

Discussion of “The power of monitoring: how to make the most of a contaminated multivariate sample”

Citation for published version (APA):

Raymaekers, J., Rousseeuw, P. J., & Vranckx, I. (2018). Discussion of “The power of monitoring: how to make the most of a contaminated multivariate sample”. *Statistical Methods and Applications*, 27(4), 589-594. <https://doi.org/10.1007/s10260-018-0425-3>

Document status and date:

Published: 01/12/2018

DOI:

[10.1007/s10260-018-0425-3](https://doi.org/10.1007/s10260-018-0425-3)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the “Taverne” license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license


Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Discussion of “The power of monitoring: how to make the most of a contaminated multivariate sample”

Jakob Raymaekers¹ · Peter J. Rousseeuw¹  · Iwein Vranckx¹

Accepted: 11 March 2018 / Published online: 17 March 2018
© Springer-Verlag GmbH Germany, part of Springer Nature 2018

Abstract In this comment on the discussion paper “The power of monitoring: how to make the most of a contaminated multivariate sample” by A. Cerioli, M. Riani, A. Atkinson and A. Corbellini, we describe how the hard rejection property of the MCD method can be mimicked by an S-estimator with appropriate rho-function. We also point the reader to fast and deterministic algorithms for the MCD, S- and MM-estimators that are specifically suited for monitoring experiments. They were made available a few years ago and successfully used for monitoring in our papers. Finally, the question is raised how monitoring can be applied or extended for increasing numbers of cases, variables and tuning parameters.

Keywords Deterministic algorithm · minimum covariance determinant · nearby contamination · S-estimator

1 Introduction

We would like to congratulate Cerioli, Riani, Atkinson and Corbellini (henceforth CRAC) on their well-written and lavishly illustrated exposition about the usage and benefits of monitoring, and thank the editors for inviting us to comment on this interesting work.

✉ Peter J. Rousseeuw
peter@rousseeuw.net

Jakob Raymaekers
jakob.raymaekers@kuleuven.be

Iwein Vranckx
iwein.vranckx@kuleuven.be

¹ Department of Mathematics, University of Leuven, Celestijnenlaan 200b, 3001 Louvain, Belgium

2 The problem of nearby contamination

The leading example in the paper is the geyser (Old Faithful) dataset. From the scatterplot of this bivariate dataset we see that it consists of two clusters, the smaller of which contains about 30–35% of the observations. If one interprets the smaller cluster as contamination this is a relatively high contamination level, though it should not be prohibitive since the estimators considered in the paper can be tuned to a breakdown value well above 35%. However, the contamination happens to lie quite close to the inlying data. Having a large fraction of contamination located fairly close by makes the geyser data particularly challenging, as illustrated by CRAC. If we use a scatter estimator with a breakdown value of e.g. 40%, replacing any 35% of clean data by data points positioned anywhere cannot completely destroy the scatter matrix (in the sense of making its first eigenvalue arbitrarily large or its last eigenvalue arbitrarily close to zero), but that does not imply that the scatter matrix will have a small bias. Indeed, it is known that the bias of the estimators under study is the largest for nearby contamination, as shown by Hubert et al. (2014).

The other real data example in the paper is the cows dataset with 4 variables. Our first instinct was to carry out a PCA to get some idea about the shape of the data. Figure 1 shows the first two principal components of the cows data, which explain 96% of the total variance. (Here we used the ROBPCA method of Hubert et al. (2005), but classical PCA gave a very similar picture.) The plot shows that this dataset is equally challenging. CRAC interpret it as a well-behaved point cloud plus contamination. Again most of the contamination is nearby, and then it fans out. An alternative interpretation is that it could be a sample from a skewed distribution, and in that model none of these points need to be considered outlying.




Fig. 1 The first two principal components of the cows data

The paper starts by analyzing the geyser data through monitoring an S-estimator, the subsequent MM-estimator, and the MCD. Each of these estimators is tuned by its breakdown value. For the S-estimator this yields the surprising conclusion that tuning for a 50% breakdown value achieves the intended result, whereas 49% does not. A second conclusion is that the tuning of the MM-estimator appears more stable, but that is only because the MM-estimator starts from the S-estimator with 50% breakdown value, and it fails when starting from the 49% version.

These conclusions would appear to suggest that S-estimators are not suitable in such challenging situations. In fact, using the Matlab code kindly provided by CRAC we carried out an additional experiment, confirming that if the small cluster is moved a little bit in the direction of the larger one, even the S-estimator tuned for 50% breakdown no longer detects it, resulting in an uninformative monitoring plot consisting of horizontal lines only. Following the S-estimate by the MM step does not change that of course.

In both the geyser and cows examples, as well as the two simulated data sets in Section 6 whose contamination is equally nearby, CRAC find that the MCD gives better results due to its hard rejection.

