StaRSaC: Stable Random Sample Consensus for Parameter Estimation

Jongmoo Choi and Gérard Medioni
Institute of Robotics and Intelligent Systems
University of Southern California, Los Angeles, CA 90089
{jongmooc, medioni}@usc.edu

Abstract

We address the problem of parameter estimation in presence of both uncertainty and outlier noise. This is a common occurrence in computer vision: feature localization is performed with an inherent uncertainty which can be described as Gaussian, with unknown variance; feature matching in multiple images produces incorrect data points. RANSAC is the preferred method to reject outliers if the variance of the uncertainty noise is known, but fails otherwise, by producing either a tight fit to an incorrect solution, or by computing a solution which includes outliers. We thus propose a new estimator which enforces stability of the solution with respect to the uncertainty bound. We show that the variance of the estimated parameters (VoP) exhibits ranges of stability with respect to this bound. Within this range of stability, we can accurately segment the inliers, and estimate the parameters, the variance of the Gaussian noise. We show how to compute this stable range using RANSAC and a search. We validate our results by extensive tests and comparison with state of the art estimators on both synthetic and real data sets. These include line fitting, homography estimation, and fundamental matrix estimation. The proposed method outperforms all others.

1. Introduction

We address the problem of parameter estimation in presence of both uncertainty and outlier noise. This is a common occurrence in computer vision: feature detection introduces location uncertainty which can be described as Gaussian, with unknown variance; feature matching in multiple images produces incorrect data points [9]. RANSAC has been successfully applied in computer vision applications including range data segmentation and structure from motion, in presence of a number of outliers [5, 6]. The objective function to be maximized in RANSAC is the number of data points that lie within a distance threshold (i.e. inliers). This random selection is repeated a number of times and the hypothesis with the maximum value represents the fit. A critical assumption in the RANSAC algorithm is that the inlier bound, or the variance of the noise is known [5]. Image features (e.g. interesting points) are measured with inaccuracy caused by camera noise, the inaccuracy of feature extraction and matching, and moving objects. Unfortunately, degradation of performance is inevitable if we use a predefined threshold (noise model). For instance, outliers might be taken as inliers when a large bound is employed. As a result, the estimation result, even computed with a single outlier, can be biased significantly. Conversely, if a small bound is applied, the number of inliers may not be large enough to correctly estimate parameters [1].

Here, we propose a new estimator which enforces stability of the solution with respect to the uncertainty bound. We show that the variance of the estimated parameters (VoP) exhibits ranges of stability with respect to this bound. Within this range of stability, we can accurately, segment the inliers, and estimate the parameters, the variance of the Gaussian noise.

We show how to compute this stable range using RANSAC and a search in the bound space. The minimum variance of estimated parameters (VoP) is related to the optimal threshold of RANSAC. The performance of RANSAC depends on the threshold. For a correct threshold which captures the inlier structure, since the mode of Gaussian acts like an attractor in which estimation will converge to the maximum value of the Gaussian, the estimated results produced by random trials can be close to the ground truth and the VoP will be small. For a low threshold, the sample taken by RANSAC is too small to represent the inlier data, resulting a biased output. For a high threshold, outliers are included and the estimate may deviate far from the desired parameters. Both small and high threshold give a higher VoP than the correct threshold case. Based on this, we can define the stability of an estimator, which describes the VoP with respect to the changes of threshold.

The proposed algorithm can be summarized as follows: First, an initial range is defined between 0 and the value computed by a least-square fit to all available data. Second, we estimate a proper threshold which has the mini-

978-1-4244-3991-1/09/$25.00 ©2009 IEEE

mum VoP in the search region. This provides the estimate of parameters. We then choose the highest bound consistent with these parameters, which in turn allow us to estimate the variance of the Gaussian noise.

The reminder of the paper is organized as follows. Section 2 shows the related work. In Section 3, we present the StaRSaC algorithm. Section 4 provides the analysis of the method. This is followed by Section 5 where the proposed idea is validated experimentally.

2. Related Work

Robust estimation methods can be classified into four categories: random sampling algorithms, scale estimation, case deletion diagnostics, and voting methods

RANSAC (1981) [5] is a minimal subset random sampling search techniques. The objective function to be maximized is the number of data points that lie within a distance threshold. This random selection is repeated a number of times and the hypothesis with the maximum value represents the fit. LMedS (Least-Median-Squares, 1996) [17] calculates for each solution the median distance between the points and model. MLESAC (Maximum LikElihood SAmple Consensus, 2000) [14] chooses the solution that maximizes the likelihood rather than the number of inliers. MAPSAC (Maximum A Posteriori SAmple Consensus, 2002) [12] use Bayesian probabilities to improves MLESAC being more robust against outliers. M-Estimators [6] reduces the effect of outliers weighting the residual of individual point.

