



UNIVERSITÀ
DEGLI STUDI
DI UDINE

Università degli studi di Udine

Reliable exterior orientation by a robust anisotropic orthogonal Procrustes Algorithm

Original

Availability:

This version is available <http://hdl.handle.net/11390/865150> since 2016-12-01T09:41:31Z

Publisher:

Published

DOI:10.5194/isprsarchives-XL-5-W1-81-2013

Terms of use:

The institutional repository of the University of Udine (<http://air.uniud.it>) is provided by ARIC services. The aim is to enable open access to all the world.

Publisher copyright

(Article begins on next page)

RELIABLE EXTERIOR ORIENTATION BY A ROBUST ANISOTROPIC ORTHOGONAL PROCRUSTES ALGORITHM

A. Fusiello^a, E. Maset^b, F. Crosilla^b

^a Dipartimento di Ingegneria Elettrica, Gestionale e Meccanica,
University of Udine, Via Delle Scienze, 208 - 33100 Udine, Italy
andrea.fusiello@uniud.it

^b Dipartimento di Ingegneria Civile e Architettura,
University of Udine, Via Delle Scienze, 208 - 33100 Udine, Italy
fabio.crosilla@uniud.it, maset.eleonora@spes.uniud.it

KEY WORDS: Robust Methods, Anisotropic Orthogonal Procrustes Analysis, Camera Exterior Orientation

ABSTRACT:

The paper presents a robust version of a recent anisotropic orthogonal Procrustes algorithm that has been proposed to solve the so-called camera exterior orientation problem in computer vision and photogrammetry. In order to identify outliers, that are common in visual data, we propose an algorithm based on Least Median of Squares to detect a minimal outliers-free sample, and a Forward Search procedure, used to augment the inliers set one sample at a time. Experiments with synthetic data demonstrate that, when the percentage of outliers is greater than 30% or the data size is small, the proposed method is more accurate in detecting outliers than the customary detection based on median absolute deviation.

1 INTRODUCTION

In a recent paper (Garro et al., 2012) a new method for the exterior orientation problem solution of a calibrated camera has been proposed. The method is fundamentally based on the anisotropic extension of the Orthogonal Procrustes Analysis with different scaling of the coordinate values for each point.

The problem of estimating the position and orientation of the perspective point of a camera, given its intrinsic parameters and a set of image and object correspondences, is very well known in computer vision and photogrammetry. Many algorithms have been proposed; these are mainly based on direct and indirect procedures. Direct procedures are usually faster but less accurate, as they do not minimize a significant cost function (Fischler and Bolles, 1981, Quan and Lan, 1999, Fiore, 2001, Ansar and Dailidis, 2003, Lepetit et al., 2009), whereas iterative methods, that explicitly minimize a meaningful geometric error, are more accurate but slower (Lowe, 1991, Haralick et al., 1989, Kumar and Hanson, 1994, Horaud et al., 1997). The proposed method (Garro et al., 2012) reaches the best trade-off between speed and accuracy and is easier to implement than its closest competing algorithm (Hesch and Roumeliotis, 2011).

On the other hand, Orthogonal Procrustes Analysis (OPA) is a very useful tool to perform the direct least squares solution of similarity transformation problems in any dimensional space. At first, it was used for the multidimensional rotation and scaling of different matrix configuration pairs (Schönemann, 1966, Schönemann and Carroll, 1970). Successively, the solution was generalized for the simultaneous least squares matching of more than two corresponding matrices (Gower, 1975, Ten Berge, 1977). In photogrammetry OPA has been already applied for matching different 3D object models from images (Crosilla and Beinat, 2002) and from 3D laser point clouds (Beinat and Crosilla, 2001) respectively. In both cases the authors applied the so called generalized version of the OPA problem to simultaneously match more than two coordinate matrices of corresponding points expressed in different reference systems.

