Exploiting Uncertainty in Random Sample Consensus

Rahul Raguram[1], Jan-Michael Frahm[1] and Marc Pollefeys[1,2]

[1]Department of Computer Science
The University of North Carolina at Chapel Hill
{rraguram,jmf}@cs.unc.edu

[2]Department of Computer Science
ETH Zürich
marc.pollefeys@inf.ethz.ch

Abstract

In this work, we present a technique for robust estimation, which by explicitly incorporating the inherent uncertainty of the estimation procedure, results in a more efficient robust estimation algorithm. In addition, we build on recent work in randomized model verification, and use this to characterize the 'non-randomness' of a solution. The combination of these two strategies results in a robust estimation procedure that provides a significant speed-up over existing RANSAC techniques, while requiring no prior information to guide the sampling process. In particular, our algorithm requires, on average, 3-10 times fewer samples than standard RANSAC, which is in close agreement with theoretical predictions. The efficiency of the algorithm is demonstrated on a selection of geometric estimation problems.

1. Introduction

The Random Sample Consensus (RANSAC) algorithm [8] is a widely used robust estimation technique, finding application in a variety of computer vision problems. The algorithm is simple and works well in practice, providing robustness even for substantial levels of data contamination. The basic RANSAC algorithm operates in a hypothesize-and-verify framework, where a minimal subset of the input data points is randomly sampled and used to hypothesize model parameters. The verification step then involves evaluating this model against all data points, and determining its support (the number of data points consistent with the model). This hypothesize-and-verify loop is terminated when the probability of finding a model with larger support than the current best model falls below some user-specified threshold (typically 1%-5%).

In its standard formulation, RANSAC operates with minimal assumptions on the data, relying purely on the fact that with enough iterations, a 'good' sample (the generally accepted definition being a minimal subset that is free from

outliers) will eventually be drawn, resulting in a model that has largest support. Since it is computationally infeasible to try every possible minimal subset, the standard termination criterion in RANSAC determines the number of trials required to ensure with some confidence η_0, that at least one minimal subset is drawn which contains all inliers. It can be shown [2] that for a 95% confidence in the solution, approximately three uncontaminated samples are drawn on average before the confidence in the solution is achieved.

The implicit assumption in the above formulation is that a model produced from an all-inlier minimal subset will be consistent with all other inliers in the data and, similarly, that a model produced from a contaminated sample will be consistent only with a few points. In practice, both these assumptions may be violated, leading to either increased runtimes, or incorrect solutions. In particular, it has been shown [5] that in practice, the standard stopping criterion is overly optimistic: due to noise in the data, a model generated from an all inlier sample may not be consistent with all other inliers. As a consequence of this decreased support, the number of samples drawn in RANSAC typically increases by a factor of two to three compared to the theoretically expected number. In addition, when the minimal sample forms a degenerate configuration for the entity being estimated (for instance, points on a plane in the case of fundamental matrix estimation), even a contaminated sample may result in a solution with large support. In this case, an incorrect solution will be returned by the algorithm.

In recent years, techniques have been proposed to deal with the cases described above. For instance, the Locally Optimized RANSAC (Lo-RANSAC) technique [5] aims to address the case where an all-inlier sample is supported only by a subset of all the inliers. Given a hypothesis with the largest support so far, Lo-RANSAC performs an 'inner RANSAC' loop, where a fixed number of models are generated by sampling non-minimal subsets from within the support of the current best solution, and verifying these models against all data points. The RANSAC algorithm then resumes sampling on all data points, carrying out the local optimization step every time a hypothesis with better sup-

port is found. The effect of the local optimization step is to mitigate the effects of noisy data points by using non-minimal subsets to generate hypotheses. Since this often results in an increase in the support, the net effect is that the overall number of iterations required by the algorithm reduces by a factor of two to three, thus moving closer to the theoretically predicted number. In addition, techniques have been proposed to deal with the case of degenerate data configurations, either specifically for epipolar geometry estimation [6] or for linear estimation problems in general [9].

In this work, we introduce a new framework for robust estimation, which addresses some of the issues mentioned above, by *explicitly* characterizing the uncertainty of the models hypothesized in the estimation procedure, in terms of their covariance. This uncertainty arises from imprecise localization of the data points from which the model is computed, as well as from their spatial distribution, which influences the conditioning of the problem. We show that by explicitly accounting for these uncertainties, it is possible to account for the effects of noise in the data. In other words, we show that knowledge of the uncertainties can be leveraged to develop an inlier classification scheme that is dependent on the model uncertainty, as opposed to a fixed threshold. In terms of compensating for noise in the estimation process, our work is closely related to the Lo-RANSAC approach, which provides a solution to a very similar problem. However, while Lo-RANSAC attempts to deal with the *effect* of imperfect model estimation (i.e., reduced support), the formulation we propose addresses precisely the *cause* itself. This strategy has the advantage that, given a model (with associated uncertainty), we can immediately identify a set of points that could potentially support this solution, if the imperfect model estimation were taken into account. The 'true' inliers then form a subset of this set of points, and we are able to efficiently extract this subset.

