

# ANSAC: Adaptive Non-Minimal Sample and Consensus

Victor Fragoso<sup>1</sup>  
victor.fragoso@mail.wvu.edu

Christopher Sweeney<sup>2</sup>  
csweeney@cs.washington.edu

Pradeep Sen<sup>3</sup>  
psen@ece.ucsb.edu

Matthew Turk<sup>3</sup>  
mturk@cs.ucsb.edu

<sup>1</sup> West Virginia University  
Morgantown, WV, USA

<sup>2</sup> University of Washington  
Seattle, WA, USA

<sup>3</sup> University of California  
Santa Barbara, CA, USA

---

## Abstract

While RANSAC-based methods are robust to incorrect image correspondences (outliers), their hypothesis generators are not robust to correct image correspondences (inliers) with positional error (noise). This slows down their convergence because hypotheses drawn from a minimal set of noisy inliers can deviate significantly from the optimal model. This work addresses this problem by introducing ANSAC, a RANSAC-based estimator that accounts for noise by adaptively using more than the minimal number of correspondences required to generate a hypothesis. ANSAC estimates the inlier ratio (the fraction of correct correspondences) of several ranked subsets of candidate correspondences and generates hypotheses from them. Its hypothesis-generation mechanism prioritizes the use of subsets with high inlier ratio to generate high-quality hypotheses. ANSAC uses an early termination criterion that keeps track of the inlier ratio history and terminates when it has not changed significantly for a period of time. The experiments show that ANSAC finds good homography and fundamental matrix estimates in a few iterations, consistently outperforming state-of-the-art methods.

## 1 Introduction

The robust estimation of geometrical models from pairs of images (*e.g.*, homography, essential and fundamental matrices) is critical to several computer vision applications such as structure from motion, image-based localization, and panorama stitching [0, 1, 2, 3, 4, 5, 6, 7, 8, 9]. Several applications use robust estimators to compute these transformations from candidate image correspondences that contain outliers (*i.e.*, incorrect image correspondences). While RANSAC-based estimators are robust to outliers, they lack a hypothesis generator that can produce an accurate model from a minimal set of noisy inliers (*i.e.*, correct image correspondences with positional error). Since existing hypothesis generators are sensitive to the noise in the data, the generated hypotheses can deviate significantly from the optimal model, yielding a delay in the estimator's convergence.

Although there exist hypothesis-generation improvements [1, 2, 3, 4, 5, 6] that accelerate the convergence of RANSAC-based estimations, there is no previous work addressing the noise in the correspondences *at the hypothesis-generation stage*. The effect of noisy

inliers in a hypothesis generated from a minimal sample (a set with the least amount of correspondences to generate a hypothesis) can be reduced by using more samples (non-minimal). This is because more inliers provide additional constraints that help the estimator to produce a more accurate model. However, using a non-minimal sample to generate a hypothesis increases the chances of adding an outlier, which can lead to wrong hypotheses. For this reason, existing methods [9, 17, 24] that tackle the effect of noisy inliers operate as a refinement stage rather than operating directly in the hypothesis generation phase. While these methods reduce the effect of noise in hypotheses, they can increase the computational overhead since they run an additional inner-refinement stage in the hypothesis-and-test loop of a RANSAC-based estimator.

To reduce the noise effect on hypotheses produced in the hypothesis-generation phase, this work presents a novel Adaptive Non-Minimal Sample and Consensus (ANSAC) robust estimator. Unlike previous work [9, 8, 9, 13, 14, 16, 17, 24, 31], ANSAC produces hypotheses from non-minimal sample sets in the hypothesis generation stage of a RANSAC-based estimator. The homography and fundamental matrix estimation experiments show that ANSAC returns a good estimate consistently and quickly.

## 1.1 Related Work

The key element of existing methods that accelerate the estimation process is their ability to assess the correctness of every candidate correspondence. Different approaches compute this correctness by means of probabilities, rankings, and heuristics. The probability-based methods (*e.g.*, Guided-MLESAC [31], BEEM [16], BLOGS [9], and EVSAC [24]) aim to calculate the probability of correctness of each candidate correspondence using the matching scores. They use then the computed probabilities to form a discrete distribution over the candidate correspondences to sample and generate hypotheses.

Instead of randomly sampling candidate correspondences from a probability distribution, PROSAC [8] creates an initial subset with candidate correspondences ranked by quality. PROSAC generates a hypothesis by randomly sampling a minimal number of candidate correspondences from this subset. PROSAC iteratively expands the subset by including lower ranked candidate correspondences and keeps sampling a minimal number of candidate correspondences from this subset until convergence.

While previous methods accelerate the estimation process, they still use minimal samples, which are sensitive to noise. To alleviate this issue, Chum *et al.* [9, 17] proposed LO-RANSAC, an estimator that adds a refinement stage after finding a good hypothesis. This refinement stage is another RANSAC process in which non-minimal samples are generated from the inliers supporting the new best hypothesis.

