FINDING MULTIVARIATE OUTLIERS WITH FASTPCS

BY KAVEH VAKILI AND ERIC SCHMITT

The Projection Congruent Subset (PCS) is a new method for finding multivariate outliers. Like many other outlier detection procedures, PCS searches for a subset which minimizes a criterion. The difference is that the new criterion was designed to be insensitive to the outliers. PCS is supported by FastPCS, a fast and affine equivariant algorithm which is also detailed. Both an extensive simulation study and a real data application from the field of engineering show that FastPCS performs better than its competitors.

1. Introduction. Outliers are observations that do not follow the pattern of the majority of the data [6]. Outlier identification is a major part of data analysis for at least two reasons. First, because a few outliers, if left unchecked, will exert a disproportionate pull on estimated parameters, and we generally do not want our inferences to depend on such observations. In addition, we may want to find outliers to set them aside and study them as objects of interest in their own right. In any case, detecting outliers in settings involving more than two variables is difficult because we can not inspect the data visually and have to rely on algorithms instead.

Formally, this paper concerns itself with the simplest, most basic variant of the multivariate outlier detection problem. The general setting is that of a sample of n observations $x_i \in \mathbf{R}^p$ (with $n > p$), at least $h = \lfloor (n+p+1)/2 \rfloor$ of which are drawn from a multivariate elliptical distribution. The objective is to identify reliably the index of the remaining ones. A more complete treatment of this topic can be found in textbooks [2, for example].

In this article we introduce PCS, a new procedure for finding multivariate outliers. We also detail FastPCS, a fast algorithm for computing it. The main output of FastPCS is an outlyingness index measuring how much each observation departs from the pattern set by the majority of the data.

The PCS outlyingness index is affine equivariant (meaning that the outlyingness ranking of the observations is not affected by a linear transformation of the data) and can be computed efficiently for moderate values of p and large values of n .

To derive this index, FastPCS proceeds in two steps. First, it strives to select among many possible h -subsets of observations one devoid of outliers.

Keywords and phrases: Outlier detection, multivariate statistics, computational statistics

Then, the outlyingness index is simply the distance of each observation to this subset.

For easier outlier detection problems, we find that our approach produces results similar to state-of-the-art outlier detection algorithms. When considering more difficult cases however we find that the solution we propose leads to significantly better outcomes.

In the next section we motivate and define the PCS outlyingness and FastPCS. Then, in Section 3 we compare FastPCS to several competitors on synthetic data. In Section 4 we conduct a real data comparison. In Section 5, we offer closing thoughts and discuss directions for further research.

2. The PCS outlyingness index.

2.1. Motivation. Throughout this note, we will denote as H_m a subset of the indexes $\{1, 2, \dots, n\}$ (The subscript m indexing such sets will be used later on). For any H_m we denote its (sample) mean and covariance matrix as

$$(2.1) \quad (t_m, S_m) = \left(\text{ave}_{i \in H_m} x_i, \text{cov}_{i \in H_m} x_i \right)$$

and we will write the squared Mahalanobis distance of an observation x_i as

$$(2.2) \quad d_{MD,i}^2(t_m, S_m) = (x_i - t_m)' S_m^{-1} (x_i - t_m) .$$

FastPCS starts selects among many such subsets an h -subset of observations that minimizes a criterion. Other affine equivariant outlier detection algorithms that proceed in this way are the FastMVE [2], the FastMCD [8] and the SDE [9]. For all four procedures we will denote the chosen h -subset as H_* .

Each of these algorithms relies on a different criterion to select the optimal h -subset. In all cases, there is a straightforward relationship between H_* and the resulting outlyingness index. For example, for FastMVE and FastMCD the outlyingness index is the vector of distances $d_{MD,i}^2(t_*, S_*)$. For SDE the outlyingness index is the vector of values of $P_M(x_i)$

$$(2.3) \quad P_M(x_i) = \sup_{a_m \in B_{M_p}} \frac{|x_i' a_m - \text{med}_{j=1}^n(x_j' a_m)|}{\text{mad}_{j=1}^n(x_j' a_m)} ,$$

where B_{M_p} is a set of M_p directions orthogonal to hyperplanes spanned by a p -subset of $\{1 : n\}$ and H_* contains the observations with smallest values of $P_M(x_i)$.

In all cases, if H_* itself is contaminated, then the resulting outlyingness index can no longer be depended upon to find the outliers. For FastMVE

and FastMCD, H_* is the subset of h observations with smallest covariance matrix determinant and enclosed by the ellipsoid with smallest volume, respectively. Given enough outliers, an adversary can easily confound these two characterizations (for example, by placing outliers on a subspace of $\mathbf{R}^p$) and ensure that a contaminated h -subset will always be chosen over h -subsets formed of genuine observations.