Very recently, an interesting extension of the OPA analysis with anisotropic scaling has been proposed (Bennani Dosse and Ten Berge, 2010), also in its generalized form (Bennani Dosse et al., 2011). The algorithm operates with anisotropic scaling along space dimension and is based on an iterative block relaxation technique (de Leeuw, 1994) that starts with a uniform initialization. The iterative algorithm proposed by (Garro et al., 2012) applies instead anisotropic scaling for each measurement and is based on the same kind of relaxation techniques. Although this algorithm proved to reach a good trade-off between speed and accuracy, it has the drawback of a Procrustes method, i.e. it is a least squares estimation technique failing to cope with outliers.

To overcome this problem, we present a robust version of the above mentioned algorithm, that can tolerate a minority of arbitrarily large outliers. At the beginning, Least Median of Squares (LMedS) (Rousseeuw, 1984) is used in order to detect an initial minimal subset of data, free from outliers. The initial subset is then augmented one sample at a time with Forward Search (FS) (Hadi and Simonoff, 1993, Atkinson and Riani, 2000), an iterative scheme that proved to be effective in dealing with the well-known masking and swamping problems that can occur when multiple outliers are present in the data set. An empirical comparison of FS against the customary classification based on the Median Absolute Deviation (MAD) scale estimation has been carried out. It was noticed that, up to 20% outliers, MAD and FS produce comparable results, but when the percentage of outliers increases, FS becomes more reliable for outliers detection.

2 PROCRUSTEAN EXTERIOR IMAGE ORIENTATION

In this section we will briefly summarize the procrustean approach to the exterior image orientation problem presented in (Garro et al., 2012). The reader is referred to the original paper for more details.

Given a number of 2D-3D point correspondences $\mathbf{m}_j \leftrightarrow \mathbf{M}_j$ and the intrinsic camera parameters K , the exterior image orientation problem requires to find a rotation matrix R and a translation

vector \mathbf{t} (which specify attitude and position of the camera) such that:

$$\zeta_j \tilde{\mathbf{m}}_j = K[R|\mathbf{t}]\tilde{\mathbf{M}}_j \quad \text{for all } i. \quad (1)$$

where ζ_j denotes the depth of $\tilde{\mathbf{M}}_j$, and the $\tilde{\cdot}$ denotes homogeneous coordinates (with a trailing "1").

After some rewriting, (1) becomes:

$$\underbrace{\begin{bmatrix} \mathbf{M}_1^T \\ \vdots \\ \mathbf{M}_n^T \end{bmatrix}}_S = \underbrace{\begin{bmatrix} \zeta_1 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \zeta_n \end{bmatrix}}_Z \underbrace{\begin{bmatrix} \tilde{\mathbf{p}}_1^T \\ \vdots \\ \tilde{\mathbf{p}}_n^T \end{bmatrix}}_P R + \underbrace{\begin{bmatrix} \mathbf{c}^T \\ \vdots \\ \mathbf{c}^T \end{bmatrix}}_{\mathbf{1c}^T}. \quad (2)$$

where $\tilde{\mathbf{p}}_j = K^{-1}\tilde{\mathbf{m}}_j$, $\mathbf{c} = -R^T\mathbf{t}$, and $\mathbf{1}$ is the unit vector. The previous equation can be written more compactly in matrix form:

$$S = ZPR + \mathbf{1c}^T \quad (3)$$

where

P is the matrix by rows of (homogeneous) image coordinates defined in the camera frame,

S is the matrix by rows of point coordinates defined in the exterior system,

Z is the diagonal (positive) depth matrix,

\mathbf{c} is the coordinate vector of the projection centre,

R is the orthogonal rotation matrix.

One can recognize an instance of an *anisotropic* OPA problem with *data scaling* (Bennani Dosse and TenBerge, 2010), where the usual (isotropic) scale factor is substituted by an anisotropic scaling characterized by a diagonal matrix Z of different scale factors.