In contrast to the previous methods, ANSAC uses subsets with a high predicted inlier ratio to draw *non-minimal samples* and generate hypotheses that account for noise at the *hypothesis-generation phase*. Thus, ANSAC does not require an inner-RANSAC refinement process as in LO-RANSAC, which can be expensive [17]. Consequently, it avoids the extra computational overhead that an inner refinement process incurs. The key element in ANSAC is to predict the inlier ratio of the subsets to determine the size of a non-minimal sample.

## 2 ANSAC

The goal of ANSAC is to increase the likelihood of finding an accurate hypothesis in the early iterations by generating hypotheses from outlier-free non-minimal samples. To this end, ANSAC prioritizes the use of subsets with high estimated inlier ratios to generate hypotheses. The inlier ratio can be considered as the probability of selecting a correct correspondence from a set at random, since it is the fraction of correct correspondences in a given




Figure 1: Overview of the pipeline in ANSAC. First, it ranks the correspondences by a correctness quality value. Then, ANSAC builds subsets, and estimates the inlier-ratio prior using a correctness quality value for every subset. Then, the hypothesis-and-test loop uses the subsets to draw hypotheses and tests them to produce a final estimate. ANSAC uses a Kalman filter to refine the inlier-ratio estimate of the current subset from which it generates hypotheses. To do so, the filter combines the inlier ratio prior and the one calculated by testing hypotheses. The non-minimal sampler uses the refined inlier ratio to generate a non-minimal or minimal sample. ANSAC stops when a termination criterion is met.

set. As such, ANSAC uses the inlier ratio to assess the odds of producing an outlier-free non-minimal sample, since it measures the purity of a set. The estimation pipeline of ANSAC (illustrated in Fig. 1) has two main stages: initialization and hypothesis-and-test loop.

## 2.1 Initialization

The main goal of the initialization is to compute subsets of correspondences and their estimated inlier ratio priors. This stage has three main components: correspondence ranking, subset computation, and inlier ratio prior calculation. Given the candidate image correspondences, the initialization stage outputs the ranked subsets  $\{\mathcal{S}_i\}_{i=1}^{n-m+1}$  and their initial inlier ratio estimates/priors  $\{\varepsilon_i^0\}_{i=1}^{n-m+1}$ , where  $n$  is the total number of candidate correspondences, and  $m$  is the minimal number of correspondences to generate a hypothesis.

### 2.1.1 Ranking Correspondences and Subset Generation

The main goal of this stage is to rank high the correspondences that are likely to be correct. This is because ANSAC aims to generate hypotheses from outlier-free samples as early as possible. Thus, ANSAC first calculates their correctness or quality value (e.g., using SIFT ratio [20] or Meta-Recognition Rayleigh (MRR) [13]) and ranks the correspondences. Then, ANSAC builds an initial subset with the minimal number of best-ranked correspondences to draw a hypothesis. This initial subset is the base to generate the remaining subsets. ANSAC generates each subsequent subset by adding the next ranked correspondence to the previously generated subset. In this way, ANSAC keeps growing the most recently generated subset until it uses all the correspondences. The initialization stage computes the  $i$ -th subset  $\mathcal{S}_i$  efficiently by including all the top  $j$  ranked correspondences. Mathematically, the  $i$ -th subset is  $\mathcal{S}_i = \{(\mathbf{x} \leftrightarrow \mathbf{x}')_j : 1 \leq j < m + i, 1 \leq i \leq n - m + 1\}$ , where  $(\mathbf{x} \leftrightarrow \mathbf{x}')_j$  is the  $j$ -th ranked candidate correspondence.

### 2.1.2 Computing the Inlier Ratio Prior

The last step in the initialization stage is the calculation of the inlier ratio prior for every subset. These inlier ratio priors are key to initialize a Kalman filter that will refine the inlier ratio estimate of every subset over time; see Sec. 2.2.1.

The inlier ratio  $\varepsilon_i$  of  $\mathcal{S}_i$  is defined as the probability of drawing a correct correspondence from it, i.e.,  $\varepsilon_i = \Pr(y = +1 \mid \mathcal{S} = \mathcal{S}_i)$ , where  $y \in \{+1, -1\}$  is a random variable indicating correct or incorrect correspondence, respectively. This probability can be computed as follows:  $\varepsilon_i = \frac{1}{|\mathcal{S}_i|} \sum_{j=1}^{|\mathcal{S}_i|} I(y_j = +1)$ , where  $I(e)$  is the indicator function which returns 1 when  $e$  is true, and 0 otherwise. The ideal indicator function is an oracle that identifies correct correspondences, which is challenging to obtain. Nevertheless, it can be approximated by mapping the correctness or ranking quality values to the  $[0, 1]$  range. For instance, using