It is not always recognized that the SDE outlyingness index is also sensitive to the outliers. Given enough outliers, they can always be placed in such a way that the denominator in Equation (2.3) will be smaller for some direction a_* along which the numerator is smaller for outliers than for genuine observations (for example, consider a direction a_* close to the principal axis of the lighter, orange ellipse in the lower-left panel in Figure 1). Then, for many observations the maximum outlyingness will be attained at $a_m = a_*$. Repeated over a large number of such directions, this will cause the value of the outlyingness index in Equation (2.3) to be smaller for many outliers than for genuine observations.

Consider the following example. The four panels in Figure 1 all depict the same 70 draws from a standard bivariate normal distribution together with 30 draws from a second normal distribution with smaller variance and centered at $(5, -1)$. Light, orange ellipses in the first three plots show the (t_*, S_*) obtained using the best subset found by the FastMCD, FastMVE and SDE. These were computed with the R package `rrcov` [11], using default settings; 500 starting p -subsets (for the first two), and 500 directions a_m (for SDE); and $h = 51$ (see Section 3.4) in all cases.

In all three cases, the fitted ellipses (drawn in solid orange lines), fail to adequately fit any h -subset of the data, including the subset they enclose. In particular, their centers of symmetry (orange stars) are not located in areas of highly concentrated observations. In all cases, the model fitted on H_* appears visually distinct from the distribution governing the good part of the data (drawn as a dashed, dark blue ellipse). This is confirmed by the biases (a dimensionless measure of dissimilarity between two subsets, see Section 3) of 7, 5 and 3.5 for the three algorithms.

The algorithm we propose differs from these methods in that it uses a new measure of multivariate congruence (which we detail in the next section) to select the optimal h -subset. The main advantage of this new characterization lies in its insensitivity to the outliers. As we argue below, this makes the outlyingness index derived from FastPCS both quantitatively and qualitatively more reliable.

FIG 1. The four plots depict (left right, top, bottom) the same configuration of observations. The darker blue dashed ellipses show the contours of the model governing the distribution of the majority –here 70 out of 100– of the observations. The light orange, solid, ellipses show, respectively, the FastMCD, FastMVE, SDE and FastPCS estimates of scatter and location.

2.2. Construction of the PCS Outlyingness index. PCS looks for the h -subset of multivariate observations that is most *congruent* along many univariate projections. In this context, we measure the congruence of a given h -subset along a given projection by the size of its overlap with a second subset that is optimal (in a sense we make precise below) on that projection. The PCS criterion is based on the observation that a spatially cohesive h -subset will tend to be congruent with these optimal h -subsets, and a spatially disjoint one will not.

More precisely, denoting by A_{mk} a $p \times p$ matrix formed of p observations (we detail below how we pick these p observations), a_{mk} the hyperplane $\{a_{mk} : A_{mk}a_{mk} = 1_p\}$ and $d_{P,i}$ the (squared) orthogonal distance of x_i to a_{mk} :

$$(2.4) \quad d_{P,i}^2(a_{mk}) = (x_i' a_{mk} - 1)^2 / \|a_{mk}\|^2.$$

The set of h observations with smallest $d_{P,i}^2(a_{mk})$ will be denoted as H_{mk} . For a given subset H_m and direction a_{mk} we define the *incongruence index*

of H_m along a_{mk} as

$$(2.5) \quad I(H_m, a_{mk}) := \log \operatorname{ave}_{i \in H_m} d_{P,i}^2(a_{mk}) - \log \operatorname{ave}_{i \in H_{mk}} d_{P,i}^2(a_{mk}) .$$

This index is always positive and will have small value if the projection of the members of H_m along a_{mk} greatly overlaps with the members of H_{mk} . To remove the dependence of Equation (2.5) on a_{mk} we measure the incongruence of H_m by considering the average over many directions:

$$(2.6) \quad I(H_m) := \operatorname{ave}_{a_{mk} \in B(H_m)} I(H_m, a_{mk}) ,$$

where $B(H_m)$ are all directions orthogonal to a hyperplane spanned by a p -subset of H_m . We call the H_m with smallest $I(H_m)$ the *projection congruent subset*. In essence, the I index measures the spatial cohesion of an h -subset H_m in terms of how much the projections of its members overlap with those of the H_{mk} over many projections. In practice, it would be too laborious to evaluate Equation (2.6) over all members of $B(H_m)$. A practical solution is to take the average over a sample of K random directions $\tilde{B}_K(H_m)$ instead.

The I index of a spatially disjoint h -subset tends to be higher than that of a spatially cohesive h -subset. This is because when H_m forms a spatially cohesive set of observations, $\#\{H_m \cup H_{mk}\}$ tends to be larger.

This is illustrated on an example in Figure 2. Both panels depict the same set of $n = 100$ points. These points form two disjoint groups of observations. The main group contains 70 points and is located on the left-hand side. Each panel illustrates the behavior of the I index for a given h -subset of observations. H_1 (left) forms a set of spatially cohesive observations, all belonging to the main group. H_2 , in contrast, forms a spatially disjoint set of observations and contains members drawn from both groups.