Figure 1: The orientation of the camera and the depth of the points are estimated in such a way to minimize the length of the Δ s for all the points, in a least squares sense.

To obtain the least squares solution for Eq. (3) – along the same line as in (Schönemann and Carroll, 1970) – let us make explicit the residual matrix Δ :

$$S = ZPR + \mathbf{1c}^T + \Delta. \quad (4)$$

The geometric interpretation, with reference to Figure 1, is that the estimated depth defines a 3D point along the optical ray of the

image point \mathbf{p} , and the segment (perpendicular to the optical ray) joining this point and the corresponding reference 3D point \mathbf{M} is the residual. Therefore, Δ is a matrix whose rows are difference vectors between reference 3D points (S) and the back-projected 2D points (P) based on their estimated depths (Z) and the estimated camera attitude and position (R, \mathbf{c}).

The solution of the anisotropic OPA problem finds Z, R and \mathbf{c} in such a way to minimize the sum of squares of the residual Δ , i.e., the sum of the squared norm of the difference vectors mentioned above. This can be written as

$$\min \|\Delta\|_F^2 \quad \text{subject to } R^T R = I \quad (5)$$

The problem is equivalent to the minimization of the Lagrangian function

$$F = \text{tr}(\Delta^T \Delta) + \text{tr}(L(R^T R - I)) \quad (6)$$

where L is the matrix of Lagrangian multipliers. This can be solved by setting to zero the partial derivatives of F with respect to the unknowns R, \mathbf{c} and the diagonal matrix Z . The derivation of the formulae is reported in (Garro et al., 2012). It turns out that in the anisotropic case the unknowns are entangled in such a way that one must resort to the so called “block relaxation” scheme (de Leeuw, 1994), where each variable is alternatively estimated while keeping the others fixed. Algorithm 1 summarizes the procedure.

Algorithm 1 PROCURSTEAN EXTERIOR ORIENTATION

Input: a set of 2D-3D correspondences (P, S)

Output: position \mathbf{c} and attitude R of the camera

1. Start with $Z = 0$ (or any guess on the depth);
 2. Compute $R = U \text{diag}(1, 1, \det(UV^T)) V^T$ with $UDV^T = P^T Z (I - \mathbf{1} \mathbf{1}^T/n) S$;
 3. Compute $\mathbf{c} = (S - ZPR)^T \mathbf{1}/n$;
 4. Compute $Z = \text{diag}(PR(S^T - \mathbf{c} \mathbf{1}^T)) \text{diag}(PP^T)^{-1}$;
 5. Iterate over steps 2, 3, 4 until convergence.
-

In order to work toward the robust version of Algorithm 1, please note that:

- step 2 can be seen as the solution of a simple OPA problem, namely finding R such that $ZP = R(I - \mathbf{1} \mathbf{1}^T/n) S$;
- step 3 is simply the mean along the columns of the matrix $S - ZPR$.

We will return to this observation later.

3 ROBUST ESTIMATION TECHNIQUES

Almost inescapably image measurements contain gross errors (*outliers*) that may be arbitrarily large and cannot be “averaged out”, as is typically done with small-scale noise. They lead to incorrect results in parameter estimation; therefore robust estimation techniques have been developed over the last years in both the computer vision and the statistics literature (Ben-Gal, 2005, Stewart, 1999),




Figure 2: The $(s + 1)^{th}$ residual (diamonds) and the value of the threshold $(t_{(1-\alpha/2(s+1),s-k)}\sigma_s)$ (red line) at each iteration of FS.

