained by careful manual selection in a number of real-world datasets, *e.g.* KUSVOD2 and ADELAIIDRMF [36], that leads to correct, however, inevitably noisy points. Even if we assume that the annotator did a perfect work and all found points are supposedly noise-free, the discrete nature of photography (*i.e.*, the scene is projected to a grid of pixels) prevents having perfect points with no noise. Also, correcting points is very useful for the user who does not want to run bundle adjustment. We therefore propose to minimize the noise by correcting a point so it has a zero error distance (*e.g.*, reprojection distance for \mathbf{H} or geometric error for \mathbf{F}) to the final model.

Homography. A way to correct correspondences is by introducing “half” homography \mathbf{A} , where $\mathbf{H} = \mathbf{A}\mathbf{A}$ [12]. The middle point averaging out the points in this reference frame is calculated as follows: $\mathbf{m} = (\phi(\mathbf{A}\mathbf{x}) + \phi(\mathbf{A}^{-1}\mathbf{x}'))/2$, where ϕ is a mapping, normalizing a point by its homogeneous coordinate. The points projected to the manifold are $\tilde{\mathbf{x}} \sim \mathbf{A}^{-1}\mathbf{m}$ and $\tilde{\mathbf{x}}' \sim \mathbf{A}\mathbf{m}$. The up-to-scale relation \sim can be removed via mapping ϕ . The corrected correspondence $(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}')$ has zero error to \mathbf{H} , because the elimination of \mathbf{m} from the two equations implies $\tilde{\mathbf{x}}' \sim \mathbf{A}\mathbf{A}\tilde{\mathbf{x}} = \mathbf{H}\tilde{\mathbf{x}}$. The square root \mathbf{A} exists if the eigenvalues of \mathbf{H} have positive real parts. Planar homographies from image datasets satisfy this condition in the experiments. In general, a homography should be checked before applying the proposed point correction.

Epipolar geometry. The corrected points must lie *perfectly* on the epipolar lines. A fast procedure was presented by Lindstrom in [21]. Moreover, if the intrinsic camera matrices are known, [21] enables to efficiently obtain the triangulated 3D points as well. Since [21] corrects the correspondences even they are incorrect matches, we use the oriented epipolar constraint [10] to remove some of the incorrect ones. Results are put in the appendix.

7. Experiments

To test VSAC and each of the new techniques proposed in this paper, we have downloaded the EVD (11 pairs; avg.

Problem		# of RANSAC iterations			
		$\sim 10^2$	$\sim 10^3$	$\sim 10^4$	$\sim 10^5$
H	w	100%	100%	100%	100%
	w/o	96%	93%	86%	81%
F	w	100%	99%	99%	100%
	w/o	85%	82%	79%	78%

Table 1: Percentage of detected failures by the proposed criterion (1) with (w) and without (w/o) removing dependent inliers on homography \mathbf{H} and fundamental matrix \mathbf{F} estimation problems when trying to match image pairs without a common field-of-view, *i.e.*, they do not match. In total, 500 image pairs tested.

Dataset	Method	Error (px)	t_{total} (ms)	t_{DEG} (ms)
PhotoTourism	DEG	0.64	54.7	19.8
	DEG ⁺	0.44	31.6	1.7
Kusvod2	DEG	2.34	19.4	5.7
	DEG ⁺	1.83	11.6	1.4

Table 2: Avg. geometric error (px), RANSAC run-time (t_{total} ; ms), and time for recovering from degenerate solutions (t_{DEG} ; ms) of DEGENSAC⁺ and DEGENSAC [11] on image pairs from datasets KUSVOD2 (7 pairs) and PhotoTourism [19, 31] (100) where 1/3 of the correspondences are consistent with a homography, *i.e.*, dominant plane.

inlier ratio: 0.24 ± 0.16) [26] and HPATCHES (142; ratio: 0.51 ± 0.24) [1] datasets for \mathbf{H} , and the KUSVOD2 (7; ratio: 0.52 ± 0.18) [19] and PHOTOTOURISM (500 pairs are chosen randomly; ratio: 0.37 ± 0.12) [31] for \mathbf{F} estimation.

DEGENSAC⁺ is compared to DEGENSAC on images from the KUSVOD2 and PHOTOTOURISM datasets where at least 1/3 of the correspondences are consistent with a homography, *i.e.*, dominant plane. For PHOTOTOURISM, we used the ground truth camera calibration. For KUSVOD2, the intrinsic parameters are approximated as proposed in Section 3. The avg. error (px), overall RANSAC time, and the run-times of the DEGENSAC versions (ms) are reported in Table 2. DEGENSAC⁺ leads to the lowest errors while being faster than the original DEGENSAC algorithm.