For each H_m -subset, $m = \{1, 2\}$, we drew two hyperplanes a_{m1} (dark blue, dashed) and a_{m2} (light orange). The dark blue dots show the members of $\{H_m \cup H_{m1}\}$. Similarly, light orange dots show the members of $\{H_m \cup H_{m2}\}$. The diamonds (black squares) show the members of H_{m1} (H_{m2}) that do not belong to H_m . After just two projections, the number of non-overlapping observations (i.e. $\{H_m \setminus H_{m1}\} \cap \{H_m \setminus H_{m2}\}$) is 8 ($m = 1$) and 18 ($m = 2$) respectively. As we increase the number of directions a_{mk} , this pattern repeats and the difference between a spatially cohesive and a disjoint subset grows steadily.

The I index measures the size of this overlap. For a direction a_{mk} , the members of H_{m1} and H_{m2} not in H_m (shown as diamonds and black squares in Figure 2) will decrease the denominator in Equation (2.5) without affecting the numerator, increasing the overall ratio. Consequently, h -subsets formed

FIG 2. *Incongruence index for a subset H_1 of spatially cohesive observations (left) and a subset H_2 of spatially disjoint observation (right)*

of spatially disjoint groups of observations will have larger values of the I index. Crucially, the I index characterizes a cohesive h -subset of observations independently of the spatial configuration of the outliers. For example, the pattern shown in Figure 2 would still hold if the cluster of outliers were more concentrated. This is also illustrated in the fourth quadrant of Figure 1 where the optimal h -subset found by FastPCS is not unduly attracted by members of the smaller cluster of observations located on the right.

In Sections 3 and 4, we show that this new characterization allows FastPCS to reliably select uncontaminated h -subsets. This includes many situations where other algorithms fail to do so. First though, the following section details the FastPCS algorithm.

2.3. A Fast Algorithm for the PCS Outlyingness. To compute the PCS outlyingness index, we propose the FastPCS algorithm (2.7). It combines ideas from FastMCD (the use of random p -subsets as starting points and an iterative concentration step) with some new ones. Throughout, M_p denotes the number of starting p -subsets.

An important characteristic of our algorithm is that it can detect exact-fit situations: when h or more observations lie exactly on a subspace, FastPCS will return the indexes of an h -subset of those observations and the fit given by the observations in H_* will coincide with the subspace. Since FastPCS is affine equivariant, this behavior suggests that it has maximal finite sample breakdown point, as discussed in [5, remark 1, pg 123] and [7, 2]. Intuitively, this is because in any affine equivariant metric the $n - h$ observations not in

the subspace do in effect lie arbitrarily far away from the ones that are.

Step *b* increases the size of H_m from $p+1$ to its final size $h = \lfloor (n+p+1)/2 \rfloor$ in L steps, rather than in one as is done in FastMCD. This improves the robustness of the algorithm when outliers are close to the good data. We find that increasing L does not improve performance much if L is greater than 3 and use $L = 3$ as default.

(2.7) Algorithm FastPCS

for $m = 1$ to M_p do:

a: $H_m^0 \leftarrow \{p+1 \text{ observations drawn at random from } 1:n\}$

b: for $l = 1$ to L do:

$$D_i(H_m^l) \leftarrow \frac{\text{ave}_{k=1}^K d_{P,i}^2(a_{mk})}{\text{ave}_{j \in H_m^l} d_{P,j}^2(a_{mk})} \quad \text{for all } i = 1, \dots, n$$

set $q \leftarrow \lfloor (n-p-1)l/(2L) \rfloor + p+1$

set $H_m^l \leftarrow \{i : D_i(H_m^l) \leq D_{(q)}(H_m^l)\}$ ('concentration step')

end for

$H_m \leftarrow H_m^L$

c: compute $I(H_m) \leftarrow \text{ave}_{k=1}^K I(H_m, a_{mk})$

end for

Keep H_* , the subset H_m with lowest $I(H_m)$. The final outlyingness index of each observation is given by $d_{MD,i}(t_*, S_*)$.

Empirically also, we found that small values for K are sufficient to achieve good results and that we do not gain much by increasing K above 25, so we set $K = 25$ as the default (this is the value we use throughout). That such a small number of random directions suffice to reliably identify the outliers is remarkable. This is because FastPCS uses directions generated by p -subsets of H_m .

Compare this to SDE algorithm which needs a much larger number of projections to reliably find the outliers. This is because in SDE the data is projected along directions given by hyperplanes passing through p points drawn indiscriminately from the entire set of observations. Consequently, when the contamination rate is high, most of these p -subsets will be contaminated, yielding directions that can end up almost parallel to each other. In contrast, our choice always ensures a wider spread of directions when H_m is uncontaminated and thus yields better results.

Like FastMCD and FastMVE, FastPCS uses many random p -subsets as starting points. The number of initial p -subsets, M_p , must be large enough to ensure that at least one of them is uncontaminated. For FastMCD and

FastMVE, for each starting p -subset, the computational complexity scales as $O(p^3 + np^2)$ whereas for FastPCS it is $O(p^3 + np)$. Computing the SDE outlyingness index also costs $O(p^3 + np)$ for each hyperplane and here also the number of such hyperplanes must be of order M_p to ensure that at least one of them does not pass through any of the outliers.