Going further, the framework we propose builds on another observation: all-inlier minimal subsets are sampled multiple times in RANSAC before the algorithm terminates. As mentioned above, on average, three all-inlier samples are drawn by RANSAC before a 95% confidence in the solution is achieved. In addition, due to the effects of noise, an all-inlier sample may not gather sufficient support, which further increases the number of samples drawn. By leveraging knowledge of the uncertainties, we can correctly estimate the support for any given all-inlier model, since we are explicitly accounting for the noisy estimation process. This implies that once we have found the first all-inlier sample, we can robustly find its support, and terminate the algorithm at that point. Note that here, we make the assumption that there is only one model in the data, which is indeed the case for a large variety of estimation problems.

To summarize the main contributions of this work:

a. We show how the uncertainties of the estimation pro-

cess and the points may be integrated into a random sampling framework, allowing us to robustly estimate the support for a model generated from a noisy minimal subset. Both the model and point uncertainties are factored into the inlier identification step, which results in a set of *potential* inliers, from which the true inliers can be efficiently extracted.

b. Given that we can estimate this support, we demonstrate that it is possible to terminate the random sampling algorithm once the *first* good sample has been drawn. Since we are stopping early as well as compensating for imperfect estimation, this should theoretically result in a factor of approximately 6-9 fewer samples than required by standard RANSAC. We show results that demonstrate this speed-up.

c. The framework we develop is flexible and can be integrated with other techniques, such as those for optimized model verification [4]. We leverage this idea both to identify good samples, as well as achieve an additional computational benefit. In addition, our algorithm may also be combined with techniques such as Least Median of Squares [18] and algorithms that handle degeneracy [6, 9].

2. Uncertainty Analysis

In general, there exists a wide body of work in photogrammetry and computer vision that deals with error analysis and propagation with respect to image measurements. A survey of this literature can be found in [12, 15, 10]. In this work, we restrict our attention to some common geometric entities in order to specifically demonstrate how knowledge of uncertainties may be integrated into a random sampling framework. Before describing the proposed robust estimation algorithm, we first provide a discussion of uncertainty analysis as applied to the estimation of two common geometric relations, the homography and fundamental matrix, and show how point uncertainties are transformed under the covariance of an estimated relation. The discussion in this section refers to [7] and [20]; the former develops the theory behind characterizing the uncertainty for homography estimation in terms of its covariance, while the latter deals with fundamental matrix estimation.

2.1. Uncertainty for homography estimation

Given a set of corresponding points $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$, where $\mathbf{x}_i$ and $\mathbf{x}'_i$ are in $\mathbb{P}^2$, the 2D homography $\mathbf{H}$ is the 3×3 projective transformation that takes each $\mathbf{x}_i$ to $\mathbf{x}'_i$. Since each point correspondence provides two independent equations in the elements of the matrix $\mathbf{H}$, a total of four corresponding points are required to compute $\mathbf{H}$ up to scale.

Expressing $\mathbf{x}_i$ and $\mathbf{x}_i'$ in terms of homogeneous 3-vectors, the pair of equations for correspondence i is given by:

$$\begin{bmatrix} w_i'\mathbf{x}_i^T & \mathbf{0}^T & -x_i'\mathbf{x}_i^T \\ \mathbf{0}^T & -w_i'\mathbf{x}_i^T & y_i'\mathbf{x}_i^T \end{bmatrix} \mathbf{h} = \mathbf{0}$$

where $\mathbf{x}_i' = (x_i', y_i', w_i')^T$ and $\mathbf{h}$ is the 9-vector made up of the entries of the matrix $\mathbf{H}$. By stacking the equations corresponding to four point correspondences, we then obtain a set of equations $\mathbf{A}\mathbf{h} = \mathbf{0}$, where $\mathbf{A}$ is an 8×9 matrix. The Direct Linear Transformation (DLT) algorithm for determining the solution then consists of computing the nullspace of $\mathbf{A}$ using the SVD. The unit singular vector corresponding to the smallest singular value is the solution $\mathbf{h}$ (up to scale).

In practice, the points used in the computation of the homography are imperfectly localized, and these errors in position cause uncertainty in the estimated transformation. This uncertainty is characterized by the covariance of the transform, which depends on the error in point locations, their spatial distribution and the method of estimation. As in [7], we assume that the DLT is used to compute $\mathbf{H}$. It is common to consider the computation points as having errors modeled by a homogeneous, isotropic Gaussian noise process with known variance. Note that the standard RANSAC algorithm also uses an estimate of this quantity to set the threshold for inlier classification. Note that while we do make the simplifying assumption of a common, isotropic covariance matrix for every point, as discussed in [13], this is a reasonable assumption to make in practice.