The Least Median of Squares is a robust regression method that achieves the maximum theoretical breakdown point of 0.5, i.e., it can tolerate up to 50% percent of outliers in the data without being affected. LMedS estimates the parameters of the model by minimizing the median of the absolute residuals. Let $\mathbf{X} = \{\mathbf{x}_i\}$ be a set of data points and let \mathbf{a} be a k -dimensional parameter vector to be estimated, while $r_{i,\mathbf{a}}$ is the Euclidean distance between a point \mathbf{x}_i and the model determined by \mathbf{a} . The LMedS estimate is:

$$\hat{\mathbf{a}} = \underset{\mathbf{a}}{\operatorname{argmin}} \operatorname{median}_{\mathbf{x}_i \in X} r_{i,\mathbf{a}}^2 \quad (7)$$

The value of \mathbf{a} is determined using a random sampling technique. A number of k -points subsets of the n data points is randomly selected and a parameter vector \mathbf{a}_S is fitted to the points in each subset S . Each \mathbf{a}_S is tested by calculating the squared residual distance r_{i,\mathbf{a}_S}^2 of each of the $n - k$ points in $X - S$ and finding the median. The \mathbf{a}_S corresponding to the smallest median is chosen as the estimate $\hat{\mathbf{a}}$. If n is the number of points in the initial set of data, there are $\binom{n}{k}$ different subsets of k points that should be considered. In order to speed up the process, the number of subsets to test can be reduced to

$$M = \frac{\log(1-p)}{\log(1-(1-\epsilon)^k)}, \quad (8)$$

where p is the probability of calculating the correct estimate and ϵ is the percentage of outliers.

Since LMedS is known to have a low statistical efficiency, it is important to refine the fit on a “clean” (i.e., outliers-free) set of data point. This is customarily done by computing a robust scale estimate σ^* based on the median absolute deviation (MAD):

$$\sigma^* = 1/\phi^{-1}(0.75)(1 + \frac{5}{n-k}) \sqrt{\operatorname{median}_{\mathbf{x}_i \in X-\hat{S}} r_{i,\hat{\mathbf{a}}}^2} \quad (9)$$

where \hat{S} is the subset used to calculate $\hat{\mathbf{a}}$, ϕ^{-1} is the inverse of the cumulative normal distribution, and $1 + 5/(n - k)$ is a correction factor. Then $\hat{\mathbf{a}}$ is refined with a least squares estimate using only the data points such that:

$$r_{i,\hat{\mathbf{a}}}^2 < (\theta\sigma^*)^2 \quad (10)$$

where θ is a constant that depends on the confidence level chosen for the test (typically around 2 or 3).

According to the taxonomy set forth in (Davies and Gather, 1993), this is a univariate single-step outlier detection procedure, that will be referred to as MAD in the following. Sequential procedures, on the contrary, perform successive elimination or addition of points. In particular, *inward testing* starts with the whole set of data points and iteratively test the point with the highest residual




Figure 3: Residuals at the last iteration of FS arranged in ascending order. The red horizontal line is the threshold that has been exceeded.

for being an outlier; it stops when the first inlier is encountered. On the contrary, *outward testing* starts with an outlier-free set of points and iteratively tests the remaining points for inclusion in the clean subset, based on statistics computed on the clean subset; it stops when all the points have been processed.

Forward Search (FS) (Hadi and Simonoff, 1993, Atkinson and Riani, 2000) is an outlier detection method of the latter category. The algorithm starts from an initial subset of observations of size $m < n$ that is known to be outlier free. In our case this is the set \hat{S} from LMedS, and $m = k$. Then the iterative process starts and at every i^{th} iteration the s samples ($s = m + i - 1$) with the lowest residual are used to estimate the parameters of the model \mathbf{a}_s and the $s + 1$ residual r_{s+1,\mathbf{a}_s} is monitored to detect the point when the iteration starts including outliers.

Several monitoring heuristics have been proposed in literature. We consider here one of the simplest, that dates back to (Hadi and Simonoff, 1993); the forward search stops when:

$$r_{s+1,\mathbf{a}_s}^2 \geq (t_{(1-\alpha/2(s+1),s-k)}\sigma_s)^2 \quad (11)$$

where $t_{(1-\alpha/2(s+1),s-k)}$ the $1 - \alpha/2(s + 1)$ quantile of the t -distribution with $s - k$ degrees of freedom, and σ_s is the standard deviation of the least s residuals (the residual divided by σ_s is also called “studentized” residual).