The value of M_p (and therefore the computational complexity of all four algorithms) grows exponentially with p . The actual run times will depend on implementation choices but in our experience FastMCD is slightly faster than FastPCS and both are noticeably slower than either FastMVE or SDE. In practice this means all four procedures become impractical for values of p much larger than 25. Furthermore, all four procedures belong to the class of so called ‘embarrassingly parallel’ algorithms, i.e. their time complexity scales as the inverse of the number of processors meaning that they are particularly well suited to benefit from modern computing environments. To enhance user experience, we implemented FastPCS in C++ code wrapped in a portable R package [3] distributed through CRAN (package `FastPCS`).

3. Empirical Comparison: Simulation Study. In this section we evaluate FastPCS numerically and contrast its performance to that of the SDE, FastMCD and FastMVE algorithms. The last three were computed using the R package `rrcov` with default settings (except, respectively, for the number of random directions and starting subsets which for all algorithms we set according to Equation (3.5)). Each algorithm returns a subset H_* of h observations with the smallest outlyingness index. Our evaluation criteria are the bias and the misclassification rate of these h observations. Below, we briefly describe these.

3.1. Asymptotic Bias. Given a central model $\mathcal{F}_u$ and an arbitrary contaminating distribution $\mathcal{F}_c$, consider the ε -contaminated model

$$(3.1) \quad \mathcal{F}_\varepsilon = (1 - \varepsilon)\mathcal{F}_u(\mu_u, \Sigma_u) + \varepsilon\mathcal{F}_c(\mu_c, \Sigma_c) .$$

For a subset of observations H_* , the (asymptotic) bias measures the deviation of (t_*, S_*) from the true (μ_u, Σ_u) . Formally, denoting $G = |S_*|^{-1/p}S_*$ and $\Gamma = |\Sigma_u|^{-1/p}\Sigma_u$, we have that for an affine equivariant scatter matrix S , all the information about the bias is contained in the matrix $\Gamma^{-1/2}G\Gamma^{-1/2}$ or equivalently its condition number [13]:

$$(3.2) \quad \text{bias}(S) = \log \lambda_1(\Gamma^{-1/2}G\Gamma^{-1/2}) - \log \lambda_p(\Gamma^{-1/2}G\Gamma^{-1/2}) ,$$

where λ_1 (λ_p) are the largest (smallest) eigenvalues of $\Gamma^{-1/2}G\Gamma^{-1/2}$. Evaluating the bias of (t_*, S_*) is an empirical matter. For a given sample, the bias will depend on the dimensionality of the data, the rate of contamina-

tion and the distance separating the outliers from the good part of the data. Finally, the bias also depends on the spatial configuration of the outliers (the choice of $\mathcal{F}_c$). Fortunately, for affine equivariant algorithms the outlier configurations causing the largest biases are known and so we can focus on these cases.

3.2. Misclassification rate. We can also compare the outlyingness indexes in terms of rate of contamination of their final h -subsets. Denoting by I_c the index set of the contaminated observations and $\mathcal{I}(\bullet)$ the indicator function, the misclassification rate is:

$$(3.3) \quad \text{Mis.Rate}(I_c, H) = \frac{\sum_{i \in H} \mathcal{I}(i \in I_c)}{\sum_{i=1}^n \mathcal{I}(i \in I_c)}.$$

This measure is always in $[0, 1]$, thus yielding results that are easier to compare across dimensions and rates of contamination. A value of 1 means that the H contains all the outliers. The main difference with the bias criterion is that the misclassification rate does not account for how disruptive the outliers are. For example, when the distance separating the outliers from the good part of the data is small, it is possible for $\text{Mis.Rate}(I_c, H_*)$ to be large without a commensurate increase in $\text{bias}(S_*)$.

3.3. Outlier configurations. We generate many contaminated datasets X_ε of size n with $X_\varepsilon = X_u \cup X_c$ where X_u and X_c are, respectively, the genuine and outlying part of the sample. The bias depends on the distance between the outliers and the genuine observations which we will measure by

$$(3.4) \quad \nu = \min_{i \in I_c} \sqrt{d_{MD,i}^2(t_u, S_u) / \chi_{0.99,p}^2}.$$

The bias also depends on the spatial configuration of X_c . For affine equivariant algorithms, the worst-case configurations (those causing the largest bias) are known. In increasing order of difficulty these are:

- **Shift configuration.** If we constrain the adversary to (a) $|\Sigma_c| \geq |\Sigma_u|$ and (b) place X_c at a distance ν of X_u , then, to maximize the bias, the adversary will set $\Sigma_c = \Sigma_u$ (Theorem 1 in [4]) and set μ_c in order to satisfy (b). Intuitively, this makes the components of the mixture the least distinguishable from one another.
- **Point-mass configuration.** If we omit the constraint (a) above but keep (b), the adversary will place X_c in a single point so $|\Sigma_c| = 0$ (Theorem 2 in [4]).
- If we omit both constraints (a) and (b), the adversary may set $\mu_c = \mu_u$ and choose Σ_c to obtain a large bias. The Barrow wheel contamination

[10] does this.