MINPRAN (1995) [3], MUSE (1996) [10], ASSC (2004) [15] explicitly estimate of the unknown scale of the inliers noise to extract the inliers from the data. MINPRAN assumes the outliers are randomly distributed within a certain range. MUSE evaluates a hypothesized fit over potential inlier sets via in objective function of unbiased scale estimates. MUSE is limited to handle extreme outliers. In ASSC, mean shift finds a local peak close to zero on the residual density curve, then mean shift valley algorithm finds the valley (local minimum) next to that peak [15]. It is usually assumed that a large number of data samples are available to estimate the scale. As a result, a scale estimator might fail to get an unbiased result in cases where only a moderate number of data samples are available [8].

Case deletion method has been proposed to remove outliers, which identifies outliers based on some measures such as Cook's distance (1977) in statistics [13] . The influence of each observation can be evaluated by the measuere. The results may depend on the initial estimation since the methods delete outliers incrementally.

In the Hough transform [2], the parameter space is divided into cells, and the parameters are selected based on the values of the cells which have accumulated numbers by each datum. One of important problem is the dimensionality with respect to the problem space. Tensor voting [11] can handle a very large percentage of outliers, but does not take Gaussian noise into account.

3. Stable Random Sample Consensus

3.1. Problem formulation

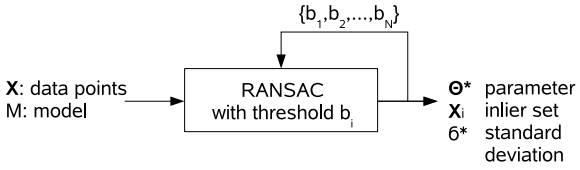

Figure 1. Problem formulation

Our method produces a good estimate θ^* of the parameter vector θ, and the RANSAC threshold, from which can be derived the inlier set, and an estimate σ^* of the Gaussian noise standard deviation σ as illustrated in Fig 1.

3.2. The algorithm

Our key observation is that, as we change the RANSAC threshold, the variance of the estimated parameter $\hat{\theta}$, (VoP) falls into 3 categories.

If the threshold is too small: RANSAC produces a tight fit to an unstable solution. This instability increases with both the variance of the uncertainty and with the number of outliers.

If the threshold is too large: Here again, the fit produces a set of possible solutions, and is affected by outliers.

Finally, we experimentally observe that there exists a wide range for which the fit is stable.

Once we have found a stable estimate $\theta^*$of the parameters, we can choose the largest RANSAC threshold consistent with θ^*, which in turn provide us with the set of inliers.

Finally, if needed, we can provide a good estimate of the variance of the Gaussian noise from the set of inliers. This is presented in Algorithm 1.

3.3. Illustration: 2-D line fit

We show a simple 2-D line estimation problem. First, we study the behavior of RANSAC with respect to the threshold. We define a 2-D line model and 160 inlier points with Gaussian noise $\mathcal{N}(0, 2)$. We add 40 outlier points into the data set randomly distributed. The total number of points is 200. Fig. 2 shows the behavior of RANSAC which provides a good result if the threshold is the same as the pre-defined standard deviation ($b = \sigma = 2$). Both a large threshold ($b = 200$) and a small threshold ($b = 0.0002$) yield significant biased results in this example.

Algorithm 1 StaRSaC: Stable Random Sample Consensus

Input: measurement $\mathbf{X}$, a model M
Output: estimate θ^*, inliers $\mathbf{X}^*_{in}$, bound b^*
Initialization of the range: $[b_L, b_U]$
  Lower bound: $b_L = \epsilon$
  Upper bound (b_U) using least-square method
Define the search range $\{b_i\} \in [b_L, b_U]$
For $b' = [b_1, b_2, \cdots, b_n]$
  Set $S = \emptyset$
  For j=1 to K
    Compute parameter $\hat{\theta}_j = f(b', \mathbf{X})$ using RANSAC
    $S = S \bigcup \hat{\theta}_j$
  End
  $var(\hat{\theta}|b_i) = E\{(\bar{\theta} - \hat{\theta}_j)^2\}, \theta_j \in S$
End
b_m, θ^* using $\arg\min_i var(\hat{\theta}|b_i)$
Choose b^* in $\{b|b > b_m\}$ consistent with θ^*
$\mathbf{X}^*_{in} = f(b^*, \mathbf{X})$
Compute σ^* using $\mathbf{X}^*_{in}$

Fig. 3 (a) shows the VoP and VoI with respect to the threshold. While the VoI decreases asymptotically when the threshold decreases, the VoP is bounded and stays in a stable range. Fig. 3 (b) shows the number of inliers with respect to the threshold. In the stable range, more inliers might provide better results since the estimation accuracy is related to the number of inliers. Fig. 3 (c) shows that the estimation results in the stable region produce the correct parameters. The error is computed as the inner product between the parameters of the ground truth 2-D line and the estimated parameters.