Figure 2 plots the $(s + 1)^{th}$ residual at each iteration: please note how the threshold adapts to the variance of the clean subset and how the residual of the first outlier pops out. Same analysis can be made on Figure 3, which plots the ordered residuals of the whole point set: the outliers are clearly separated and well above the threshold.

The i^{th} iteration of the algorithm can be summarized as follows:

1. Estimate the parameters of the model \mathbf{a}_s using the $s = m + i - 1$ points in the clean subset;
2. Calculate the residuals r_{s+1,\mathbf{a}_s} of all the n points and arrange them in ascending order;
3. Test the $(s + 1)^{th}$ residual:

If $r_{s+1,\mathbf{a}_s}^2 \geq (t_{(1-\alpha/2(s+1),s-k)}\sigma_s)^2$ then declare all the remaining $n - s$ observations as outliers and stop.

Otherwise increment i (include the $(s + 1)^{th}$ point in the inliers).

4. If $m + i - 1 = n$, then stop, otherwise go to step 1.

4 ROBUST PROCRUSTEAN EXTERIOR IMAGE ORIENTATION

Resuming to the observations at the end of Section 2, one can immediately figure out that the procrustean exterior image orientation algorithm can be made robust against outliers by substituting:

- step 2 with a robust solution (e.g., LMedS) to the problem of finding the rotation between the two point sets ZP and $(I - \mathbf{1}\mathbf{1}^T/n)S$;
- step 3 with the *median* instead of the mean along the columns of $S - ZPR$.

This works pretty straightforward, and being an iterative process, convergence can be improved by re-using the estimate of the inliers set from one iteration to the next. As customary, we used a threshold based on MAD (as described in Sec. 3) to reject outliers at each iteration.

However this simple strategy cannot recover the inlier set with high reliability for *every* condition. There are cases where FS, albeit more computationally expensive, should be used because of its superior accuracy, as experiments will show. Therefore we propose to refine the clean subset produced by the LMedS procedure with FS; the resulting robust procrustean exterior orientation procedure is summarized in Algorithm 2.

Algorithm 2 ROBUST PROCRUSTEAN EXTERIOR ORIENTATION

Input: a set of 2D-3D correspondences (P, S)

Output: position \mathbf{c} and attitude R of the camera

- Find a clean subset of three points¹:
 1. Start with $Z = 0$ (or any guess on the depth);
 2. Compute R s. t. $P^T Z = R(I - \mathbf{1}\mathbf{1}^T/n)S$ using LMedS (on current inliers);
 3. Compute $\mathbf{c} = \text{mean}_{\text{inliers}}((S - ZPR))^T$;
 4. Compute $Z = \text{diag}(PR(S^T - \mathbf{c}\mathbf{1}^T)) \text{diag}(PP^T)^{-1}$;
 5. Declare as inliers points that pass the test of Eq. (10);
 6. Iterate over steps 2, 3, 4, 5 until convergence;
 - Refine inliers estimate using FS;
 - Compute R, \mathbf{c}, Z (algorithm (1)) considering the set of points determined in the previous step.
-

5 EXPERIMENTAL VALIDATION

We run synthetic experiments, using $n = \{20, 40, 60, 80, 100\}$ 3D points. The image coordinates obtained from the projection of the points have been perturbed with gaussian noise with standard deviation $\sigma = 2.0$ and different percentages of outliers

¹Exterior orientation requires at least three points.

{10%, 20%, 30%, 40%, 50%} have been added to the data. Outliers are generated by adding a random value with uniform distribution in $[0, 100]$ to both coordinates of a point. The focal length of the camera was set to 500 pixel and the image size is $\approx 600 \times 600$ (Figure 4 shows an example of synthetic image). For each value of n and percentage of outliers the test has been run 100 times.