Assuming that the points in each image are perturbed by Gaussian noise with variance σ (we make the simplifying assumption that the variance is the same for points in both images; this does not affect the analysis), the uncertainty of the homography is encapsulated in a 9×9 covariance matrix Λ_h, which is given by

$$\Lambda_h = \mathbf{J}\mathbf{S}\mathbf{J} \tag{1}$$

where $\mathbf{J} = -\sum_{k=2}^{9} \mathbf{u}_k\mathbf{u}_k^T/\lambda_k$, with $\mathbf{u}_k$ representing the k^{th} eigenvector of the matrix $\mathbf{A}^T\mathbf{A}$ and λ_k the corresponding eigenvalue. $\mathbf{S}$ is a 9×9 matrix with a closed-form expression; we refer the reader to [7] for details on the derivation and the coefficients of $\mathbf{S}$.

In standard RANSAC, the computed homography is then typically used to determine the (symmetric) transfer error for each point correspondence in the data, given by

$$d_\perp = d(\mathbf{x}_i', \mathbf{H}\mathbf{x}_i)^2 + d(\mathbf{x}_i, \mathbf{H}^{-1}\mathbf{x}_i')^2 \tag{2}$$

where $\mathbf{H}$ and $\mathbf{H}^{-1}$ represent forward and backward transformations, respectively. If this error is below a predefined threshold t, the correspondence is classified as an inlier. Note that while the threshold setting, in theory, is related to the variance of the point measurement noise, the contribution of this noise to the uncertainty of the transform itself is not incorporated into standard RANSAC techniques.
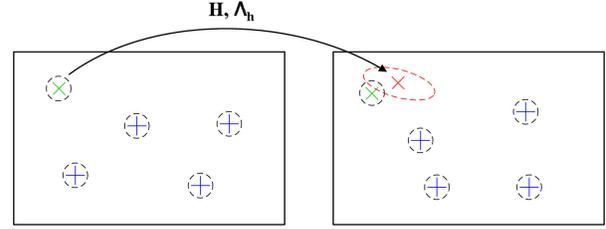
Figure 1. Using the covariance matrix of the estimated transformation in point transfer.

In order to determine the effect that the uncertainty of the transform has on point transfer, we consider the forward transformation $\hat{\mathbf{x}}_i' = \mathbf{H}\mathbf{x}_i$ (the formulation is symmetric in the reverse direction). Given the (assumed) 3×3 covariance matrix Λ_x for point $\mathbf{x}$, the covariance of the transferred point $\hat{\mathbf{x}}'$ may no longer be an isotropic Gaussian distribution. The uncertainty in the transferred point is then given by [7]

$$\Lambda_{\hat{\mathbf{x}}'} = \mathbf{B}\Lambda_h\mathbf{B} + \mathbf{H}\Lambda_x\mathbf{H} \tag{3}$$

where B is the 3×9 matrix

$$\begin{pmatrix} \mathbf{x}^T & \mathbf{0}^T & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{x}^T & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{0}^T & \mathbf{x}^T \end{pmatrix}$$

The first term in equation (3) represents the uncertainty in $\hat{\mathbf{x}}'$ given an exact point $\mathbf{x}$ and uncertain $\mathbf{H}$. The second term represents the uncertainty given an exact $\mathbf{H}$ and uncertain point measurement $\mathbf{x}$. The 3×3 matrix $\Lambda_{\hat{\mathbf{x}}'}$ thus defines a covariance ellipse that represents the uncertainty in point transfer due to both uncertainty in point measurement as well as uncertainty in the homography estimated from a minimal sample of four noisy correspondences. This result is illustrated in Figure 1, where a homography $\mathbf{H}$ is computed from a set of four point correspondences (denoted by the blue '+' symbols), along with its associated covariance matrix Λ_h. This covariance matrix is then used to transfer the uncertainty of another point (denoted by the green 'x') from the left image into the right. It can be observed that the uncertainty of the transferred point is now represented in terms of the red uncertainty ellipse, which represents a level set of the probability density function describing the noise spread around the ellipse centre. Since the two error regions (of the measured point in the right image and the transformed point) intersect in the second image, this implies that by accounting for the various uncertainties, this correspondence can be explained as a true match.

2.2. Uncertainty for fundamental matrix estimation

Given a set of corresponding points $\mathbf{x}_i \leftrightarrow \mathbf{x}_i'$, the fundamental matrix $\mathbf{F}$ is the 3×3 matrix that satisfies $\mathbf{x}_i'^T \mathbf{F}\mathbf{x}_i = 0$. Each correspondence provides one linear equation in the entries of $\mathbf{F}$. Since $\mathbf{F}$ is defined up to a scale factor, a simple

linear solution may be found from 8 point matches. In fact, by enforcing the rank constraint, $\mathbf{F}$ may be computed from 7 point matches. Due to the relative simplicity of the linear 8-point algorithm in terms of the discussion of its uncertainty analysis, we focus on this algorithm in the presented work. Thus, given 8 point correspondences, an appropriately normalized [11] matrix $\mathbf{M}$ may be constructed, where

$$\mathbf{M} = \begin{pmatrix} x_1'x_1 & x_1'y_1 & x_1' & y_1'x_1 & y_1'y_1 & y_1' & x_1 & y_1 & 1 \\ x_2'x_2 & x_2'y_2 & x_2' & y_2'x_2 & y_2'y_2 & y_2' & x_2 & y_2 & 1 \\ . & . & . & . & . & . & . & . & . \\ . & . & . & . & . & . & . & . & . \\ x_n'x_n & x_n'y_n & x_n' & y_n'x_n & y_n'y_n & y_n' & x_n & y_n & 1 \end{pmatrix}$$

Solving $\mathbf{Mf} = \mathbf{0}$, where $\mathbf{f}$ is a 9-vector made up of the entries of $\mathbf{F}$, followed by rank 2 constraint enforcement then provides the solution.