4. Analysis of the method

4.1. RANSAC and the threshold

In RANSAC estimation, the cost function is defined as

$$C = \sum_i \rho(e_i^2), \tag{1}$$

where

$$\rho(e^2) = \begin{cases} 0 & e^2 < b^2 \\ constant & e^2 \geq b^2. \end{cases} \tag{2}$$

If we know the standard deviation σ of the Gaussian distribution, the optimal threshold b^* can be computed by using the inverse cumulative chi-squared (χ^2_n) distribution [6]. In real applications, however, the prediction of noise may not be possible.

4.2. VoP and the RANSAC threshold

We show that the minimum variance of parameters (VoP) is related to the optimal threshold of RANSAC. The performance of RANSAC depends on this threshold. To study

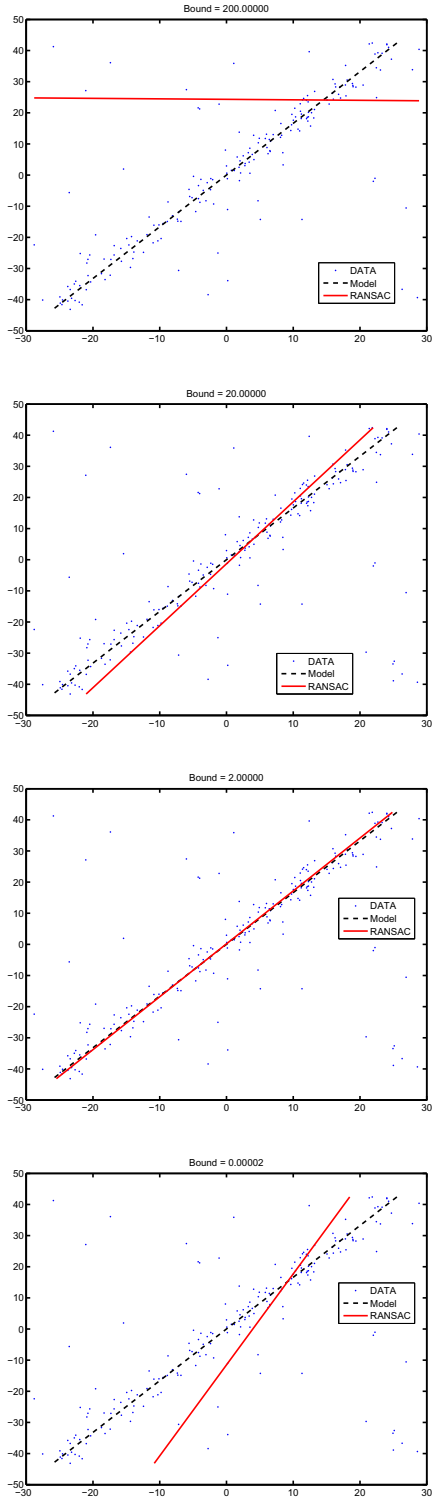

Figure 2. 2-D line estimation results with different thresholds (200, 20, 2, 0.0002. Threshold 0.2 provides a result similar to 20.

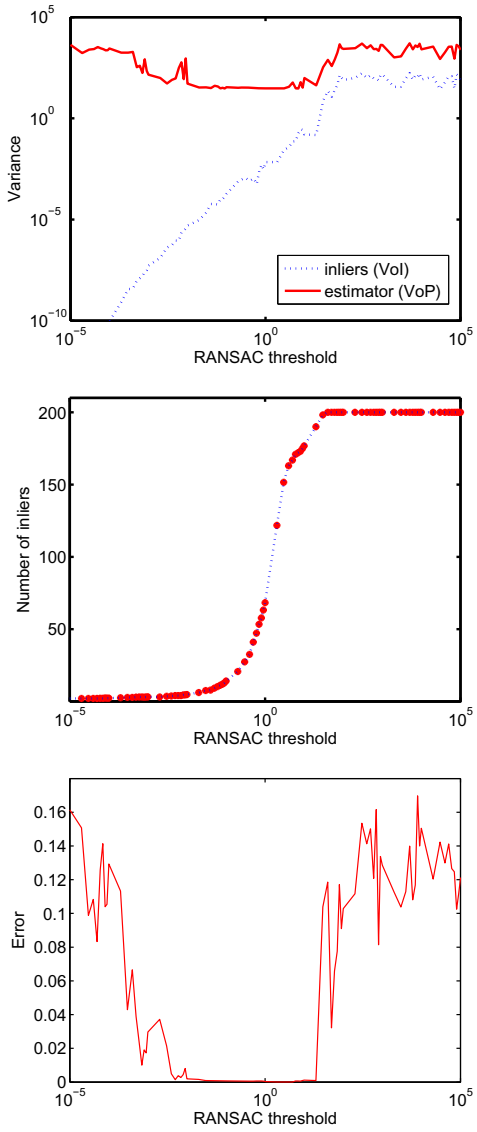

Figure 3. Results of 2-D line estimation. (a) Variance of estimator and variance of inliers. (b) Number of inliers with respect to the bound changes. (c) Error in 2-D line estimation.