Figure 4: An example of the synthetic image data. The green squares are the inliers, the red stars are the outliers.

First we tested the first part of the algorithm, the one responsible for extracting a clean subset, based on LMedS, and, as expected, it provides a clean subset (with up to 50% of outliers) in most of the cases (76%). The percentage reaches 92% if one considers only a maximum of 40% of outliers. Figure 5 reports the false negative rate (see below) for the first part of the algorithm.




Figure 5: False negative rate vs percentage of outliers and number of points for the first part of the algorithm.

Then we compared empirically the robust procrustean exterior orientation method (with FS) against the baseline method with the MAD-based test only, in order to assess whether FS brings some advantage in detecting outliers.

The outlier detection stage post LMedS can be considered as a classification task, where each point has to be labelled as inlier or outlier (i.e., we test the null hypothesis that a point is an outlier) θ and α were set to 2.0 and 0.0001 respectively, so as to yield a similar rejection threshold for a chosen reference case, corresponding to 100 points, 10% outliers, $\sigma = 2.0$.

Let c be the number of *false negatives*, i.e., actual outliers that are erroneously deemed as inlier (type I error), a be the number

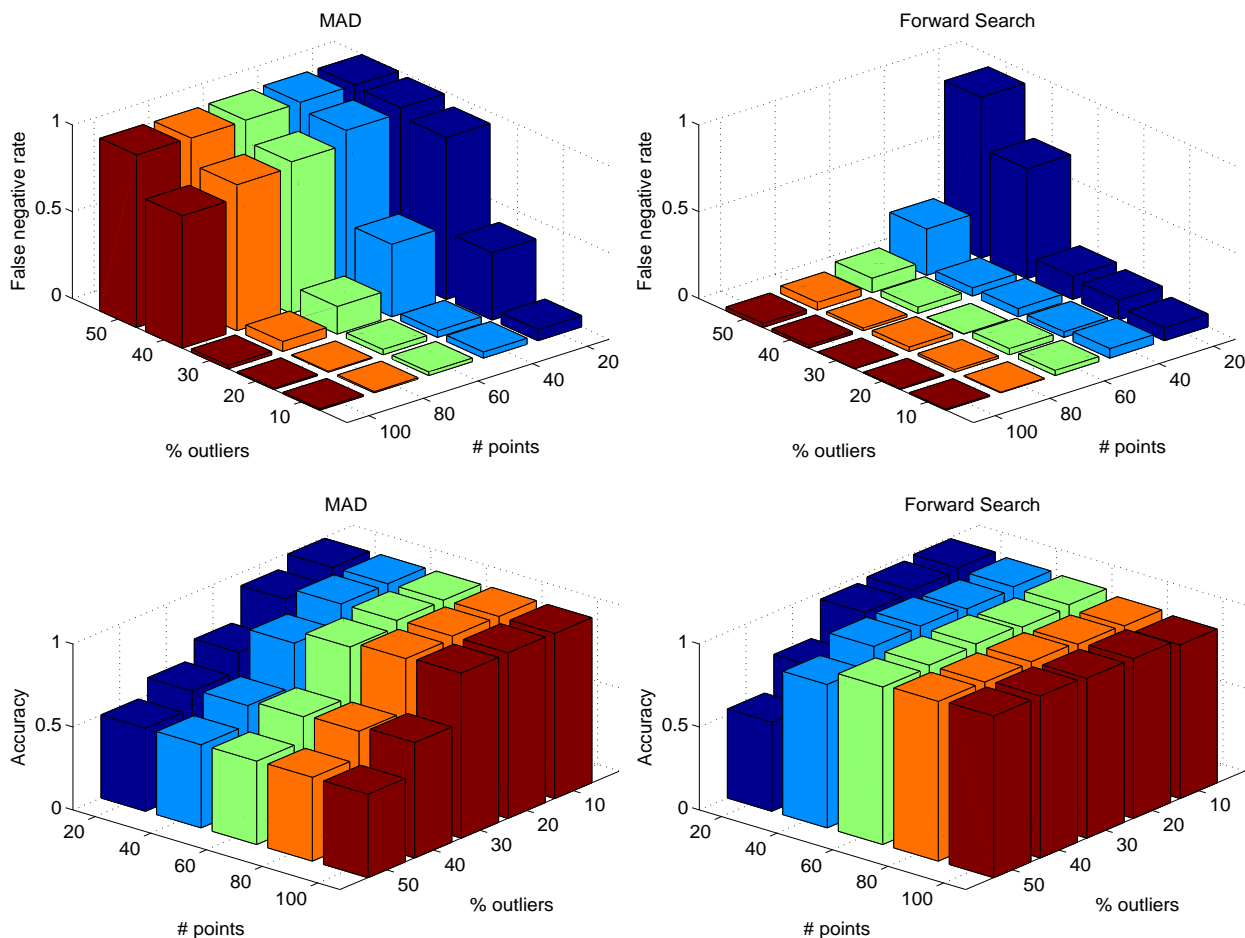


Figure 6: Top row: false negative rate vs percentage of outliers and number of points for MAD (left) and FS (right). Bottom row: accuracy vs percentage of outliers and number of points for MAD (left) and FS (right). Gaussian noise has $\sigma = 2.0$.

of *true positives*, i.e., outliers that are correctly detected and d be the number of *true negatives*, i.e., inliers that are correctly detected.

In order to compare the two algorithms, we monitored two figures of merit:

- the *false negative rate* defined as $c/(a + c)$, and
- the *accuracy*, defined as $(a + d)/n$.