The uncertainty analysis for the 8-point algorithm was developed recently in [20]. Given the measurement noise in the pixel coordinates, the goal is to characterize the uncertainty of $\mathbf{F}$ in terms of its covariance matrix $\Lambda_\mathbf{F}$. If we denote the solution to the linear system of equations as $\tilde{\mathbf{f}}$, the covariance of $\tilde{\mathbf{f}}$ can be shown to be:

$$\Sigma_{\tilde{\mathbf{f}}}(X) = \sigma^2 \mathbf{J}_X \mathbf{J}_X^T \tag{4}$$

where $\mathbf{J}_X$ is an 8×32 matrix that represents the Jacobian of the transformation that generates $\tilde{\mathbf{f}}$ from the vector of point matches X, and σ is the standard deviation of the noise in point measurements. The rank 2 constraint is then enforced by computing

$$\mathbf{F} = \mathbf{UD} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \mathbf{V}^T$$

where $\mathbf{U}$, $\mathbf{D}$ and $\mathbf{V}$ are obtained from the singular value decomposition $\tilde{\mathbf{F}} = \mathbf{UDV}^T$. An efficient method for computing the Jacobian of the SVD is described in [16], and we use this result to compute the 9×9 Jacobian matrix $\mathbf{J}_{SVD}$. The final covariance matrix for $\mathbf{F}$ is then obtained as

$$\Sigma_F = \mathbf{J}_{SVD} \begin{pmatrix} \Sigma_{\tilde{\mathbf{f}}} & \mathbf{0}_{8,1} \\ \mathbf{0}_{1,8} & 0 \end{pmatrix} \mathbf{J}_{SVD}^T$$

Along the lines of the discussion in Section 2.1, the uncertainty of the transform may be used to map the covariance ellipse of a point in one image into the other. However, since the fundamental matrix defines a point-line mapping, the interpretation of this uncertainty is now altered. Specifically, given the fundamental matrix $\mathbf{F}$ and point $\mathbf{x}$, we have the epipolar line $\mathbf{l} = \mathbf{Fx}$ corresponding to $\mathbf{x}$. The uncertainty of this line is represented in terms of a covariance matrix

$$\Sigma_\mathbf{l} = \mathbf{J}_F \Sigma_F \mathbf{J}_F^T + \sigma^2 \mathbf{J}_\mathbf{x} \mathbf{J}_\mathbf{x} \tag{5}$$

where $\mathbf{J}_F$ and $\mathbf{J}_\mathbf{x}$ are the Jacobian of the point-line mapping with respect to $\mathbf{F}$ and $\mathbf{x}$. This covariance matrix defines a

conic C_k, which corresponds to a hyperbola in the second image. As in the case of the homography, the validity of a match can be determined in terms of whether the point $\mathbf{x}'$ lies within the mapped uncertainty area, or if their error regions intersect.

3. Incorporating uncertainty into RANSAC

In this section, we describe how knowledge of the uncertainty of the estimation may be leveraged in RANSAC to obtain a more efficient algorithm.

Inlier classification:

The discussion in Section 2 introduced the idea behind characterizing the uncertainty of the estimated transformations, and hinted at a way in which this could be used to perform the classification of points into inliers and outliers. For instance, consider the case of homography estimation - given the covariance $\Lambda_\mathbf{h}$ of the homography, we can compute the error ellipse for a transferred point. Now, instead of using the distance between the transferred point $\hat{\mathbf{x}}'$ and the measured point $\mathbf{x}'$ (as in equation (2)) as a metric, we can use this ellipse to identify potential inliers. In particular, if the point $\mathbf{x}'$ falls within this area, or if the covariance ellipse of the point $\mathbf{x}'$ intersects the mapped ellipse, then this could be a potential inlying match, since by accounting for the various uncertainties, these points indeed do conform to each other. If this process is repeated for all corresponding points, a set of *potential* inlying matches can be obtained. A similar interpretation holds true for the case of the fundamental matrix, where the covariance $\Lambda_\mathbf{F}$ maps the uncertainty of a point in one image to a hyperbola in the other.