**Algorithm 1** ANSAC

---

**Require:** The correspondences  $\mathcal{C} \leftarrow \{(\mathbf{x} \leftrightarrow \mathbf{x}')_{j=1}^n\}$

**Ensure:** Estimate  $H^*$

- 1:  $\{(\mathcal{S}_i, \varepsilon_i^0)\}_{i=1}^{n-m+1} \leftarrow \text{Initialize}(\mathcal{C})$
- 2: Initialize:  $s \leftarrow 1, \varepsilon_s \leftarrow \varepsilon_s^0, \varepsilon'_s \leftarrow \varepsilon_s, \varepsilon \leftarrow 0, \varepsilon_H \leftarrow 0, L \leftarrow \infty$  // See Section 2.1.
- 3: **for** iteration = 1 **to**  $L$  **do**
- 4:  $(\mathcal{X}, M) \leftarrow \text{GenerateSample}(\mathcal{S}_s, \varepsilon_s)$  // Each subset generates at most  $M$  hypotheses according to Eq. (8).
- 5:  $H \leftarrow \text{GenerateHypothesis}(\mathcal{X})$
- 6:  $(\varepsilon_H, \varepsilon'_{s,H}) \leftarrow \text{TestHypothesis}(H)$
- 7: **if**  $\varepsilon_H > \varepsilon$  // The new hypothesis has more inliers than the current best hypothesis. **then**
- 8:  $H^* \leftarrow H, \varepsilon \leftarrow \varepsilon_H, \varepsilon'_s \leftarrow \varepsilon'_{s,H}$
- 9:  $L \leftarrow \text{UpdateMaxIterations}(\varepsilon)$  // See Sec. 2.2.4.
- 10: **end if**
- 11:  $s \leftarrow \text{UpdateSubset}(\text{iteration}, M)$  //  $s \leftarrow s + 1$  when  $M$  hypotheses were generated, else  $s \leftarrow s$ .
- 12:  $\varepsilon_s \leftarrow \text{RefineInlierRatio}(\mathcal{S}_s, \varepsilon_s')$
- 13: **if**  $\text{Terminate}(\varepsilon)$  **then**
- 14:  $\text{Terminate loop}$
- 15: **end if**
- 16: **end for**

---

MRR [13] probabilities or mapping Lowe’s ratio to the  $[0, 1]$  range using a radial basis function. The initialization stage computes the inlier ratio prior  $\varepsilon_i^0$  for every subset  $\mathcal{S}_i$  using this approximation to the indicator function  $\hat{I}$ .

## 2.2 Hypothesis-and-test Loop

The goal of this stage is to robustly estimate a model from the correspondences as quickly as possible. To this end, this stage iteratively draws non-minimal samples when possible from the ranked subsets  $\{\mathcal{S}_i\}_{i=1}^{n-m+1}$  and estimates their inlier ratios  $\{\varepsilon_i\}_{i=1}^{n-m+1}$ . Similar to PROSAC [8], ANSAC progressively utilizes all the subsets according to the rank in order to generate hypotheses. However, unlike PROSAC, which generates hypotheses from minimal samples, ANSAC estimates the subsets’ inlier ratios to decide the size of the non-minimal sample for hypothesis generation. Thus, since ANSAC assumes that the subsets at the top of the ranking are less contaminated it will likely use most of the candidate correspondences and generate hypotheses from non-minimal samples when possible. As ANSAC goes through the lowest-ranked subsets, it will likely use minimal samples to generate hypotheses. Similar to PROSAC, ANSAC uses a new subset once it reached the maximum number  $M$  of hypotheses it can generate from the current subset. However, unlike PROSAC, which uses a recursive formula to determine the maximum number  $M$  of hypotheses to draw from a given subset, ANSAC uses the estimated inlier ratio of a given subset to calculate  $M$ .

Specifically, ANSAC operates as detailed in Alg. 1. Given the input subsets  $\{\mathcal{S}_i\}_{i=1}^{n-m+1}$  and their inlier ratio priors  $\{\varepsilon_i^0\}_{i=1}^{n-m+1}$ , ANSAC sets its current-subset index  $s$  pointing to the first subset, *i.e.*,  $s \leftarrow 1$ . Subsequently, it generates a sample  $\mathcal{X}$  from the current subset  $\mathcal{S}_s$  in step 4. However, unlike existing estimators that always draw minimal samples, ANSAC uses the current inlier ratio estimate  $\varepsilon_s$  to determine the size of the sample; see Sec. 2.2.2. Given the non-minimal sample, ANSAC generates a hypothesis  $H$  in step 5. In step 6, the hypothesis tester identifies two sets of correspondences: 1) those that support the generated hypothesis  $H$  across all the input correspondences to compute its overall inlier ratio  $\varepsilon_H$ ; and those in the current subset to calculate the subset’s inlier ratio  $\varepsilon'_{s,H}$ . When the inlier ratio  $\varepsilon_H$  is larger than the inlier ratio  $\varepsilon$  of the current best hypothesis, then ANSAC updates several variables (steps 7-10): the current best hypothesis  $H^* \leftarrow H$ , the overall inlier ratio  $\varepsilon \leftarrow \varepsilon_H$ , the “observed” inlier ratio  $\varepsilon'_s \leftarrow \varepsilon'_{s,H}$  (the fraction of correspondences from the current subset

$\mathcal{S}_s$  supporting the new best hypothesis  $H^*$ ), and the maximum number  $M$  of hypotheses to draw from  $\mathcal{S}_s$ . In step 11, ANSAC uses a new subset by increasing  $s \leftarrow s + 1$  if it has generated  $M$  hypotheses from the current subset. Then, a Kalman filter refines the inlier ratio estimate  $\varepsilon_s$  of the current subset in step 12; see Sec. 2.2.1. Lastly, ANSAC checks if the current best hypothesis  $H^*$  satisfies a termination criteria and stops if satisfied (steps 13-15).