Local Optimization criteria proposed in Section 5 are tested on the four downloaded datasets. Table 3 reports the average number of samples drawn, *so-far-the-best* models encountered and LO runs. While the original LO-RANSAC [9] applies LO always when a new *so-far-the-best* model is found, the proposed criteria leads to running LO significantly less often, *i.e.*, once per problem, causing a significant speed-up with no deterioration in the accuracy.

Dependent inlier removal helps in modelling the probability of a model being random as shown in Fig. 2. We estimated the parameters of the Poisson distribution (see Section 2) on two scenes from the PHOTOTOURISM and EVD

Dataset	Sample #	Sftb #	LO #
HPatches	77.6	4.41	1.00
EVD	534.1	5.66	1.20
PhotoTourism	14.4	3.49	1.00
Kusvod	442.7	4.07	1.08

Table 3: The avg. number of samples drawn, *so-far-the-best* (sftb) models encountered and LO runs on four datasets by VSAC. For the original LO-RANSAC, the numbers of sftb and LO runs are identical. VSAC runs notably fewer LO.

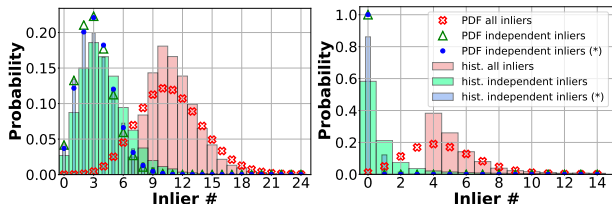


Figure 2: Histograms of the numbers of all (red bar) and independent (green bar) inliers of models estimated in RANSAC; and the expected inlier numbers from the Poisson distribution with its parameters calculated from all (red crosses) and from the independent inliers (green triangle) on a scene from EVD (left) an PHOTOTOURISM (right) datasets. The blue histogram (*) reports inlier counts of models generated from corrupted samples. Thus, the blue PDF shows the true probability of points being consistent with a bad model.

datasets from all inliers (red crosses) and, also, from only the independent ones (green triangles). We also show the histogram of the actual inlier numbers of models estimated inside RANSAC. The green histogram is closer to the values of the corresponding Poisson distribution than the red one. Thus, removing dependent inliers helps in recognizing and rejecting random models.

Failure detection results are shown in Table 1. Images with no common scene were matched with SIFT [22] detector having, thus, only incorrect correspondences, and therefore, no good model exists. We calculated the percentage of cases when the proposed criterion (1) detects if the currently matched image pair has no common field-of-view, *i.e.*, they do not match, with and without removing the dependent inliers. Criterion (1) with removing the dependent inliers almost *always* detects if two images do not match.

Adaptive SPRT is compared to the original algorithm proposed in [25] and, also, to textbook RANSAC without preemptive verification on the HPATCHES and PHOTOTOURISM datasets. We chose a total of 480 image pairs that have inlier ratio lower than 0.5. This was done to test the techniques on cases where RANSAC is required to do many iterations and, thus, the preemptive model verification is essential.

The CDFs of the estimation errors (in pixels) and pro-

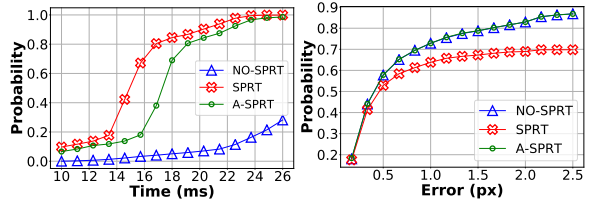


Figure 3: CDFs of the processing times (ms) and geometric errors (px) of the proposed A-SPRT, SPRT, and RANSAC.

cessing times (in milliseconds) are shown in Fig. 3. Being accurate is interpreted by a curve close to the top-left corner. It can be seen that, while the proposed A-SPRT leads to similar processing time as the original one (left plot), it is significantly more accurate (right).

Geometric Accuracy and Speed. The proposed VSAC incorporating A-SPRT, Calibrated DEGENSAC⁺, fast LO, PROSAC, and final covariance-based polishing (VSAC) or IRLS from [4] (VSAC_{MGSC}) is compared with the following *state-of-the-art* robust estimators. (1) USACv20 – framework from [17] with SPRT, GC-RANSAC, DEGENSAC and P-NAPSAC. (2) default OpenCV’s RANSAC implementation. (3) USAC framework from [29] with SPRT, LO-RANSAC, DEGENSAC and PROSAC. (4) GC-RANSAC from [3] with SPRT, PROSAC and DEGENSAC. (5) MAGSAC++ from [4] with P-NAPSAC and DEGENSAC. (6) ORSA – RANSAC with a contrario approach [13]. (7) Neural-Guided RANSAC [7] for epipolar geometry.