FIG 3. The three outlier configurations (Shift, Point-mass and Barrow wheel). The outliers are depicted as light orange dots.

We also considered other configurations such as radial outliers as well as cases where ν was set to extremely large values (i.e. $\nu \geq 10000$), but they posed little challenge for any of the algorithms, so these results are not shown.

3.4. Simulation parameters. We can generate both the uncontaminated data X_u and the contaminated data X_c from the standard normal distribution since all methods under consideration are affine equivariant. For the shift and point configurations, we will also generate data from the standard Cauchy distribution in order to quantify the sensitivity of each method to the tail index of the data.

For the shift and point configurations we generated the outliers by setting Σ_c as either D_p or $10^{-4}D_p$ (D_p denotes a rank p diagonal matrix) and set μ_c so that Equation (3.4) is satisfied. We generated the Barrow wheel configuration using the **robustX** package [10] with default parameters. The three configuration types are depicted in Figure 3 for $n = 100$, $p = 2$, $\varepsilon = 0.4$, $\nu = 2$ and $\mathcal{F}_u$ is the bivariate normal. The outlying observations are the lighter orange circles. The complete list of simulation parameters follows:

- the dimension p is one of $\{4, 8, 12, 16\}$,
- the sample size is $n = 25p$,
- the contamination fraction ε is one of $\{0.1, 0.2, 0.3, 0.4\}$,
- the configuration of the outliers is either shift, point, or Barrow wheel,
- for shift and point contamination, the distance ν comes from the uniform distribution on $(0, 10)$. The Barrow wheel contamination does not depend on ν .

- $\alpha \in (0.5, 1)$ is a parameter determining the size of the subset of observations assumed to follow a model (with $h \approx \lceil \alpha n \rceil$). In Section 3.5 (Section 3.6) we consider the case where we set $\alpha = 0.5$ ($\alpha = 0.75$).
- the number of initial $(p + 1)$ -subsets M_p is given by [9]

$$(3.5) \quad M_p = \frac{\log(0.01)}{\log(1 - (1 - \varepsilon_0)^{p+1})},$$

with $\varepsilon_0 = 4(1 - \alpha)/5$ so that the probability of getting at least one uncontaminated subset is always at least 99 percent.

In Figures 4 to 10 we display the bias and the misclassification rate (left and right plots, respectively) for discrete combinations of the dimension p , and contamination rate ε . In all cases, we expect the outlier detection problem to become monotonically harder as we increase p and ε , so not much information will be lost by considering a discrete grid of a few values for these parameters. For the Barrow wheel, p and ε are the only parameters we have and so we can chart the results as trellises of boxplots [1].

For the shift and point contamination models, the configurations also depend on the distance separating the data from the outliers. Here, the effects of ν on the bias are harder to foresee: clearly nearby outliers will be harder to detect but misclassifying distant outliers will increase the bias more. Therefore, we will test the algorithms for many values (and chart the results as a function) of ν . For both the bias and the misclassification rates curves, for each algorithm, a solid colored line will depict the median and a dotted line (of the same color) the 75th percentile. Here, each panel will be based on 1000 simulations.

3.5. Simulation results (a). The first part of the simulation study covers the case where there is no information about the extent to which the data is contaminated. Then, we have to set the size of the subset of observations having positive weight to $h = \lfloor (n + p + 1)/2 \rfloor$, corresponding to the lower bound of slightly more than the majority of the observations.

In Figure 4 we display the bias ($\text{bias}(S_*)$) and the misclassification rate ($\text{Mis.Rate}(I_c, H_*)$) curves of each algorithm as a function of ν for different values of p and ε for the shift configuration when $\mathcal{F}_u$ is the standard normal. All the methods perform well until $\varepsilon = 0.3$ and $p \leq 8$. For larger values of p and ν , the bias curves and misclassification rate of FastMVE and SDE clearly stand out for values of ν between 2 and 6. Eventually, as we move the outliers further away, SDE and FastMVE can again find them reliably.

When $\mathcal{F}_u$ is Cauchy (Figure 5), FastMVE performs significantly worse than the other algorithms starting already at $\varepsilon = 0.2$ and $p \geq 12$. Eventually,

FIG 4. Bias (left) and misclassification rate (right) due to shift contamination for $\varepsilon = \{0.1, \dots, 0.4\}$, $p = \{4, \dots, 16\}$, $h = \lfloor (n + p + 1)/2 \rfloor$ and Normal-distributed observations shown as a function of ν . FastMCD (dotdash lines), FastMVE (longdash lines), SDE (twodash lines), FastPCS (solid lines).

($\varepsilon = 0.3$ and $p \geq 12$) FastMCD also breaks away from Σ_u . SDE performs better until $\varepsilon = 0.4$, where it is also noticeably affected by the thicker tails of the Cauchy distribution and fails to identify the outliers. FastPCS maintains constant and low bias (and misclassification rates) throughout.