### 2.2.1 Estimating the Inlier Ratio with a Kalman Filter

The Kalman filter [27, 30] iteratively refines the estimate of the inlier ratio of the current subset  $\varepsilon_s$ . It does so by combining the inlier ratio prior  $\varepsilon_s^0$  and the observed inlier ratio  $\varepsilon'_s$ . This is a crucial component since the inlier ratio is the value that determines the size of the sample and the progressive use of subsets. To use this filter, ANSAC assumes that the inlier ratio varies linearly as a function of the ranked subsets and it has Gaussian noise.

The prediction step of the Kalman filter estimates the inlier ratio  $\varepsilon_k^k$  at time  $k$  given the inlier ratio  $\varepsilon_s^{k-1}$  at previous time and the inlier ratio prior  $\varepsilon_s^0$ . The filter uses the following linear inlier ratio “motion” model:

$$\varepsilon_{k|k-1} = \alpha_k \varepsilon_s^{k-1} + \beta_k \varepsilon_s^0 \quad (1)$$

$$\sigma_{k|k-1}^2 = \alpha_k^2 \sigma_{s,k-1}^2 + \sigma_p^2, \quad (2)$$

where  $\varepsilon_{k|k-1}$  is the predicted inlier ratio,  $\sigma_{k|k-1}$  is the standard deviation for the predicted inlier ratio,  $\alpha_k$  and  $\beta_k$  are weights that control the contribution of their respective inlier ratio terms,  $\sigma_{s,k-1}$  is the standard deviation of the previous estimate  $\varepsilon_s^{k-1}$ , and  $\sigma_p$  is a standard deviation parameter that enforces stochastic diffusion. The weights  $\alpha_k$  and  $\beta_k$  must satisfy  $1 = \alpha_k + \beta_k$  since the inlier ratios range between  $[0, 1]$ .

To initialize the filter, ANSAC sets  $\varepsilon_s^{k-1} \leftarrow \varepsilon_s^0$  when it is the first iteration. However, when ANSAC uses a new subset, then ANSAC sets  $\varepsilon_s^{k-1} \leftarrow \varepsilon_{s-1}$ , where  $\varepsilon_{s-1}$  is the estimate of the previous set. This is a good initial value since subsets only differ by one sample. ANSAC uses the following function to set  $\alpha_k = \alpha^{I(s \leftarrow s+1)}$ , where  $I(s \leftarrow s+1)$  returns 1 when a new current subset is used, and 0 otherwise; and the parameter  $\alpha$  satisfies  $0 \leq \alpha \leq 1$ . This function forces the filter to use the inlier ratio prior only when ANSAC uses a new subset. This is because the inlier ratio of a subset can be refined over multiple iterations and the prior must contribute once to the refinement.

The update step aims at refining the output of the prediction step. To this end, it combines the observed inlier ratio  $\varepsilon'_s$ , which is obtained by testing the best current hypothesis on the current subset. The update step performs the following computations:

$$g = \frac{\sigma_{k|k-1}^2}{\sigma_{k|k-1}^2 + \sigma_u^2} \quad (3)$$

$$\varepsilon_s^k = \varepsilon_{k|k-1} + g(\varepsilon'_s - \varepsilon_{k|k-1}) \quad (4)$$

$$\sigma_{s,k}^2 = (1 - g)\sigma_{k|k-1}^2, \quad (5)$$

where  $\sigma_u^2$  is a parameter enforcing stochastic diffusion,  $g$  is the Kalman gain,  $\varepsilon_s^k$  is the refined inlier ratio for the current subset at time  $k$ , and  $\sigma_{s,k}^2$  is its standard deviation. At the end of this stage, ANSAC updates the inlier ratio of the current subset  $\varepsilon_s = \varepsilon_s^k$ .

### 2.2.2 Adaptive Non-minimal Sampling

Unlike many RANSAC variants (e.g., [3, 8, 11, 13, 14, 16, 23, 24]) which use minimal samples to generate hypotheses, ANSAC computes the size  $q$  of a sample as follows:

$$q(\varepsilon_s) = (1 - \lambda(\varepsilon_s))q_{\min} + \lambda(\varepsilon_s)q_{\max}, \quad (6)$$

where  $q_{\min} = m$  is the minimum sample size to generate a hypothesis;  $q_{\max}$  is the maximum sample size, which is upper-bounded by the size of the current subset  $\mathcal{S}_s$ ; and

$$\lambda(\varepsilon_s; \omega, \mu) = (1 + \exp(-\omega(\varepsilon_s - \mu)))^{-1}, \quad (7)$$

is a logistic curve with parameters  $\omega$  (steepness) and  $\mu$  (inflection point). The value constraints for these parameters are:  $0 < \omega$  and  $0 < \mu < 1$ .