Note that this does not imply that there is a consistent transformation that gives rise to this set of inliers; rather, this set represents points that could potentially be classified as inliers by accounting for uncertainties and considering each correspondence independently. A more subtle point is that since the covariance of the transform is estimated from the minimal sample and used to transfer additional points (that were not used in its computation), outliers may also fall into the set of potential inliers. This is because, as discussed in [12] (Section 5.2.6), extrapolation far beyond the set of points used to compute the transform may be unreliable. However, we note that if the assumptions of the point measurement uncertainties is accurate, then *all* inliers should fall into this set of potential inliers, since we are explicitly accounting for noise in the estimation. We thus make the observation that this set of potential inliers should contain all true inliers, and is likely to have a significantly higher inlier ratio than the original set of correspondences. This observation is important as it helps the algorithm terminate much quicker than standard RANSAC, particularly when the contamination level of the data is high.

Incorporating the covariance test into RANSAC:

Given that we can compute a set of potential inliers for a particular hypothesized transform, we then address the issue of when this covariance check should be applied. Clearly, it does not make sense to apply the uncertainty analysis to a sample that is contaminated, since the formulation described in Section 2 makes the assumption that the only source of error is noise in the point locations (and does not account for outlying matches). Thus, if we can identify minimal samples that are likely to be uncontaminated, then the covariance test can be applied to narrow down all potential inliers in the data. Identifying these uncontaminated samples is not trivial, since merely observing their support is not indicative - for instance, when the inlier ratio of the data is low, even a true sample may have a low absolute number of inliers. One way to address this problem is to estimate how likely it is that a configuration of points that form the support for a particular solution, could have occurred by chance. By choosing a model that has sufficiently low probability of being random, it will be possible to identify good samples. This is the idea behind the Minimum Probability of Randomness (MINPRAN) algorithm [19] which estimates this probability by assuming that outlier residuals follow a uniform distribution. However, as noted in [21], applying this technique to the estimation of geometric entities such as the fundamental matrix is not straightforward.

More recently, work in optimized model verification [14, 1, 4] has focused on minimizing the number of data points tested in the verification step of RANSAC. These techniques make use of the observation that most of the models hypothesized in RANSAC are likely to be contaminated, and are consistent with only a small fraction of the data points. Thus, it is often possible to discard bad hypotheses early on in the verification process. In particular, an optimal randomized model verification strategy is developed in [4], based on Wald's theory of sequential decision making. The evaluation step is cast as an optimization problem which aims to decide whether a model is good (H_g) or bad (H_b), while simultaneously minimizing the number of verifications performed per model, as well as the probability that a good model is incorrectly rejected early on.

Wald's Sequential Probability Ratio Test (SPRT) is based on the likelihood ratio

$$\lambda_j = \prod_{r=1}^{j} \frac{p(x_r|H_b)}{p(x_r|H_g)} \qquad (6)$$

where x_r is equal to 1 if the r^{th} data point is consistent with a given model, and 0 otherwise. Thus, $p(1|H_g)$ denotes the probability that a randomly chosen data point is consistent with a good model, and this can be approximated by the inlier ratio ε. Similarly, $p(1|H_b)$ is the probability that a randomly chosen data point is consistent with a bad

model, approximated by a parameter δ, which is initialized based on geometric considerations and may be updated as the algorithm progresses, in terms of the average fraction of data points consistent with bad models. If, after evaluating j data points, the likelihood ratio becomes greater than some threshold A, the model is rejected. As shown in [4], the probability α of rejecting a good model in the SPRT is $\alpha \leq 1/A$. Thus, the larger the value of A, the smaller the probability of rejecting a good model. However, the number of points verified per model increases as $log(A)$. An optimal value of A can be found to satisfy these two constraints.

While the verification procedure described above was developed with a view to minimizing the number of evaluations for bad models, we make the observation that as a byproduct, we also obtain an idea of the models that are likely to be uncontaminated. Since contaminated models will be rejected early on in the SPRT, models that survive this evaluation procedure are likely to be those that are good. Adopting this strategy to identify good models has the benefit that an additional computational speedup comparable to [4] is achieved by the early termination of contaminated hypotheses.

3.1. Algorithm

With the above discussion in mind, we now describe a modified RANSAC algorithm that incorporates knowledge of the uncertainties into the estimation process. The technique is described in Algorithm 1.

3.2. Algorithm description

Algorithm 1 describes the Cov-RANSAC technique which accounts for uncertainty in the estimation process. The hypothesis generation phase proceeds as in standard RANSAC. Note that we do not assume that prior information about the validity of correspondences is available. If available, this information can be easily integrated into the sampling stage to achieve non-uniform sampling [3]. Hypothesis verification is then performed using the SPRT. There are two possible outcomes at this stage: either a model is rejected early on, or the model survives this evaluation stage and is accepted. In the first case, we move on to the next hypothesis. In the second case, we perform the covariance estimation procedure to identify the set of potential inliers I'. With a view to increasing the performance of the algorithm, note that we first perform a check on the covariance of the transform, to estimate how well constrained the solution is. The covariance of the transform depends on the spatial distribution of the points used in its computation; thus, if the points fall into an unstable configuration, the uncertainty of the transform could potentially be very large. While this in itself would not cause the algorithm to fail, it could cause the size of the set of potential inliers I' to be close to the size of the dataset, implying that the speed-up