FIG 5. Bias (left) and misclassification rate (right) due to shift contamination for $\varepsilon = \{0.1, \dots, 0.4\}$, $p = \{4, \dots, 16\}$, $h = \lfloor (n + p + 1)/2 \rfloor$ and Cauchy-distributed observations shown as a function of ν . FastMCD (dotdash lines), FastMVE (longdash lines), SDE (twodash lines), FastPCS (solid lines).

In Figure 6, we show the results for the more difficult case of point-mass contamination when $\mathcal{F}_u$ is the standard normal. Starting at $\varepsilon = 0.2$, the bias

curves and misclassification rate of FastMCD become very high (except for $p = 4$). Starting at $(\varepsilon = 0.3$ and $p \geq 8$) all the algorithms except FastPCS fail to reliably find the outliers: looking at the misclassification rate, the optimal h -subsets for SDE, FastMVE and FastMCD even contain a higher contamination rate than ε .

The Cauchy case shown in Figure 7 is also interesting. It suggests, again, that FastMCD and FastMVE are very sensitive to fat tails in the distribution of the good part of the data. The behavior of SDE is roughly in line with Figure 6. Again, we see that FastPCS is the only algorithm that can reliably find the outliers. Furthermore, the misclassification rates reveal that the H_* found by SDE, FastMCD and FastMVE often contains a proportion of outliers higher than ε .

FIG 6. Bias (left) and misclassification rate (right) due to point-mass contamination for $\varepsilon = \{0.1, \dots, 0.4\}$, $p = \{4, \dots, 16\}$, $h = \lfloor (n+p+1)/2 \rfloor$ and Normal-distributed observations shown as a function of ν . *FastMCD* (dotted lines), *FastMVE* (longdash lines), *SDE* (twodash lines), *FastPCS* (solid lines).

In Figure 8 we show the bias curves for the Barrow wheel as boxplots. FastMCD performs badly compared to the other methods, and FastMVE deteriorates for ε larger than 30 percent. It is at $\varepsilon = 0.4$ that FastPCS demonstrably exhibits superior performance, with SDE producing results in between the other three algorithms.

Across the tests, FastPCS consistently maintains low and stable biases and misclassification rates. Furthermore, the performance of FastPCS is relatively insensitive to whether ν , ε and p are low or high. Additionally, in contrast with the other algorithms, FastPCS is unaffected by the thickness of the tails of $\mathcal{F}_u$. The consistent ability of FastPCS to detect outliers is a highly desirable feature, which we characterize as a form of qualitative robustness.

FIG 7. Bias (left) and misclassification rate (right) due to point-mass contamination for $\varepsilon = \{0.1, \dots, 0.4\}$, $p = \{4, \dots, 16\}$, $h = \lfloor (n+p+1)/2 \rfloor$ and Cauchy-distributed observations shown as a function of ν . *FastMCD* (dotdash lines), *FastMVE* (longdash lines), *SDE* (twodash lines), *FastPCS* (solid lines).

FIG 8. Bias (left) and misclassification rate (right) due to the Barrow wheel contamination for $\varepsilon = \{0.1, \dots, 0.4\}$ and $p = \{4, \dots, 16\}$, $h = \lfloor (n+p+1)/2 \rfloor$. *FastMCD*, *FastMVE*, *SDE*, *FastPCS* (far right).

3.6. *Simulation results (b).* In this section, we consider the case, important in practice, where the user can confidently place an upper bound on the rate of contamination of the sample. To fix ideas, we will set h , the number of observations assumed to follow a model, to $h \approx 3n/4$. For *FastPCS*, *FastMVE* and *FastMCD* we adapt the algorithms by setting their respective α parameter to 0.75. *SDE* does not have a corresponding parameter, so we take as member of H_* the (approximately) $3n/4$ observations with smallest

outlyingness. Furthermore, we reduce the number of starting subsets (FastPCS, FastMCD, FastMVE) and random directions (SDE), by using Equation (3.5), but this time setting $\varepsilon_0 = 0.2$. Then, as before, we measure the effect of our various configurations of outliers on the estimators (but this time only considering values of $\varepsilon \in \{0.1, 0.2\}$).

In Figures 9 and 10, we show the simulation results for shift and point-mass contamination. The results for Normal-distributed observations are shown in the first two rows, while the last two rows show the results for observations drawn from the multivariate Cauchy distribution. The shift contamination results in Figure 9 show that when we increase the size of the active subsets, FastPCS still maintains a low bias and misclassification rate, but at equivalent rates of contamination, the other algorithms exhibit noticeably weaker performance than in Figures 4 and 5. Results in Figure 10 illustrate that under point-mass contamination, FastPCS again reports good results. The other algorithms continue to exhibit larger biases and misclassification rates (at equivalent rates of contamination) than in Figures 6 and 7. We do not show the results for the Barrow wheel because all the algorithms perform equivalently.