Please note that false negatives (type I error) are way more dangerous than false positives, as they may skew the final least squares estimate, whereas false positives (type II error) can only impact on the statistical efficiency, as they reduce the number of “good” measurements that are considered in the final least squares estimate. For this reason, the main indicator is the false negative rate, which should be as small as possible. The accuracy indeed is monitored in order to check that the method is not achieving a low false negative rate by being too “fussy”, i.e. rejecting too many samples as outliers.

Results are reported in Fig. 6. The bar corresponding to 20 points and 50% outliers can be considered as an extreme case and should probably left out from general considerations. The opposite case, corresponding to 100 points and 10% outliers, is the “easy” one.

The rundown of the experiments is that with a fairly large set of point (> 60) and moderate outliers contamination ($\leq 30\%$)

MAD and FS produce comparable results, but when the size of the data is smaller or the contamination increases, FS performs better than MAD with respect to all the indicators. This makes sense, as one could expect that the advantage of a meticulous analysis as the FS should be more noticeable on smaller samples.

Both methods are fairly unstable and produces mixed results when $n < 20$ (results not shown here), because any statistical analysis is ineffective on such small samples.

Comparison with Fig. 5 reveals that shortfall of FS is mostly due to the first part of the algorithm failing to produce a clean subset. It also shows that FS almost never worsen the false negative rate of the clean subset and achieves high accuracy, showing that the initial clean subset has been actually extended and (almost) no outliers have been added.

6 CONCLUSIONS

We presented a robust version of a procrustean exterior orientation algorithm based on LMedS and Forward Search. Our experiments show that FS is highly effective and accurate in detecting outliers even in difficult cases corresponding to small data size or high outliers contamination.