Eq.(6) selects a sample size by linearly interpolating between  $q_{\min}$  and  $q_{\max}$ . The sample size depends directly on measuring the likelihood of not including outliers, which is measured by  $\lambda$  (see Eq.(7)). When the likelihood is high ( $> 0.5$ ), ANSAC produces a non-minimal sample. Otherwise, it produces minimal samples. ANSAC evaluates this function at every iteration since  $\varepsilon_s$  evolves over time.

### 2.2.3 Maximum Number of Hypotheses per Subset

ANSAC generates a maximum number  $M$  of hypotheses from the current subset  $\mathcal{S}_s$ . ANSAC considers the worst case scenario where the inlier ratio of the current subset is low and thus the sample size is the minimal. Thus, it calculates  $M$  using the following formula [□, □]:

$$M = \left\lceil \frac{\log(1 - \nu)}{\log(1 - p)} \right\rceil, \quad (8)$$

where  $\nu$  is the probability of producing a good hypothesis, and  $p = (\varepsilon'_s)^{q_{\min}}$  is the probability of a minimal sample being outlier-free. ANSAC adapts  $M$  as a function of  $\varepsilon'_s$  every iteration. Eq. (8) provides the least amount of hypotheses drawn from minimal samples that an estimator needs to test to achieve a confidence  $\nu$  [□, □]. The progressive use of the subsets in ANSAC depends directly on the estimated inlier ratio of each subset, unlike PROSAC [□] which uses a growth function that depends on the number of samples drawn by each subset.

### 2.2.4 Early Termination

The classical termination criterion of a RANSAC-based estimator calculates the maximum number of iterations adaptively using Eq. (8), since one iteration generates and tests a hypothesis. Even though that an estimator finds a good hypothesis in the early iterations, the estimator still iterates until it reaches the maximum number of iterations calculated with Eq. (8). In this case, the iterations after finding a good hypothesis are unnecessary.

To alleviate this problem, many approaches aim to detect if a hypothesis is likely to be bad [□, □, □] to skip the iteration, and thus avoid testing the hypothesis. Unlike these methods, ANSAC aims to detect if the estimator is unlikely to find a significantly better hypothesis than the current best one, and terminate the loop. To this end, ANSAC uses the history of the overall inlier ratio  $\varepsilon$  throughout of the iterations. Since ANSAC uses ranked correspondences to generate hypotheses, it is expected to see a rapid increase of the overall inlier ratio  $\varepsilon$  in the early iterations, and a slow or none increase in the later iterations. The proposed termination criterion detects when  $\varepsilon$  has reached the low-or-none increase zone.

The proposed early termination criterion analyzes the latest  $T$  increments of the overall inlier ratio  $\{\delta_\varepsilon(x_l)\}_{l=1}^T$ , where  $\delta_\varepsilon(l) = \varepsilon(l) - \varepsilon(l-1)$  and  $\varepsilon(l)$  is the overall inlier ratio at the  $l$ -th iteration. ANSAC decides to stop iterating if their average variation  $\bar{\delta}_\varepsilon$  is below a certain threshold  $\gamma$ . The value of  $T$  is bounded by the overall maximum number of iterations  $L$ , which is calculated using  $p = \varepsilon^m$  and Eq. (8). Thus,  $T$  has to change adaptively as well. For this reason, ANSAC uses  $T = \tau L$ , where  $0 < \tau < 1$ , to set the value for  $T$ .




Figure 2: **(a)** Accuracy of the subset-inlier-ratio estimation for Homography (top row) and Fundamental matrix (bottom row) models. The filter tends to provide a close estimation (gray curve) to the ground truth (black curve). **(b)** Accuracy of the estimation using a bad inlier-ratio prior. The filter is robust and tends to provide a close estimation in this scenario.

### 3 Experiments

This section presents a series of experiments, each demonstrating the effectiveness and efficiency of the proposed components in ANSAC. We implemented ANSAC in C++ and use the robust estimation API from the TheiaSfM library [29]<sup>1</sup>. We used the implementations of Root-SIFT [2] and CasHash [6] in TheiaSfM to match features and compute symmetric correspondences: common image correspondences obtained by matching image one as a reference and image two as a query and vice versa.

**Datasets.** We used the Oxford affine-covariant-regions dataset [22] and USAC homography-image pairs [23] for homography estimation. In addition, we used the USAC fundamental-matrix-image pairs and Strecha dataset [22, 28] for fundamental matrix estimation. The Oxford and Strecha dataset provide information to calculate ground truth. To get a ground truth for USAC dataset, we used the same procedure used by Raguram *et al.* [23], which estimates the models by running RANSAC  $10^7$  times, followed by a visual inspection of the image correspondences to ensure the models did not include outliers. See supplemental material for a visualization of the image pairs from these datasets.

#### 3.1 Kalman Filter for Inlier Ratio Estimation

The goal of this experiment is twofold: 1) measure the subset-inlier-ratio estimation accuracy of the Kalman filter; and 2) evaluate the robustness of the filter given a poor inlier ratio prior estimate. We used the Lowe’s “radius” as a correspondence quality:  $r^2 = r_1^2 + r_2^2$ , where  $r_1$  and  $r_2$  are the Lowe’s ratios obtained by matching image one to image two and viceversa. We used  $\sigma_p = \sigma_u = 0.1$  for stochastic diffusion, set  $\alpha = 0.5$  for the prediction step (see Eq. (1)), and removed the early termination criterion for this experiment.