Algorithm 1 RANSAC with uncertainty estimation: Cov-RANSAC

---

   1. Hypothesis generation:
   Sample a minimal subset of size s at random from the data
   Estimate model parameters from this minimal subset
   2. Verification:
   Evaluate the model using SPRT (refer [4] for the algorithm)
   if model accepted **then**
     3. Covariance test
     Compute the covariance Λ_T of the transform (Sec. 2)
     if trace of covariance matrix $\leq$ threshold **then**
       for $i = 1$ to number of data points **do**
         Compute the mapped error region $\Lambda_{\hat{x}'}$ or $\Lambda_{\hat{i}}$
         if point satisfies the covariance test **then**
           Add point i to set of potential inliers I'
         end if
       end for
       4. Inner RANSAC:
       Run standard RANSAC on the set I'
       Return solution and set of true inliers I
     else
       Go to hypothesis generation step
     end if
   else
     Go to hypothesis generation step
   end if

---

achieved would be minor. While it is possible to envision a set of nested calls to the algorithm, each time with a smaller set I', we use the approach in Algorithm 1 for simplicity.

Once a set of potential inliers is found, an inner RANSAC is performed (Step 4). Since the inlier ratio is now typically much higher than before, this RANSAC requires far fewer iterations to terminate. The solution found is then returned, along with the final set of inliers, I. Note that we could perform an SPRT-based RANSAC in the inner loop as well, to achieve additional computational savings. We also make the point that since we have a lower bound on the inlier ratio in the set I' (given by the ratio of the number of inliers found after the SPRT to the size of I'), we could potentially perform a least quantile of squares (similar to [18]), which would allow us to find the solution without relying on a threshold. The tradeoff in this case would be an increased number of samples in the inner loop. Finally, note that while, in general, performing rigorous variance propagation usually leads to additional costs, our algorithm typically invokes this step only once (when the first uncontaminated sample is found). Thus, variance propagation constitutes a constant additional cost which, particularly for low inlier ratio problems, is a small fraction of the total runtime.

4. Results

4.1. Covariance propagation vs. local optimization

As discussed earlier, the technique we propose in this work is related to Lo-RANSAC in that both approaches aim to deal with the fact that an uncontaminated sample may not be consistent with all inliers. Thus, one may wonder how our algorithm functions if the covariance test were simply replaced with a local optimization step. In particular, one may replace the covariance estimation and propagation steps in Algorithm 1 (i.e., Step 3) by a fixed number of RANSAC iterations, where non-minimal subsets are sampled from the set of inliers to the first model that passes the SPRT, and are then evaluated against all points. As in [5], we use 10 iterations of this inner RANSAC, with the size of the non-minimal subset being $min(size(I_k), 12)$, where I_k is the support for the hypothesis that passes the SPRT. The rest of Algorithm 1 remains the same.

The main advantage to using the covariance estimation versus local optimization is that sampling in the latter is purely local in nature and may get caught in local minima. For instance, consider the problem of estimating the homography between the images shown in Figure 2. Figure 3(a) shows the ground truth inliers superimposed on one of the images - these inliers were found by verifying $100,000$ random models and represents a close estimate of the true inlier set. For the image pair shown in Figure 2, there are 3169 putative correspondences and 1397 inliers, leading to an inlier ratio of 0.44. Running Algorithm 1 with local optimization instead of the covariance check leads to the algorithm often terminating with an inlier set that represents a local maximum; one such example is shown in Figure 3(b). On the other hand, the unmodified Algorithm 1 is consistently able to find all inliers in the data, since it compensates specifically for the noisy estimation process. This point is further illustrated in Figure 4, which compares the number of inliers returned after local optimization and after the covariance test, over 500 runs. It can be seen that in the case of local optimization a lower number of inliers is frequently returned. On the other hand, the covariance test returns, on average, a larger number on inliers. The mean and standard deviation for the number of inliers returned by the local optimization method are 975.6 and 302.2, while the corresponding numbers for the covariance test are 1274.1 and 98.9, respectively. Note that these numbers represent the raw inlier counts returned by the algorithm, without an optimization step performed *post-hoc*.