REFERENCES

- Ansar, A. and Daniilidis, K., 2003. Linear pose estimation from points or lines. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25(5), pp. 578 – 589.
- Atkinson, A. and Riani, M., 2000. *Robust Diagnostic Regression Analysis*. Springer Series in Statistics, Springer.
- Beinat, A. and Crosilla, F., 2001. Generalized procrustes analysis for size and shape 3d object reconstruction. In: *Optical 3-D Measurement Techniques*, pp. 345–353.
- Ben-Gal, I., 2005. Outlier detection. In: O. Maimon and L. Rokach (eds), *Data Mining and Knowledge Discovery Handbook*, Springer US, New York, chapter 7, pp. 131–146–146.
- Bennani Dosse, M. and TenBerge, J., 2010. Anisotropic orthogonal procrustes analysis. *Journal of Classification* 27(1), pp. 111–128.
- Bennani Dosse, M., Kiers, H. A. L. and Ten Berge, J., 2011. Anisotropic generalized procrustes analysis. *Computational Statistics & Data Analysis* 55(5), pp. 1961–1968.
- Crosilla, F. and Beinat, A., 2002. Use of generalised procrustes analysis for the photogrammetric block adjustment by independent models. *ISPRS Journal of Photogrammetry & Remote Sensing* 56(3), pp. 195–209.
- Davies, L. and Gather, U., 1993. The identification of multiple outliers. *Journal of the American Statistical Association* 88(423), pp. pp. 782–792.
- de Leeuw, J., 1994. Block-relaxation algorithms in statistics. In: *Information Systems and Data Analysis*, Springer-Verlag, p. 308325.
- Fiore, P. D., 2001. Efficient linear solution of exterior orientation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(2), pp. 140–148.
- Fischler, M. A. and Bolles, R. C., 1981. Random Sample Consensus: a paradigm model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24(6), pp. 381–395.
- Garro, V., Crosilla, F. and Fusiello, A., 2012. Solving the pnp problem with anisotropic orthogonal procrustes analysis. In: *Second Joint 3DIM/3DPVT Conference: 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT)*.
- Gower, J., 1975. Generalized procrustes analysis. *Psychometrika* 40(1), pp. 33–51.
- Hadi, A. S. and Simonoff, J. S., 1993. Procedures for the identification of multiple outliers in linear models. *Journal of the American Statistical Association* 88(424), pp. pp. 1264–1272.
- Haralick, R., Joo, H., Lee, C., Zhuang, X., Vaidya, V. and Kim, M., 1989. Pose estimation from corresponding point data. *IEEE Transactions on Systems, Man and Cybernetics* 19(6), pp. 1426–1446.
- Hesch, J. A. and Roumeliotis, S. I., 2011. A direct least-squares (dls) solution for PnP. In: *Proc. of the International Conference on Computer Vision*.
- Horau, R., Dornaika, F. and Lamiroy, B., 1997. Object pose: The link between weak perspective, paraperspective, and full perspective. *International Journal of Computer Vision* 22, pp. 173–189.
- Kumar, R. and Hanson, A. R., 1994. Robust methods for estimating pose and a sensitivity analysis. *CVGIP: Image Understanding*. 60, pp. 313–342.
- Lepetit, V., Moreno-Noguer, F. and Fua, P., 2009. *Epnip*: An accurate $o(n)$ solution to the pnp problem. *International Journal of Computer Vision* 81(2), pp. 155–166.
- Lowe, D., 1991. Fitting parameterized three-dimensional models to images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13(5), pp. 441–450.
- Quan, L. and Lan, Z., 1999. Linear n-point camera pose determination. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21(8), pp. 774–780.
- Rousseeuw, P. J., 1984. Least median of squares regression. *Journal of the American Statistical Association* 79(388), pp. pp. 871–880.
- Schönemann, P., 1966. A generalized solution of the orthogonal procrustes problem. *Psychometrika* 31(1), pp. 1–10.
- Schönemann, P. and Carroll, R., 1970. Fitting one matrix to another under choice of a central dilation and a rigid motion. *Psychometrika* 35(2), pp. 245–255.
- Stewart, C. V., 1999. Robust parameter estimation in computer vision. *SIAM Review* 41(3), pp. 513–537.
- Ten Berge, J., 1977. Orthogonal procrustes rotation for two or more matrices. *Psychometrika* 42(2), pp. 267–276.