FIG 9. Bias (left) and misclassification rate (right) due to the shift contamination for $h \approx 3n/4$, $\varepsilon = \{0.1, 0.2\}$ and $p = \{4, \dots, 16\}$. The first (last) two rows are for multivariate normal (Cauchy) distributed r.v. *FastMCD* (dotdash lines), *FastMVE* (longdash lines), *SDE* (twodash lines), *FastPCS* (solid lines).

3.7. Small Sample Accuracy. For affine equivariant outlier detection algorithms there is a trade-off between bias and accuracy [7]. This is especially problematic in small n and low ε settings because then any reduction in bias can be partially undone by losses in small sample accuracy. For

FIG 10. Bias (left) and misclassification rate (right) due to point-mass contamination for $h \approx 3n/4$, $\varepsilon = \{0.1, 0.2\}$ and $p = \{4, \dots, 16\}$. The first (last) two rows are for multivariate normal (Cauchy) distributed r.v. *FastMCD* (dotted lines), *FastMVE* (longdash lines), *SDE* (twodash lines), *FastPCS* (solid lines). The online version of this figure is in color.

FastPCS in particular, this is visible in the first rows of Figures 4 and 8: in the misclassification plots, we see that for all algorithms the optimal subsets H_* are uncontaminated by outliers. Yet, the corresponding bias curves for FastPCS tend to be somewhat higher than for the other algorithms.

To improve accuracy in small samples without conceding too much in bias, a solution is to add a so-called re-weighting step to the outlier detection algorithm. In essence, we replace H_* by a larger subset J_+ , itself derived from H_* . The motivation is that, typically, J_+ will include a greater share of the uncontaminated observations. Over the years, many such re-weighting procedures have been proposed [2]. Perhaps the simplest is the so called hard-thresholding re-weighting. Given an optimal subset H_* , the members of J_+ are:

$$(3.6) \quad J_+ = \left\{ i : d_{MD,i}^2(t_*, S_*) \leq \chi_{0.975,p}^2 \left(\frac{\text{med}_{i=1}^n d_{MD,i}^2(t_*, S_*)}{\chi_{0.5,p}^2} \right) \right\},$$

where the ratio on the rhs of the inequality is a scale factor we use to make S_* consistent at the normal distribution.

Contrary to H_* , the size of J_+ is not fixed in advance. This makes it difficult to compare algorithms on the basis of $\text{bias}(S_+)$ or $\text{Mis.Rate}(I_c, J_+)$ evaluated at contaminated sub-samples X_ε . Nonetheless, we can still compare them in terms of the bias of S_+ computed at uncontaminated data sets X_u and for various sample sizes n . In essence, we measure the accuracy of the various algorithms by considering their biases at uncontaminated

datasets as a function of n . In Figure 11, we show the results of doing this for $n = \{100, 200, \dots, 500, 599\}$ (shown as the abscissa) and $p = 8$. We set 599 observations as an upper limit because when $n \geq 600$ FastMCD uses a nested sub-sampling scheme whereby larger datasets are divided into non-overlapping sub-samples of at most 600 observations.

FIG 11. Empirical bias at uncontaminated samples for $p = 8$ as a function of sample size. *FastMCD* (dotdash lines), *FastMVE* (longdash lines), *SDE* (twodash lines), *FastPCS* (solid lines). The online version of this figure is in color.

Figure 11 depicts the median (solid) and 75th percentile (dotted) accuracy curves of the algorithms for increasing n when $\mathcal{F}_u$ is normally distributed (left) and Cauchy distributed (right). Each panel is based on 1,000 experiments. As expected, FastPCS is noticeably less accurate than the other algorithms in the normal case. Nonetheless, we maintain that this difference is small compared to the reduction in bias achieved by FastPCS on contaminated samples. Furthermore, this gap in performance depends on the distribution of the good data: when $\mathcal{F}_u$ is Cauchy the accuracy of FastPCS is similar to that of the other algorithms.

4. Empirical Comparison: Case Study. In this section, we illustrate the behavior of FastPCS on a real data problem from the field of engineering: the Concrete Slump Test Data set [12]. This dataset includes 103 data-points, each corresponding to a different type of concrete. For each observation we have 7 so-called input variables measuring, respectively, the quantity of cement (kg/m³), fly ash (kg/m³), blast furnace slag (kg/m³), water (kg/m³), super-plasticizer (kg/m³), coarse aggregate (kg/m³), and fine aggregate (kg/m³) used to make the corresponding variety of concrete. Additionally, for each observation, we have 3 so-called output variables mea-

suring attributes of the resulting concrete. These are slump (cm), flow (cm) and 28-day compressive strength (Mpa).

This dataset actually contains two groups of observations collected over separate periods. The first 78 measurements pre-date the last 25 by several years. An interesting aspect of this dataset is that these two groups largely overlap on bivariate scatter-plots of the data, forming a seemingly homogeneous group of observations. Appearances can be deceptive however: when considered along all variables jointly, the two groups are clearly distinct. For example, denoting J_O (J_N) the subset of 78 (25) members of the early (later) batch of measurements, we find that the closest member of J_N lies at a squared Mahalanobis distance of over 760 with respect to (t_O, S_O) . As a point of comparison, this is nearly 32 times larger than $\chi_{0.99,10}^2$.