the behavior of RANSAC threshold, we divide the bound space into three different regions: near optimal ($b \sim b^*$), low threshold ($b \ll b^*$), and high threshold ($b \gg b^*$) areas. Near the optimal value, the mode of Gaussian acts like an attractor in which estimation converges to the maximum value of the Gaussian. In the lower threshold area, the sample taken by small threshold in RANSAC is too small to represent the inlier data, and may produce biased results. In the high threshold area, outliers acts as inliers, the estimated results represent the whole data including outliers, which means the estimation deviates from the parameters.

4.2.1 Variance of an estimator

Parameter estimation finds the values of the parameters that lead to the best agreement between the predictions of the model (e.g. epipolar geometry) and the data (e.g. the set of corresponding points in images). Here, we do not consider the model selection problem, and use a single model (e.g. line fit, fundamental matrix $\mathbf{F}$, or homography $\mathbf{H}$).

Given a model that has a set of adjustable parameters in the vector $\theta = (\theta_1, \theta_2, \cdots, \theta_p)^T$ and measured points $\mathbf{X}$, the goal is to find θ^*. In the RANSAC algorithm, each component of the parameter vector can be defined as a function of the measurements $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_N]$, predetermined model M, and the threshold b as

$$\hat{\theta} = f(\mathbf{X}, M, b). \tag{3}$$

The error of estimated parameter is $|\theta^\dagger - \hat{\theta}|$ where $\theta^\dagger$ is the unknown ground truth. We can define the VoP as

$$var(\hat{\theta}) = E\{(\hat{\theta} - \bar{\theta})(\hat{\theta} - \bar{\theta})^T\} \tag{4}$$

where $var()$ is the covariance matrix of dimenstion $(p \times p)$. Given K estimated parameters, $\mathbf{\Theta} = [\theta_1, \theta_2, \ldots, \theta_K]$, we compute the covariance matrix $\mathbf{\Sigma} = \mathbf{\Theta}\mathbf{\Theta}^T \in R^{(P \times P)}$ and decompose it into $\mathbf{\Sigma} = \mathbf{L}\mathbf{D}\mathbf{V}^T$ using singular value decomposition. The variance of parameters (VoP) is computed as $var(\hat{\theta}) \doteq trace(\mathbf{D})$. Note that the first eigenvalue is dominant, and gives similar results in our example.

4.2.2 Case I: near optimal threshold

Suppose that we have an optimal threshold b^* and the error of our estimation is minimum with this threshold as

$$E\{\hat{\theta} - \theta^\dagger | b^*\} \leq E\{\hat{\theta} - \theta^\dagger | b\}, b \neq b^*. \tag{5}$$

First, if there is no outlier, the best estimate should represent the mode of the Gaussian. Since the Gaussian distribution is continuous, $b = b^* + \delta(\delta \to 0)$ should give the same result. Second, outliers which are far from the mode do not affect the result of estimation due to the threshold. Hence, in a certain range near the optimal threshold, we have a stable estimate $\hat{\theta}$, and $var(\hat{\theta})$ is small.

4.2.3 Case II: high threshold

Suppose that we know the optimal threshold and compute the parameters by RANSAC with the threshold. The VoP $var(\hat{\theta}|b^*)$ should be greater than zero even if we have only inliers with the optimal threshold.

The Cramér-Rao inequality shows that there is a limit on how much information about unknown parameters can be extracted from a set of measurements [16]. The score is defined by $V = \frac{\partial}{\partial \theta} \ln p_\theta(x)$. The variance of the score is

the Fisher information: $J(\theta) =< [\partial\theta \log p_\theta(x)]^2 >$. The Cramér-Rao bound (CRB) states that the mean square error of an unbiased estimator f of θ is lower bounded by the inverse of the Fisher information as

$$\sigma^2(\theta) \geq \frac{1}{J(\theta)}. \quad (6)$$

The Fisher information is related to the entropy of the distribution.

Now, consider the high threshold area. Given a set of points, $\{\mathbf{x}_i\}$, it is clear that the threshold should be between 0 and the value $b_{max} = \max ||\mathbf{x}_i - \mathbf{x}_j||$. The high threshold range is defined by a set of threshold $b_i \in [b^*, b_{max}]$.

First, the CRB explains why the VoP increases as the variance of inliers (VoI) increases.