4.2. Homography estimation

The performance of the proposed approach was evaluated against the standard RANSAC algorithm, the original Lo-RANSAC approach [5] and the SPRT-based approach [4], which represents the state of the art. Note that since

Figure 2. Image pair used for homography computation.
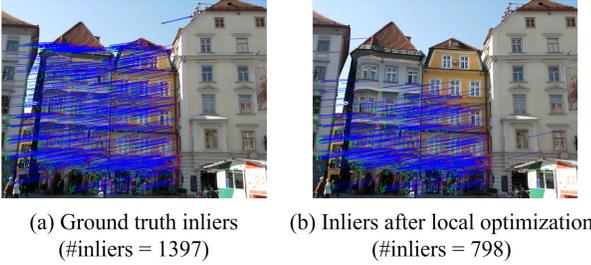
(a) Ground truth inliers
(#inliers = 1397)

(b) Inliers after local optimization
(#inliers = 798)

Figure 3. (a) Ground truth inliers and (b) inliers corresponding to a sample solution obtained using local optimization in Algorithm 1. It can be seen that the local optimization step fails to recover all true inliers.
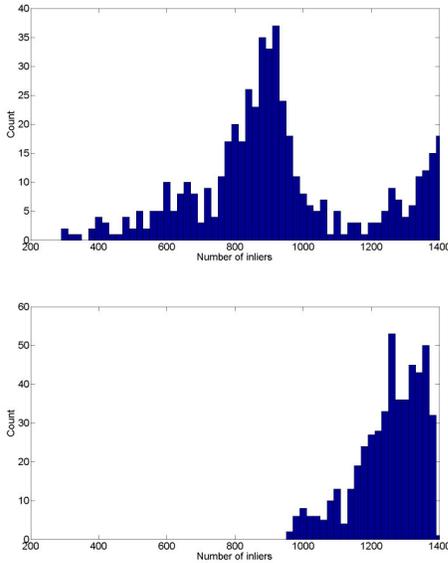
Figure 4. Histograms of inlier counts returned over 500 runs using (a) local optimization and (b) the covariance test, in Algorithm 1. Note that the average number of inliers returned is substantially lower in the former case. In addition, roughly 20% of the runs using local optimization return less than half the total number of true inliers, as compared to no such runs with the covariance test.

we do not assume the existence of prior information regarding correctness of the correspondences, we do not compare against techniques that specifically exploit this information [3, 17]. The image pairs used in the experiments are shown in Figure 5, representing a range of inlier ratios. For these experiments, the assumed standard deviation for the mea-

Figure 5. Image pairs used in the homography experiments. Inlier ratios: $A - 0.33$, $B - 0.23$, $C - 0.44$, $D - 0.19$

surement error was 1 pixel. For the SPRT, the parameters were initialized conservatively with $\varepsilon = 0.1$ and $\delta = 0.01$. The threshold parameter for the trace of the covariance matrix was set to 500, and this was held constant over the images tested. The results of the evaluation are tabulated in Table 1, which shows average numbers over 500 runs for inliers found (I), samples evaluated (k) and verifications per model (vpm) for the four RANSAC techniques. It can be observed that using the covariance information in RANSAC leads to a 3-10 fold reduction in the number of samples evaluated compared to standard RANSAC. In addition, since the algorithm leverages the SPRT, the number of verifications per model is also reduced by a factor of 4-18 compared to standard RANSAC. In contrast, the Lo-RANSAC technique reduces the number of samples by 2-3 fold, with a slight increase in the number of inliers found (this is due to the use of non-minimal samples). The SPRT-based RANSAC significantly reduces the number of verifications per model, but requires the evaluation of the same order of samples as in standard RANSAC. In summary, the performance improvement of the Cov-RANSAC technique in terms of samples evaluated is in line with theoretical expectations. Furthermore, an additional speed-up is obtained due to the optimized model verification incorporated in the algorithm.

4.3. Fundamental matrix estimation

A similar performance comparison was performed for the task of estimating the epipolar geometry for image pairs. The images used in the experiments are shown in Figure 6. The standard deviation of the measurement error was again assumed to be 1 pixel, and the SPRT parameters were initialized as $\varepsilon = 0.2$ and $\delta = 0.05$. The results of the evaluation are shown in Table 2, which tabulates average numbers over 500 runs, for all techniques. The results again demonstrate the advantage of using the uncertainty information. It can be seen that the proposed technique leads to a 4-10 fold reduction in the number of samples evaluated, with an additional computational benefit resulting from the early termination of bad hypotheses.

		RANSAC	SPRT	Lo-RANSAC	Cov-RANSAC
A	I	146	147	152	145
$\varepsilon \sim 0.33, N = 449$	k	689	972	278	102
	vpm	449	31	449	64
B	I	1278	1280	1367	1274
$\varepsilon \sim 0.44, N = 3169$	k	215	272	94	65
	vpm	3169	367	3169	688
C	I	296	299	310	298
$\varepsilon \sim 0.23, N = 1327$	k	2411	2893	1004	303
	vpm	1327	29	1327	127
D	I	428	437	441	432
$\varepsilon \sim 0.19, N = 2239$	k	3127	3561	1809	302
	vpm	2239	39	2239	120