Fig. 2 (a) shows the estimation accuracy results for homography and fundamental matrix estimation. It presents the ground truth inlier-ratio (black curve), the estimated inlier-ratio (gray curve), and the inlier ratio prior (dashed) computed from correspondence qualities (see Sec. 2.1.2) as a function of the sorted subsets. The estimate curve tends to be close to the ground truth even when the inlier-ratio priors are not accurate. This experiment confirms the efficacy of the filter to estimate the inlier-ratio of the subsets in ANSAC. To test the

<sup>1</sup>Source code available at <http://vfragoso.com/>



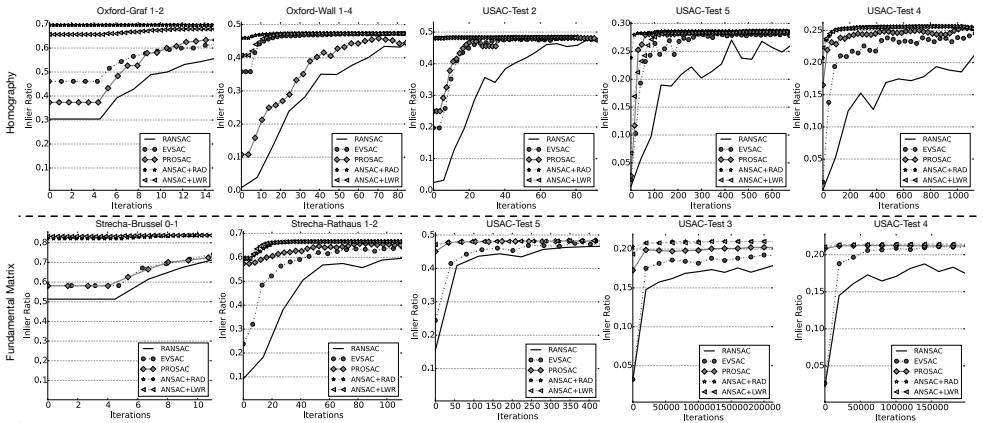


Figure 3: Convergence analysis for homography (top row) and fundamental matrix (bottom row) estimation. ANSAC tends to converge in the earliest iterations, while the competing methods (EVSAC, PROSAC, and RANSAC) require more iterations.

robustness of the filter to bad inlier-ratio priors, we replaced the inlier-ratio priors from the correctness quality values with synthetic estimates that deviate significantly from the ground truth. The results of this experiment are shown in Fig. 2 (b). The filter is able to “track” the ground truth for both models, even though it used priors that deviated significantly from the ground truth. See supplemental material for larger figures.

### 3.2 Convergence Analysis

The goal of this experiment is to analyze the benefits that the adaptive non-minimal sampler brings to the convergence of ANSAC. We did not use an early termination criterion for this experiment. The experiment was run with the following parameters:  $\alpha = 0.5$  and  $\sigma_p = \sigma_u = 0.1$  for the filter,  $\omega = 20$  and  $\mu = \frac{3}{4}$  which are the parameters of  $\lambda$  (see Eq. (7)), and  $q_{\max} = 4 \times q_{\min}$  (see Eq. (6)). The experiment considered the following competing estimators: RANSAC [14] (the baseline), EVSAC [15], and PROSAC [8]. We used the TheiaSfM implementations of the aforementioned estimators. The experiment used ANSAC with two different correctness correspondence quality measures: 1) Lowe’s radius (ANSAC+RAD), introduced above; and 2) the widely used Lowe’s ratio (ANSAC+LWR). Also, the experiment used correspondences ranked using Lowe’s ratio for PROSAC, and calculated the correctness probabilities from descriptor distances for EVSAC. A trial in this experiment measured the estimated overall inlier ratio of every estimator as a function of the iterations in their hypothesis-and-test loop. The experiment performed 1000 trials, and fitted a curve to all its trials for every estimator in order to get a smooth summary.

The top and bottom row in Fig. 3 show the results for homography and fundamental matrix estimation, respectively. The plots are sorted from left-to-right according to their ground truth overall inlier ratio. The plots show that ANSAC tends to converge faster than the competing methods. ANSAC converges faster than the competing methods when the inlier ratio is high ( $> 0.5$ ). However, when the inlier ratio is  $< 0.5$ , ANSAC tends to converge faster or comparable than PROSAC. This is because when the overall inlier ratio of an image pair is high, ANSAC uses non-minimal samples which yields a faster convergence than that of the competing methods. On the other hand, when the overall inlier ratio is low, ANSAC tends to use minimal samples more often, which yields a comparable performance than that of PROSAC. See supplemental material for more results about convergence.