Second, the VoP increases if we add outliers into the data set. It can be explained by the Fisher information inequality [4] which states the Fisher information of a mixture is equal to or greater than the individual Fisher information as

$$\frac{1}{J(X+Y)} \geq \frac{1}{J(X)} + \frac{1}{J(Y)}. \quad (7)$$

If we add outliers and assume that the Fisher information is non-zero, $\frac{1}{J(Y)} > 0$, we can get the relationship $\sigma^2(\theta_{X+Y}) > \sigma^2(\theta_X)$. If the bound b^* corresponding to the $\frac{1}{J(X)}$ is an optimal bound, a larger bound $b^* + \delta(\delta > 0)$ gives $\frac{1}{J(X+Y)}$ because the Y has uniform distribution. Therefore, we can get the relationship:

$$var(\hat{\theta}|\mathbf{X}_{inliers}) < var(\hat{\theta}|\mathbf{X}_{outliers}). \quad (8)$$

So, for case II, $var(\hat{\theta}) > var(\theta^*)$.

4.2.4 Case III: low threshold

Case I shows the relation between VoE and the Fisher information when uniform noise is added to the inliers with the Gaussian distribution. Case III shows a different case in which the sample in the Gaussian distribution are restricted by new threshold $b^* - \delta(\delta > 0)$.

We start with inlier data $\mathbf{X}_{inliers} \sim \mathcal{N}(0, \sigma^2)$ and consider a subset of the data, $\mathbf{S} \subset \mathbf{X}_{inliers}$. We link two conditions: $var(\hat{\theta}|\mathbf{X}_{inliers})$ and $var(\hat{\theta}|\mathbf{S})$.

We can consider two cases. In the first case, RANSAC consistently converges close to the correct solution $\hat{\theta}$. The variance of the solution is then given by the CRB. For a Gaussian distribution, this bound is $2(\sigma^2)^2/n$, where n is the number of samples. Here, the number of samples is smaller than the full number of inliers, therefore the variance is higher,

$$var(\hat{\theta}|\mathbf{X}_{inliers}) < var(\hat{\theta}|\mathbf{S}). \quad (9)$$

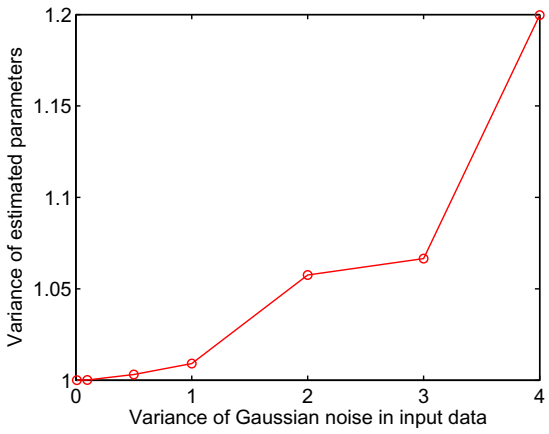

Figure 4. The relationship between VoI and VoP in 2-D line estimation.

In the second case, RANSAC may converge to different solutions on different trials, each with a small variance (because of the tight bound). Then, the variance of the solution is large. So, for case III, we have $var(\hat{\theta}) > var(\theta^*)$.

In summary, $var(\hat{\theta})$ is small and stable in case I, and larger in both case II and case III.

4.3. The relation between the RANSAC threshold b and the variance of the Gaussian noise σ^2

The direct way to solve the problem would be to estimate σ at the same time as we estimate θ. Unfortunately, there is no simple way to perform this evaluation. Instead, we evaluate the RANSAC threshold b and the parameter θ.

It is therefore important to establish the relationship between b and σ. In one direction, if we know σ, we can estimate b using the inverse cumulative chi-squared (χ^2_n) distribution [6]. The reverse is not simple. We can verify experimentally that they are correlated as shown in Fig 4.

Knowing θ^* defines a range for $b = \{b_i, \ldots, b_{i+k}\}$ in which the parameter θ is stable. We choose the corresponding b^* as the largest such value producing θ. This is in accordance to the CRB, which states that the minimum variance solution is the one with more samples. Referring to Fig 3, the range for b is $[0.02, 3]$, and we choose $b^* = 3.0$, which valid gives us 152 inliers (out of 160).

4.4. Influence of the number of RANSAC trials

To compute the VoP, RANSAC is repeated K times. We use a constant number $K = 30$ in all examples. We experimented with different set of values from $K = 10$ to $K = 100$ and obtained similar results.

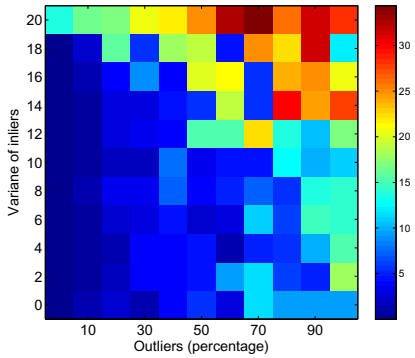

Figure 5. Stability analysis of 2-D line estimation.