Table 1. Comparison of the different RANSAC techniques for homography estimation on real data. The table shows, for each image pair and estimation algorithm, the average number of inliers (I), samples evaluated (k) and verifications per model (vpm). Note that Cov-RANSAC achieves a **3-10 fold reduction** in the number of samples evaluated.
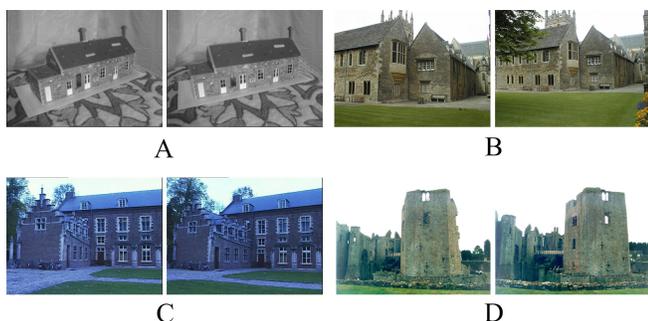
Figure 6. Image pairs used in the fundamental matrix experiments. Inlier ratios: $A - 0.65$, $B - 0.53$, $C - 0.5$, $D - 0.59$

		RANSAC	SPRT	Lo-RANSAC	Cov-RANSAC
A	I	292	294	296	294
$\varepsilon \sim 0.65, N = 449$	k	159	242	92	44
	vpm	449	39	449	61
B	I	1276	1280	1298	1278
$\varepsilon \sim 0.53, N = 2399$	k	762	819	539	132
	vpm	2399	79	2399	191
C	I	661	661	670	660
$\varepsilon \sim 0.5, N = 1327$	k	1737	1791	788	312
	vpm	1327	47	1327	144
D	I	1248	1251	1261	1251
$\varepsilon \sim 0.59, N = 2109$	k	384	422	212	38
	vpm	2109	41	2109	133

Table 2. Comparison of the different RANSAC techniques for fundamental matrix estimation on real data. The table shows, for each image pair and estimation algorithm, the average number of inliers (I), samples evaluated (k) and verifications per model (vpm). Note that Cov-RANSAC achieves a **4-10 fold reduction** in the number of samples evaluated.

5. Conclusion

In this paper, we demonstrated how uncertainty information may be incorporated into a random sampling framework. The resulting algorithm, Cov-RANSAC, achieves a 3-10 fold reduction in the number of samples evaluated for two common geometric estimation problems. The algorithm also uses optimized model verification, which leads to an additional computational benefit. The framework is flexible and can be combined with other techniques, for instance, using prior information to guide the sampling process. The results demonstrate the advantages of using uncertainty information in performing robust estimation.

References

[1] D. Capel. An effective bail-out test for RANSAC consensus scoring. In *Proc. BMVC*, 2005.

[2] O. Chum. Two-view geometry estimation by random sample and consensus. In *Ph.D. Dissertation, CTU*, 2005.

[3] O. Chum and J. Matas. Matching with PROSAC - progressive sample consensus. In *Proc. CVPR*, 2005.

[4] O. Chum and J. Matas. Optimal randomized RANSAC. *PAMI*, August 2008.

[5] O. Chum, J. Matas, and J. Kittler. Locally optimized RANSAC. In *Proc. DAGM*, 2003.

[6] O. Chum, T. Werner, and J. Matas. Two-view geometry estimation unaffected by a dominant plane. In *Proc. CVPR*, 2005.

[7] A. Criminisi, I. Reid, and A. Zisserman. A plane measuring device. In *Proc. BMVC*, 1999.

[8] M. Fisher and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM*, 24(6):381–395, 1981.

[9] J.-M. Frahm and M. Pollefeys. RANSAC for (quasi-)degenerate data (QDEGSAC). In *Proc. CVPR*, 2006.

[10] R. M. Haralick. Propagating covariance in computer vision. In *In Proc. Workshop on Performance Characteristics of Vision Algorithms*, 1994.

[11] R. I. Hartley. In defense of the eight-point algorithm. *PAMI*, 19(6):580–593, 1997.

[12] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. 2000.

[13] K. Kanatani. Uncertainty modeling and model selection for geometric inference. *PAMI*, 26(10):1307–1319, 2004.

[14] J. Matas and O. Chum. Randomized RANSAC with $T_{d,d}$ test. *Image and Vision Computing*, 22(10):837–842, 2004.

[15] C. McGlone, editor. *The Manual of Photogrammetry, 5th Ed.* The American Society of Photogrammetry, 2004.

[16] T. Papadopoulo and M. I. A. Lourakis. Estimating the jacobian of the singular value decomposition: Theory and applications. In *Proc. ECCV*, 2000.

[17] R. Raguram, J.-M. Frahm, and M. Pollefeys. A comparative analysis of RANSAC techniques leading to adaptive real-time random sample consensus. In *Proc. ECCV*, 2008.

[18] P. J. Rousseeuw. Least median of squares regression. *Journal of the American Statistical Association*, 79:871–890, 1984.

[19] C. V. Stewart. MINPRAN: A new robust estimator for computer vision. *PAMI*, 17(10):925–938, 1995.

[20] F. Sur, N. Noury, and M.-O. Berger. Computing the uncertainty of the 8 point algorithm for fundamental matrix estimation. In *Proc. BMVC*, 2008.

[21] P. Torr, A. Zisserman, and S. Maybank. Robust detection of degenerate configurations for the fundamental matrix. *Proc. ICCV*, 1995.