3 The choice of the ρ -function

We would like to argue that it is not the definition of S-estimators that makes them fail the stated objective in these four examples, but rather the choice of the ρ -function. In our notation, a multivariate S-estimator of a p -variate dataset $\{y_1, \dots, y_n\}$ is defined as the pair (T, C) formed by a location vector T and a PSD scatter matrix C that together minimize $\text{Det}(C)$ subject to

$$\frac{1}{n} \sum_{i=1}^n \rho(d_i(T, C)) = K \tag{1}$$

in which $d_i(T, C) = \sqrt{(y_i - T)'C^{-1}(y_i - T)}$ is the statistical distance of y_i to T relative to C . One often takes $K = E[\rho(|Z|)]$ where Z follows the multivariate standard normal distribution. Note that we write ρ as a function of the statistical distance d_i and not its square d_i^2 .

People almost exclusively use the bisquare ρ in S-estimators, which redescends slowly in order to attain a high statistical efficiency at uncontaminated data. This is what makes it so hard to detect nearby contamination. To illustrate our point, we construct a ‘custom made’ ρ -function for dealing with nearby outliers. It is given by

$$\rho_a(d) = \begin{cases} d^2/2 & \text{if } 0 \leq d \leq \sqrt{p} \\ ((1+a)(2\sqrt{p}d - p) - d^2)/(2a) & \text{if } \sqrt{p} < d \leq (1+a)\sqrt{p} \\ p(1+a)/2 & \text{if } d > (1+a)\sqrt{p} \end{cases} \tag{2}$$

where $a \geq 0$ is a tuning constant that determines where the ρ -function becomes flat. Note that the dimension p matters, because for uncontaminated Gaussian data the




Fig. 2 The custom made ρ -function ρ_a (left) and the corresponding weight function w_a (right) for $p = 2$, for different values of a

squared distance roughly follows a χ^2 distribution with p degrees of freedom which has expectation p . The corresponding ψ -function is given by $\psi_a(d) = \rho'_a(d)$:

$$\psi_a(d) = \begin{cases} d & \text{if } 0 \leq d \leq \sqrt{p} \\ (\sqrt{p}(1+a) - d) / a & \text{if } \sqrt{p} < d \leq (1+a)\sqrt{p} \\ 0 & \text{if } d > (1+a)\sqrt{p} \end{cases} \quad (3)$$

and the weight function by $w_a(d) = \psi_a(d)/d$:

$$w_a(d) = \begin{cases} 1 & \text{if } 0 \leq d \leq \sqrt{p} \\ (\sqrt{p}(1+a) - d) / (ad) & \text{if } \sqrt{p} < d \leq (1+a)\sqrt{p} \\ 0 & \text{if } d > (1+a)\sqrt{p} \end{cases} \quad (4)$$

For this ρ -function, $K = E[\rho(\|Z\|_2)]$ can be computed numerically or by Monte Carlo.

Figure 2 shows the functions ρ_a and w_a for dimension $p = 2$, for different values of a . In the right panel we see that the more a approaches 0, the harder the rejection becomes.

The left panel of Fig. 3 shows the weights assigned to the observations in the geyser data, where the statistical distances on the horizontal axis were computed from the MCD estimates of location and scatter. The MCD method (red line) assigns weights that are one for most points of the main cluster and zero for all points of the second cluster. In contrast, the bisquare (black curve) still gives quite a bit of weight to the points lying in between the two clusters. (The weights are shown as little circles on the curve.) These points pull the S-estimator towards a non-robust solution. Our custom weight function (blue curve) redescends much faster so it gives small weights to the intermediate points, thereby behaving more like the MCD. The right panel shows the weights in the cows data, with similar conclusions.

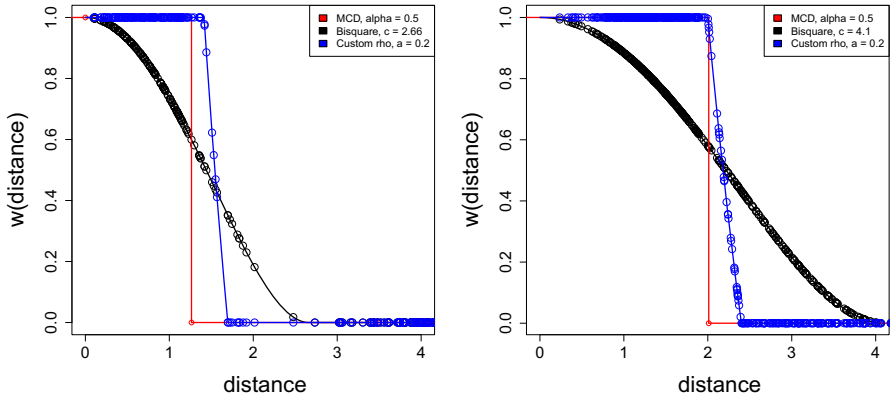


Fig. 3 Comparing the distance weights of the MCD estimator and the scatter S-estimators with bisquare ρ -function and custom ρ_a for the geyser data (left) and the cows data (right)