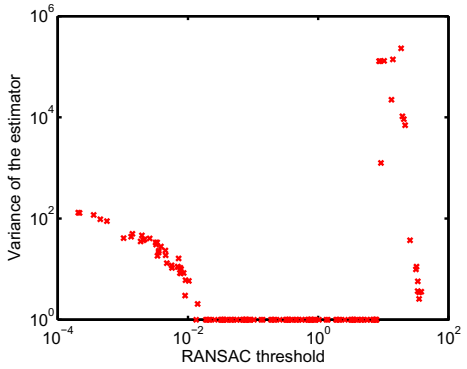

Figure 6. Variance of the estimator and bound.

4.5. Range of applicability

We investigate the range of applicability of our method using the error with respect to the ratio of outliers and the VoI. The total number of inliers points is 200. The ratio of outliers to inliers varies from 0 to 1. The inliers are subject to Gaussian noise with a variance σ^2 from 0 to 20. The error is defined as $E\{d(M, \mathbf{x}_i)\}$ where $d(M, \mathbf{x}_i)$ is the distance between the inlier point $\mathbf{x}_i$ to the ground truth line M in pixels units. Fig. 5 shows the stability of StaRSaC. The result shows that there is a large flat area in which StaRSaC can provide reasonable estimation results.

5. Experimental results

We validate our algorithm on motion estimation problems with synthetic data and real images.

5.1. Algorithm implementation

First, an initial range is defined between 0 and the value computed by a least-square method. Technically, we define a small number ϵ for the lower threshold ($\epsilon = 10^{-7}$). Second, to select the set of bounds, we use a uniform distribution in a log-scale. For instance, the range $(0.01, 0.1]$ has the same number of samples as $(0.001, 0.01]$ and $(0.1, 1.0]$. Third, to compute the variance of an estimator, RANSAC is repeated K times. We use a constant number $K = 30$ in our experiments. Fourth, to compute the VoP, we can use either the largest eigenvalue of the covariance matrix $\mathbf{\Theta}$, or the expected value of the Frobenius norm as $var(\hat{\theta}) \doteq E\{||\hat{\mathbf{H}}_i - \hat{\mathbf{H}}_j||_F^2\}, i \neq j$ where both $\hat{\mathbf{H}}_i$ and $\hat{\mathbf{H}}_j$ are normalized by setting $\mathbf{H}_{(3,3)} = 1$.

5.2. Evaluation method

Many papers show an error index as $E\{d(\hat{\theta}, \hat{\mathbf{X}}_{inliers})\}$ the distance between estimated parameters and inliers selected by the estimation. However, a result which has a small error may have significant bias [1]. We consistently evaluate our results with respect to the ground truth.

To generate synthetic points, we define a camera matrix

$$\mathbf{K} = \begin{bmatrix} 2000 & 0 & 0 \\ 0 & 2000 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

For images, corresponding points are extracted by SIFT [9] and their coordinates are normalized by the method proposed by Hartley [6].

5.3. Homography estimation with synthetic data

We define a homography matrix as $\mathbf{H} = \mathbf{K}(\mathbf{R} - \hat{\mathbf{t}}\hat{\mathbf{n}}^T/d)\mathbf{K}^{-1}$ [6]. The normal vector of the plane is $n = [0.2, 0.9, 0.3]^T$, the distance from camera to the plane is $d = f \times 10^6$, and the translation is $\mathbf{t} = [200, 200, f \times 10^4]^T$. The rotation (r_x, r_y, r_z) is defined within the range from 0° to 60°. These parameters are derived from an application where an airborne camera is observing the ground plane. The total number of points is 300. The ratio of outliers to inliers varies from 0 to 0.5. The inliers are subject to Gaussian noise with $\sigma = \{1, 3, 5, 7, 9\}$.

The error is defined as $E\{||\theta_i^\dagger - \hat{\theta}_i||^2\}$ where $\theta_i^\dagger$ is the known parameter (ground truth), $\hat{\theta}_i$ is the estimated parameter. We use only three rotation angles (r_x, r_y, r_z) for the error because rotation angles can be estimated more accurately than either the translation or the normal of the plane in our experiments. Since the decomposition does not give resonable parameters if the ratio of outliers is increased, we use the Frobenius norm as an alternative error criterion as $E\{||\mathbf{H}^\dagger - \hat{\mathbf{H}}_i||_F^2\}$ where $\mathbf{H}^\dagger$ is the ground truth homography matrix.