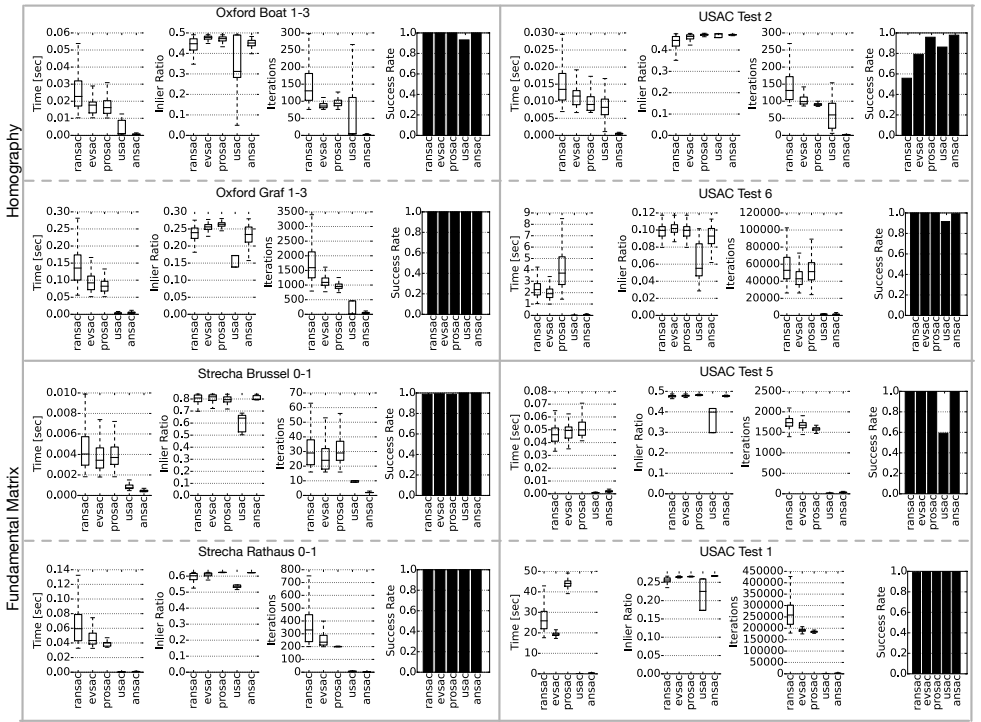


Figure 4: Box-plots measuring (from left to right) wall clock time, overall inlier ratio, iterations, and success rate of methods estimating homography (top two rows) and fundamental matrix models (bottom two rows). The plots show that ANSAC and USAC are the fastest; in most cases less than 20 msec. Unlike USAC, ANSAC achieves a similar inlier ratio and success rate that that of RANSAC, PROSAC, and EVSAC for both model estimations.

### 3.3 Estimation of Homography and Fundamental Matrix

The goal of this last experiment is to measure the speed (wall clock time) of an estimator to converge, the number of iterations, and the success rate and the overall inlier ratio of the estimated model. The experiment consisted of 1000 trials of each estimator. We measured success rate as the ratio between the number of estimates that are close enough to the ground truth and the total number of trials. This experiment considered RANSAC, PROSAC, EVSAC, and USAC [15] as part of the competing methods. We used USAC’s publicly available C++ source, and enabled PROSAC+SPRT [16]+LO-RANSAC [17]. We used ANSAC with the Lowe’s radius and enabled the termination criterion using  $\tau = 0.01$ .

The top and bottom two rows in Fig. 4 present the results for homography and fundamental matrix estimation, respectively. The plots show that ANSAC and USAC were the fastest estimators. ANSAC consistently required less than 20 msec, while USAC presented more time variation in homography estimation (see first row in Fig. 4). ANSAC computed estimates whose inlier ratio and success rates are comparable to those of RANSAC, EVSAC, and PROSAC. On the other hand, estimates computed with USAC presented a larger inlier ratio variation. These experiments confirm that ANSAC can accelerate the robust estimation of models using an adaptive non-minimal sampler, while still achieving a comparable or better performance than the state-of-the-art estimators. See the supplemental material for larger plots and extended results on all the experiments.

## 4 Conclusion

We have presented ANSAC, an adaptive, non-minimal sample and consensus estimator. Unlike existing estimators, ANSAC adaptively determines the size of a non-minimal sample to generate a hypothesis based on the inlier ratio estimated with a Kalman filter. In contrast to LO-RANSAC methods, which use non-minimal samples in a refinement process, ANSAC uses them in the hypothesis generation phase, avoiding the computational cost of the additional refinement phase. The homography and fundamental matrix estimation experiments demonstrate that ANSAC can converge in the early iterations and perform consistently better or comparable than the state-of-the-art.

**Acknowledgments.** This work was supported in part by NSF grants IIS-1657179, IIS-1321168, IIS-1619376, and IIS-1423676.