The inlier-outlier threshold is set to 2.5 pixels for \mathbf{H} and 1.5 pixels for \mathbf{F} ; and the confidence to 99%. The maximum number of iterations of the main loop of RANSAC is 3000 for homography and 5000 for fundamental matrix estimation. Table 4 shows that the proposed VSAC is *always* the fastest methods with a large margin. Its maximum run-time (t_{\max}) is 13 ms on a wide range of problems. The VSAC_{MGSC} with the IRLS from [4] as a final model optimization is slightly more accurate than VSAC while being marginally slower. Both VSAC and VSAC_{MGSC} have comparable accuracy to the *state-of-the-art* methods with many times being the most accurate methods.

For all methods, the source code provided by the authors were used. The methods are implemented in C++ except NG-RANSAC. The weight prediction part of NG-RANSAC runs in Python and CUDA on GPU. After the weight prediction, the rest of the code is in C++. All methods were run on the same computer, except NG-RANSAC. We ran it on a different machine with a GPU.

8. Conclusions

This paper presented VSAC, a novel robust geometry estimator. We introduced the concept of independent inliers that helps detecting incorrect models, and it is also

	Method	Time (milliseconds)				Error (pixels)				
		t_{med}	t_{avg}	t_{max}	$w_{\%}^t$	ε_{med}	ε_{avg}	ε_{max}	$w_{\%}^{\varepsilon}$	
Homography	HPatches (142 pairs)	VSAC	① 0.8	① 0.9	① 3.0	① 97	0.50	0.79	6.48	6
		VSAC_{MGSC}	1.8	② 1.9	② 6.9	0	① 0.43	① 0.61	① 3.47	16
		USACv20	2.8	3.3	9.1	0	0.51	8.16	509.19	6
		USAC	22.7	26.7	92.9	0	0.62	13.24	509.19	8
		OpenCV	② 1.6	7.7	43.8	② 3	0.58	1.99	141.99	9
		GC	174.3	157.8	337.6	0	① 0.43	② 0.62	② 4.34	② 23
		MGSC++	20.3	44.1	987.5	0	② 0.48	0.68	8.08	① 24
		ORSA	436.2	794.6	4692.7	0	0.82	105.69	4226.99	8
		Cross-validation error on the ground truth points:					0.63	0.63	1.15	
	EVD (11 pairs)	VSAC	① 0.5	① 0.7	① 2.9	① 100	3.18	3.54	7.65	② 18
		VSAC_{MGSC}	② 0.7	② 0.9	② 3.2	0	① 2.73	① 3.40	② 7.59	9
		USACv20	2.3	4.4	12.5	0	3.37	3.67	8.04	① 25
		USAC	7.9	11.4	29.6	0	7.23	121.44	474.01	3
		OpenCV	16.9	21.1	40.1	0	4.02	5.36	16.17	5
		GC	31.6	33.8	52.3	0	② 2.78	② 3.48	10.42	① 25
		MGSC++	28.4	27.0	63.2	0	3.48	3.78	① 7.54	13
		ORSA	56.8	68.1	218.9	0	148.79	169.66	438.45	4
		Cross-validation error on the ground truth points:					1.80	2.21	6.33	
Fundamental matrix	PhotoTour (500 pairs)	VSAC	① 2.4	① 2.5	① 6.0	① 99	① 0.15	① 0.16	② 0.64	12
		VSAC_{MGSC}	② 3.0	② 3.2	② 7.7	0	① 0.15	② 0.17	0.70	12
		USACv20	20.0	29.4	100.3	0	0.17	0.20	3.37	10
		USAC	6.3	6.7	17.5	0	0.43	0.60	9.91	2
		OpenCV	170.0	149.3	280.1	0	0.38	0.63	10.00	1
		GC	224.2	256.9	696.4	0	0.17	0.19	0.93	9
		MGSC++	216.1	273.8	10764.9	0	② 0.16	② 0.17	0.93	17
		ORSA	84.0	98.9	594.0	0	① 0.15	① 0.16	0.95	① 19
		NG-RSC	120.6	121.9	459.2	0	① 0.15	② 0.17	① 0.55	② 18
	Cross-validation error on the ground truth points:					0.06	0.06	0.16		
	Kusvod2 (15 pairs)	VSAC	① 1.4	① 2.5	① 12.7	① 45	0.53	② 1.06	① 6.37	11
		VSAC_{MGSC}	1.8	3.0	② 13.1	0	0.53	① 1.00	② 6.39	9
		USACv20	2.3	7.6	75.4	3	0.56	2.85	54.52	12
		USAC	② 1.5	② 2.6	48.3	① 45	2.11	3.37	36.46	3
		OpenCV	9.0	44.3	198.0	② 7	1.02	4.52	38.11	0
		GC	189.7	214.5	893.0	0	0.54	1.49	35.66	① 23
		MGSC++	47.4	84.4	1551.9	0	② 0.50	2.54	47.92	② 17
		ORSA	32.4	62.0	605.6	0	0.53	15.43	255.56	10
NG-RSC		105.1	110.4	599.6	0	① 0.48	3.27	39.84	② 17	
Cross-validation error on the ground truth points:					0.91	1.12	2.34			