Fig 6 shows the variance of estimated parameters with respect to the bound. We now show results for one experiment, with $\sigma = 3$, 10% outliers uniformly distributed. The result shows that there is a large flat area in which StaRSaC can provide resonable estimation results. Fig 7

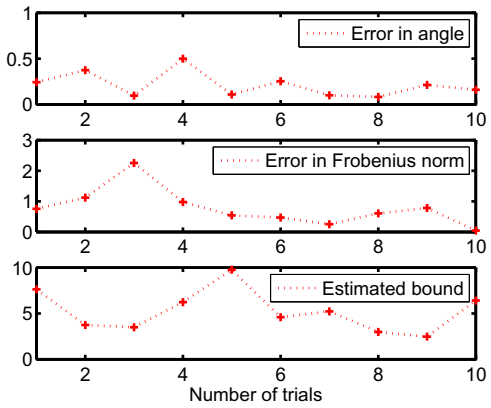

Figure 7. Estimated bound and error in rotation (and in Frobenius norm). We can estimate both the correct bound in RANSAC and accurate parameters of motion.

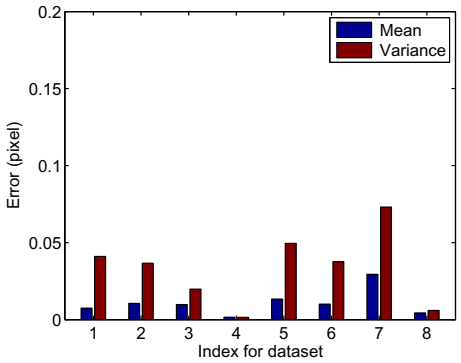

Figure 8. Results of homography estimation

shows the error in angle and Frobenius norm, and the estimated bounds. The results show that a reasonable bound can be estimated by StaRSaC.

5.4. Homography estimation with real images

We use the '2005 datasets' [7] which can be obtained from http://vision.middlebury.edu/stereo/data/. We select one image from each dataset (Art, Books, Dolls, Laundry, Moebius, Reindeer, Computer, Drumsticks, Dwarves) and generate a new image by the homography matrix $\mathbf{H}$ ($\mathbf{n} = [0.2, 0.9, 0.3]^T, d = f \times 10^6, \mathbf{t} = [200, 200, f \times 10^4]^T, \mathbf{R} = \mathbf{I}_{(3\times3)}$). The error is computed over the inlier set of points $\{(\mathbf{p}_i, \mathbf{q}_i)\}$ as $E\{||\mathbf{q}_i - \mathbf{H}\mathbf{p}_i||^2\}$.

Fig 8 shows that StaRSaC provides very accurate estimation results on this data set.

5.5. Fundamental matrix with synthetic data

We compare StaRSaC with M-estimation [6], LMedS [17], RANSAC [5], MLESAC [14], MAPSAC [12] algorithms which can be obtained from [1]. We generate 3D points ($N = 50$) randomly within a box ($2000 \times 2000 \times 50$). Two cameras are defined as $\mathbf{P}_1 = \mathbf{K}[\mathbf{I}_{(3\times3)}, [0,0,0]^T]$ and $\mathbf{P}_2 = \mathbf{K}[\mathbf{R}, \mathbf{t}]$. $\mathbf{R}$ is defined with angles $(-3°, -30°, 3°)$ and $\mathbf{t}$ is $[6000, 0, -2000]$. Points are generated by projection matrices $\mathbf{P}_1$, $\mathbf{P}_2$ and $\mathbf{F}$ is computed from the $\mathbf{P}_1$ and $\mathbf{P}_2$. The error is computed over the inlier set of points $\{(\mathbf{p}_i, \mathbf{q}_i)\}$ as $E\{d(\mathbf{F}^\dagger \mathbf{p}_i, \mathbf{q}_i)\}$ which represents the distance between an inlier and the corresponding epipolar line (the ground truth). Table 1 shows that the StaRSaC solution is consistently better than the others in all cases.

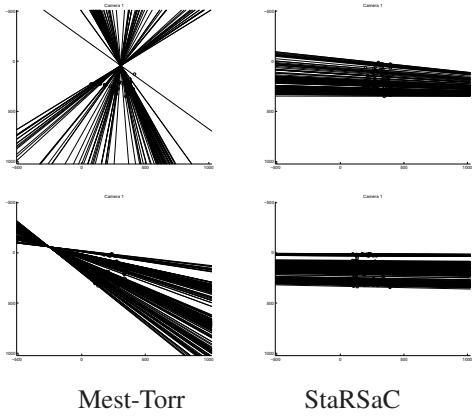

Figure 9. Examples of epipolar geometry estimation: left and right results are computed by Mest-Torr and StaRSaC, respectively, using the Aloe image (top) and the Monopoly image (bottom) in the 2006 dataset.

5.6. Fundamental matrix with real images

We again use the '2006 datasets' [7] which can be obtained from http://vision.middlebury.edu/stereo/data/. The error is computed over the inlier set of points $\{(\mathbf{p}_i, \mathbf{q}_i)\}$ as $E\{d(\mathbf{F}^\dagger \mathbf{p}_i, \mathbf{q}_i)\}$ which represents the distance between an inlier and the corresponding epipolar line (the ground truth). Fig 9 shows two examples of estimated epipolar geometries. Table 2 shows the results of fundamental matrix estimation with these real images. The StaRSaC solution is significantly better than all other methods.