## References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building Rome in a day. *Comm. of the ACM*, 54(10):105–112, 2011.
- [2] Relja Arandjelović and Andrew Zisserman. Three things everyone should know to improve object retrieval. In *Proc. of the IEEE CVPR*, 2012.
- [3] Aveek Shankar Brahmachari and Santonu Sarkar. Hop-diffusion Monte Carlo for epipolar geometry estimation between very wide-baseline images. *IEEE TPAMI*, 35(3):755–762, 2013.
- [4] Matthew Brown, Richard Szeliski, and Simon Winder. Multi-image matching using multi-scale oriented patches. In *Proc. of the IEEE CVPR*, 2005.
- [5] David P Capel. An Effective Bail-out Test for RANSAC Consensus Scoring. In *Proc. of the BMVC*, 2005.
- [6] Jian Cheng, Cong Leng, Jiayang Wu, Hainan Cui, and Hanqing Lu. Fast and accurate image matching with cascade hashing for 3D reconstruction. In *Proc. of the IEEE CVPR*, 2014.
- [7] Ondřej Chum and Jiří Matas. Randomized RANSAC with Td, d test. In *Proc. of the BMVC*, 2002.
- [8] Ondřej Chum and Jiří Matas. Matching with PROSAC-progressive sample consensus. In *Proc. of the IEEE CVPR*, 2005.
- [9] Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized RANSAC. In *Pattern Recognition*, pages 236–243. Springer, 2003.
- [10] David J Crandall, Andrew Owens, Noah Snavely, and Daniel P Huttenlocher. SfM with MRFs: Discrete-continuous optimization for large-scale structure from motion. *IEEE TPAMI*, 35(12):2841–2853, 2013.
- [11] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [12] Dieter Fox, Jeffrey Hightower, Lin Liao, Dirk Schulz, and Gaetano Borriello. Bayesian filtering for location estimation. *IEEE pervasive computing*, (3):24–33, 2003.
- [13] Victor Fragoso and Matthew Turk. SWIGS: A Swift Guided Sampling Method. In *Proc. of the IEEE CVPR*, June 2013.

- [14] Victor Fragoso, Pradeep Sen, Sergio Rodriguez, and Matthew Turk. EVSAC: Accelerating Hypotheses Generation by Modeling Matching Scores with Extreme Value Theory. In *Proc. of the IEEE ICCV*, 2013.
- [15] Jan-Michael Frahm, Pierre Fite-Georgel, David Gallup, Tim Johnson, Rahul Raguram, Changchang Wu, Yi-Hung Jen, Enrique Dunn, Brian Clipp, Svetlana Lazebnik, et al. Building Rome on a cloudless day. In *Proc. of the ECCV*. Springer, 2010.
- [16] Liran Goshen and Ilan Shimshoni. Balanced exploration and exploitation model search for efficient epipolar geometry estimation. *IEEE TPAMI*, 30(7):1230–1242, 2008.
- [17] Karel Lebeda, Jiri Matas, and Ondrej Chum. Fixing the locally optimized RANSAC. In *Proc. of the BMVC*, 2012.
- [18] Yunpeng Li, Noah Snavely, Dan Huttenlocher, and Pascal Fua. Worldwide pose estimation using 3D point clouds. In *Proc. of the ECCV*. Springer, 2012.
- [19] Yin Lou, Noah Snavely, and Johannes Gehrke. Matchminer: Efficient spanning structure mining in large image collections. In *Proc. of the ECCV*. Springer, 2012.
- [20] David G Lowe. Distinctive image features from scale-invariant keypoints. *Intl. Journal of Computer Vision (IJCV)*, 60(2):91–110, 2004.
- [21] Jiří Matas and Ondřej Chum. Randomized RANSAC with sequential probability ratio test. In *Proc. of the IEEE ICCV*, 2005.
- [22] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, Jiří Matas, Fredrik Schaffalitzky, Timor Kadir, and Luc Van Gool. A comparison of affine region detectors. *International Journal of Computer Vision*, 65(1-2):43–72, 2005.
- [23] Rahul Raguram and Jan-Michael Frahm. RECON: Scale-adaptive robust estimation via residual consensus. In *Proc. of the IEEE ICCV*, 2011.
- [24] Rahul Raguram, Jan-Michael Frahm, and Marc Pollefeys. A comparative analysis of ransac techniques leading to adaptive real-time random sample consensus. In *Proc. of the ECCV*. 2008.
- [25] Rahul Raguram, Ondrej Chum, Marc Pollefeys, Jose Matas, and Jens Frahm. Usac: a universal framework for random sample consensus. *IEEE TPAMI*, 35(8):2022–2038, 2013.
- [26] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Fast image-based localization using direct 2D-to-3D matching. In *Proc. of the IEEE ICCV*, 2011.
- [27] Christoph Strecha, Rik Fransens, and Luc Van Gool. Wide-baseline stereo from multiple views: a probabilistic account. In *Proc. of the IEEE CVPR*, 2004.
- [28] Christoph Strecha, Rik Fransens, and Luc Van Gool. Combined depth and outlier estimation in multi-view stereo. In *Proc. of the IEEE CVPR*, 2006.
- [29] Christopher Sweeney, Tobias Hollerer, and Matthew Turk. Theia: A Fast and Scalable Structure-from-Motion Library. In *Proc. of the ACM Conference on Multimedia*, 2015.
- [30] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.
- [31] Ben J Tordoff and David W Murray. Guided-MLESAC: Faster image transform estimation by using matching priors. *IEEE TPAMI*, 27(10):1523–1535, 2005.
- [32] Wei Zhang and Jana Kosecka. Image based localization in urban environments. In *Proc. of the IEEE International Symposium on 3D Data Processing, Visualization, and Transmission*, 2006.