Table 4: Comparison of speed and accuracy of the RANSAC methods. Average t_{avg} , median t_{med} , and the maximum t_{max} running time over all runs, including 10 repetitions for EVD, HPatches and Kusvod2, on four datasets, measured on a standard CPU, with the exception of the NG-RSC methods, that requires a GPU. Top and the second best results are marked and highlighted. Similarly, we report average ε_{avg} , median ε_{med} and the maximum ε_{max} error of ground truth points. To account for the randomization, the percentage of top results is given in the $w_{\%}^t$, $w_{\%}^{\varepsilon}$ columns. The errors of the top methods are very close to the accuracy of ground truth points, estimated using leave-one-out. For Kusvod2, the best methods are significantly more precise than the ground truth, making the small differences meaningless.

at the core of the new DEGENSAC⁺ method that returns a non-degenerate fundamental matrix supported by a sufficient number of non-planar independent inliers. The LO process was reformulated and sped-up while the novel final optimization is responsible for high geometric accuracy. Technical enhancements as adaptive-SPRT with automatic

parameter settings and GE for the minimal model solver provide further speed up. VSAC is as geometrically precise as the best *state-of-the-art* methods.

Acknowledgement. This research was supported by the Research Center for Informatics (project CZ.02.1.01/0.0/0.0/16_019/0000765 funded by OP VVV and by the Grant Agency of the Czech Technical University in Prague, grant No. SGS20/171/OHK3/3T/13).