5.7. Computational complexity

Although StaRSaC calls standard RANSAC algorithm several times to compute the VoP, StaRSaC can be immediately parallelized. Moreover, the search space of StaRSaC is one dimensional and independent of problems.

6. Conclusion

We propose a fully automatic algorithm for parameter estimation in computer vision applications. We show that the

Method	$\sigma = 0.5$ OL 30%	$\sigma = 0.5$ OL 50%	$\sigma = 1.0$ OL 30%	$\sigma = 1.0$ OL 50%	$\sigma = 2.0$ OL 30%	$\sigma = 2.0$ OL 50%
Mest-ls	0.680	1.414	1.161	1.465	0.989	1.265
Mest-eig	0.329	0.319	0.198	0.484	0.142	0.681
Mest-Torr	0.026	0.319	0.091	0.133	0.118	0.470
LMedS-ls	0.030	0.364	0.062	1.171	0.163	0.530
LMedS-eig	0.024	0.398	0.101	0.241	0.459	1.024
RANSAC	0.997	1.418	1.161	1.316	0.996	1.786
MLESAC	0.087	0.414	0.358	0.110	0.095	0.504
MAPSAC	0.065	0.673	0.090	0.305	0.343	0.460
StaRSaC	**0.004**	**0.154**	**0.012**	**0.027**	**0.051**	**0.381**

Table 1. Fundamental matrix estimation result (synthetic data, unit: pixel).

Method	mean(e)	mean(σ^2)	med(e)	med(σ^2)
Mest-ls	18.544	37.455	13.962	41.933
Mest-eig	14.079	25.597	4.576	25.612
Mest-Torr	11.820	22.912	1.926	8.381
LMedS-ls	9.609	14.098	1.615	4.381
LMedS-eig	9.564	13.554	1.256	4.533
RANSAC	21.658	41.448	16.814	47.172
MLESAC	12.032	12.102	1.112	3.650
MAPSAC	10.197	16.737	1.374	4.187
StaRSaC	**8.046**	**11.671**	**0.511**	**0.864**

Table 2. Fundamental matrix estimation result (21 datasets, unit: pixel).

variance of the estimated parameters (VoP) exhibits ranges of stability with respect to this bound. Within this range of stability, we can estimate both the parameters, and the variance of the Gaussian noise. We show how to compute this stable range using RANSAC. We validate our results by extensive tests and comparison with state of the art estimators on both synthetic and real data sets. These include line fitting, homography estimation, and fundamental matrix estimation. These results are very encouraging. We are working on formally establishing some of the proofs.

Acknowledgment

This research was funded in part by a grant from Airbus S.A.S.

References

[1] X. Armangué and J. Salvi. Overall view regarding fundamental matrix estimation. *Image and Vision Computing*, 21(2):205–220, 2003.

[2] D. Ballard and C. Brown. Computer vision. *Prentice-Hall: New Jersey*, 1982.

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

[4] A. Dembo, T. Cover, and J. Thomas. Information theoretic inequalities. *IEEE Transactions on Information Theory*, 37(6):1501–1518, 1991.

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

[6] R. Hartley and A. Zisserman. Multiple view geometry in computer vision, 2nd ed. *Cambridge University Press*, 2004.

[7] H. Hirschmller and D. Scharstein. Evaluation of cost functions for stereo matching. *CVPR*, 2007.

[8] R. Hoseinnezhad and A. Bab-Hadiashar. Consistency of robust estimators in multi-structural visual data segmentation. *Pattern Recognition*, 40:3677–3690, 2007.

[9] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *TPAMI*, 27(10):1615–1630, 2005.

[10] J. Miller and C. Stewart. Muse: Robust surface fitting using unbiased scale estimates. *CVPR*, pages 300–306, 1996.

[11] W. Tong, C. Tang, and G. Medioni. Epipolar geometry estimation for non-static scenes by 4d tensor voting, CVPR, 2001.

[12] P. Torr. Bayesian model estimation and selection for epipolar geometry and generic manifold fitting, International Journal of Computer Vision, 50(1):35-61, 2002.

[13] P. Torr and D. Murray. The development and comparison of robust methods for estimating the fundamental matrix. *International Journal of Computer Vision*, 24(3), 1997.

[14] P. Torr and A. Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer Vision and Image Understanding*, 78:138–156, 2000.

[15] H. Wang and D. Suter. Robust adaptive-scale parametric model estimation for computer vision. *TPAMI*, 26:1459–1474, 2004.

[16] R. Zarmir. A proof of the fisher information inequality via a data processing argument. *IEEE Transactions on Information Theory*, 44(3):1246–1250, 1998.

[17] Z. Zhang. Determining the epipolar geometry and its uncertainty: a review, International Journal of Computer Vision, 27(2):161-198, 1998.