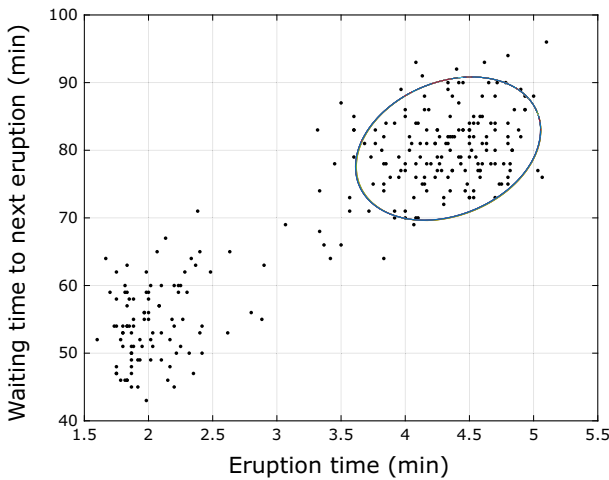


Fig. 4 Geysers data: fit obtained by the S-estimator with the custom ρ_a function

We applied the S-estimator with the custom function ρ_a for $a = 0.2$ to the geyser dataset, using the FSDA toolbox of Riani et al. (2012). This does yield a fit to the main cluster, as illustrated by the tolerance ellipse in Fig. 4. For the cows data we also obtain the desired fit.

We constructed this particular custom function ρ_a with the sole purpose of illustrating that the behavior of S-estimators does not only depend on tuning, but also on the shape of the ρ -function. The shape of this particular ρ_a makes the S-estimator use weights similar to those of the MCD, especially when a is tiny. This exemplifies the tradeoff between the ability to deal with nearby outliers and statistical efficiency at Gaussian data. But we definitely do not propose ρ_a for general use in multivariate S-estimators because it has some disadvantages. If we use $K = E[\rho(\|Z\|)]$ its breakdown value $\min(K/\max(\rho_a), 1 - K/\max(\rho_a))$ is below the 50% of the MCD

and decreases with the dimension p . We could also put $K = \max(\rho_a)/2$ but then we obtain an inconsistent estimator that needs a consistency correction. In either case we do not know how many data points will fall in the non-constant part of ρ_a . For the MCD this number is always h , and the MCD is also easier to compute.

4 Speeding up the computation

Monitoring requires the repeated calculation of a statistical procedure for various values of the parameter that needs tuning. Therefore the feasibility of monitoring depends on the computation time of the method being monitored. The Matlab code used in the paper computes the S-estimator, MM-estimator and MCD from 1000 initial elemental subsets that are kept unchanged throughout the monitoring. Computing concentration steps from 1000 subsets is computationally demanding. An alternative is to run the deterministic algorithm of Hubert et al. (2012) who start from six specifically constructed initial estimates instead of 1000 random ones, and illustrated their method by monitoring memberships in a multivariate classification. Hubert et al. (2015) extended this work to S-estimators and MM-estimators and illustrated it by monitoring the estimates of location and scatter in a flow cytometry dataset with $n = 29,486$ data points.

5 Limitations of monitoring

The review paper by CRAC clearly documents the benefits of monitoring. On the other hand the approach also has some limitations having to do with the size of the problem and the number of parameters that should be tuned. One of us works in the food sorting industry where thousands of items are inspected in milliseconds in order to detect outliers (impurities), yielding settings with tens of dimensions. In such situations monitoring could still be of use for tuning classification parameters, as long as it can be done off-line. The large sample size is the least problematic: instead of monitoring all the individual observations one can monitor a subset of them, or summary statistics. Higher dimensions are more difficult to handle due to the substantially increasing computation time. Another question is how many parameters can be tuned, as the parameters usually interact with each other. The paper illustrates tuning one parameter, and if there are more the experiment needs to be carefully designed.

References

- Hubert M, Rousseeuw P, Vakili K (2014) Shape bias of robust covariance estimators: an empirical study. *Stat Pap* 55:15–28
- Hubert M, Rousseeuw P, Vanden Branden K (2005) ROBPCA: a new approach to robust principal component analysis. *Technometrics* 47:64–79
- Hubert M, Rousseeuw P, Vanpaemel D, Verdonck T (2015) The DetS and DetMM estimators for multivariate location and scatter. *Comput Stat Data Anal* 81:64–75
- Hubert M, Rousseeuw PJ, Verdonck T (2012) A deterministic algorithm for robust location and scatter. *J Comput Graph Stat* 21(3):618–637
- Riani M, Perrotta D, Torti F (2012) FSDA: a MATLAB toolbox for robust analysis and interactive data exploration. *Chemometr Intell Lab Syst* 116:17–32