References

- [1] Vassileios Balntas, Karel Lenc, Andrea Vedaldi, and Krystian Mikolajczyk. Hpatches: A benchmark and evaluation of handcrafted and learned local descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. **2, 6**
- [2] Daniel Barath, Maksym Ivashechkin, and Jiri Matas. Progressive NAPSAC: sampling from gradually growing neighborhoods. *arXiv preprint arXiv:1906.02295*, 2019. **1**
- [3] Daniel Barath and Jiří Matas. Graph-Cut RANSAC. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6733–6741, 2018. <https://github.com/danini/graph-cut-ransac>. **1, 7**
- [4] Daniel Barath, Jana Noskova, Maksym Ivashechkin, and Jiri Matas. Magsac++, a fast, reliable and accurate robust estimator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. <https://github.com/danini/magsac>. **1, 5, 7**
- [5] Daniel Barath, Jana Noskova, and Jiří Matas. MAGSAC: marginalizing sample consensus. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. <https://github.com/danini/magsac>. **1**
- [6] Hamid Bazargani, Olexa Bilaniuk, and Robert Laganière. A fast and robust homography scheme for real-time planar target detection. *J. Real-Time Image Process.*, 15(4):739–758, Dec. 2018. **5**
- [7] Eric Brachmann and Carsten Rother. Neural-Guided RANSAC: Learning where to sample model hypotheses. In *International Conference on Computer Vision*, 2019. <https://github.com/vislearn/ngransac>. **7**
- [8] O. Chum and J. Matas. Matching with PROSAC-progressive sample consensus. In *Computer Vision and Pattern Recognition*. IEEE, 2005. **1**
- [9] O. Chum, J. Matas, and J. Kittler. Locally optimized RANSAC. In *Joint Pattern Recognition Symposium*. Springer, 2003. **1, 2, 5, 6**
- [10] O. Chum, T. Werner, and J. Matas. Epipolar geometry estimation via RANSAC benefits from the oriented epipolar constraint. In *International Conference on Pattern Recognition*, 2004. **3, 6**
- [11] Ondrej Chum, Tomas Werner, and Jiri Matas. Two-view geometry estimation unaffected by a dominant plane. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 772–779. IEEE, 2005. **1, 3, 6**
- [12] Edvin Deadman, Nicholas J. Higham, and Rui Ralha. Blocked schur algorithms for computing the matrix square root. In Pekka Manninen and Per Öster, editors, *Applied Parallel and Scientific Computing*, pages 171–182, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. **6**
- [13] Ferran Espuny, Pascal Monasse, and Lionel Moisan. A new a contrario approach for the robust determination of the fundamental matrix. In *Image and Video Technology – PSIVT 2013 Workshops*, pages 181–192, 2014. https://github.com/pmoulon/IPOL_AC_RANSAC. **1, 7**
- [14] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981. **1, 5**
- [15] D. Ghosh and N. Kaabouch. A survey on image mosaicking techniques. *Journal of Visual Communication and Image Representation*, 2016. **1**
- [16] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, USA, 2 edition, 2003. **3, 5**
- [17] M. Ivashechkin, D. Barath, and J. Matas. USACv20: robust essential, fundamental and homography matrix estimation. In *Computer Vision Winter Workshop*, 2020. **7**
- [18] Paul Jaccard. The distribution of the flora of the alpine zone. volume 11, pages 37–50, 1912. **3**
- [19] K. Lebeda, J. Matas, and O. Chum. Fixing the locally optimized RANSAC. In *British Machine Vision Conference*. Citeseer, 2012. <http://cmp.felk.cvut.cz/wbs/>. **2, 6**
- [20] Peter Lindstrom. Triangulation made easy. pages 1554–1561, 06 2010. **2**
- [21] Peter Lindstrom. Triangulation made easy. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1554–1561, 06 2010. **6**
- [22] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60:91–110, 2004. **7**
- [23] Jiayi Ma, Xingyu Jiang, Aoxiang Fan, Junjun Jiang, and Junchi Yan. Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision*, 129(1):23–79, 2021. **1**
- [24] Ezio Malis and Manuel Vargas. Deeper understanding of the homography decomposition for vision-based control. 01 2007. **4**
- [25] Jiri Matas and Ondrej Chum. Randomized RANSAC with sequential probability ratio test. In *Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1*, volume 2, pages 1727–1732. IEEE, 2005. **1, 2, 4, 5, 7**
- [26] D. Mishkin, J. Matas, and M. Perdoch. MODS: Fast and robust method for two-view matching. *Computer Vision and Image Understanding*, 2015. **2, 6**
- [27] D. R. Myatt, P. H. S. Torr, S. J. Nasuto, J. M. Bishop, and R. Craddock. NAPSAC: high noise, high dimensional robust estimation. In *In BMVC02*, pages 458–467, 2002. **1**
- [28] David Nistér. Preemptive ransac for live structure and motion estimation. In *ICCV*, pages 199–206. IEEE Computer Society, 2003. **1**
- [29] R. Raguram, O. Chum, M. Pollefeys, J. Matas, and J-M. Frahm. USAC: a universal framework for random sample consensus. *Transactions on Pattern Analysis and Machine Intelligence*, 2013. <http://www.cs.unc.edu/~rraguram/usac/USAC-1.0.zip>. **1, 7**
- [30] Debra (Dallie) Sandilands. *Univariate Analysis*, pages 6815–6817. Springer Netherlands, 2014. **3**
- [31] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Photo tourism: Exploring photo collections in 3d. In *ACM SIGGRAPH 2006 Papers*, SIGGRAPH '06, page 835–846, New York, NY, USA, 2006. Association for Computing Machinery. **2, 6**

- [32] Charles V. Stewart. Minpran: A new robust estimator for computer vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(10):925–938, 1995. 1
- [33] P. H. S. Torr and D. W. Murray. Outlier detection and motion segmentation. In *Optical Tools for Manufacturing and Advanced Automation*. International Society for Optics and Photonics, 1993. 1
- [34] P. H. S. Torr and A. Zisserman. MLESAC: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 2000. 1
- [35] P. H. S. Torr, A. Zisserman, and S. J. Maybank. Robust detection of degenerate configurations while estimating the fundamental matrix. *Computer Vision and Image Understanding*, 1998. 1
- [36] Hoi Sim Wong, Tat-Jun Chin, Jin Yu, and David Suter. Dynamic and hierarchical multi-structure geometric model fitting. In *Int. Conf. Comput. Vis.*, 2011. 6