We first ran the SDE, FastMCD, FastMVE and FastPCS algorithms on the raw dataset (for the first three, we used the `rrcov` implementations with default parameters). We set the number of random directions in SDE and the number of random p -subsets in FastMCD, FastMVE and FastPCS to $M_p = 2000$, as given in Equation (3.5). To have comparable results, we will use for each estimator $d_{MD,i}(t_*, S_*)$, the vector of statistical distances wrt H_* .

FIG 12. Mahalanobis outlyingness index $d_{MD,i}(t_*, S_*)$ for the four estimators and the four variants of the concrete dataset. In each panel, the dark blue (light orange) boxplot depicts $d_{MD,i}(t_*, S_*)$ for the members of J_O (J_N). The online version of this figure is in color.

The Mahalanobis outlyingness indexes are displayed in the first column

of Figure 12, Concrete (i). Each row corresponds to an estimator. For each estimator we drew two boxplots. The first one (dark blue) depicts the outlyingness index for the 78 members of J_O and the second (light orange) boxplot for those 25 members of J_N . All the algorithms are able to unambiguously separate the two subgroups of observations: the outlyingness values assigned to the members of the more recent batch of observations are notably larger than the outlyingness values assigned to any members of J_O .

In a second experiment, Concrete(ii), we made the original outlier detection problem harder by narrowing the distance separating the outliers from the genuine observations. To do so, we pulled the 25 outliers towards the mean of the good part of the data by replacing the original outliers with 25 observations of the form $x_{n+i}^{ii} \leftarrow (x_{n+i} + t_O)/2$ for $1 \leq i \leq 25$, $n = 78$. We denote J_N^{ii} to be these 25 members of x_{n+i} for $1 \leq i \leq 25$. The second column of Figure 12 depicts the outlyingness indexes for this second experiment. Again, for each algorithm, the first (dark blue) boxplots depict the values of the outlyingness indexes for the 78 members of J_O and the second (lighter, orange) boxplot for the 25 members of J_N^{ii} . Note that the members of J_N^{ii} still form a cluster that is genuinely separate from the main group of observations. For example, the closest member of J_N^{ii} lies at a squared Mahalanobis distance of over 190 with respect to (t_O, S_O) . For comparison, this is over 8 times larger than $\chi_{0.99,10}^2$. Here, the FastPCS and SDE outlyingness indexes continue to clearly distinguish between the two groups. The outlyingness indexes derived from FastMVE and FastMCD however fail to adequately flag the outliers.

In a third experiment, Concrete(iii), we made the original outlier detection problem harder still by increasing the contamination rate of the sample. This was done by adding an additional 25 outliers of the form $x_{n+25+j}^{iii} \leftarrow (x_{n+1} + x_{n+1+j})/2$ for $1 \leq j \leq 25$, $n = 78$. We denote J_N^{iii} the members of this third set of 50 outliers. The third column of Figure 12 depicts the outlyingness indexes for this experiment. Again for each algorithm, a (dark) blue boxplot pertains to the 78 members of J_O and a lighter orange one to the 50 members of J_N^{iii} . For SDE, FastMVE and FastMCD we can see that increasing the contamination rate of the sample causes the outlyingness indexes of the two groups to overlap while the outlyingness index produced by FastPCS continues to make a clear distinction. Remarkably, the subsets of h observations with the smallest outlyingness chosen by the first three algorithms are primarily composed of outliers.

In a final experiment, Concrete(iv), we narrowed the distance between the good data and the outliers (as in Concrete(ii)), and increased the contamination rate (as in Concrete(iii)). The fourth column of Figure 12 depicts

the outlyingness indexes for this experiment, again with for each algorithm a (darker) blue boxplot for the 78 members of J_O and an lighter, orange one for the 50 members of J_N^{iv} . Here again, the outlyingness index of SDE, FastMCD and FastMVE fails to adequately flag the outliers. However, as with each of the previous experiments, FastPCS still assigns a larger index of outlyingness to the members of the outlying group.

Overall, we see that the results of the real data experiment confirm those of the simulations, at least qualitatively. When the contamination rates are small and the outliers well separated from the good part of the data, all outlier detection methods seem to perform equally well. As we consider more difficult outlier settings however, we see that the FastPCS outlyingness index is the only one that identifies the outliers reliably.

5. Outlook. In this article we introduced PCS, a new outlyingness index and FastPCS, a fast and affine equivariant algorithm for computing it. Like many other outlier detection algorithms, the performance of FastPCS hinges crucially on correctly identifying an h -subset of uncontaminated observations. Our main contribution is to characterize this h -subset using a new measure of the congruence of a multivariate cloud of points. The main feature of this new characterization is that it was designed to be insensitive to the configuration of the outliers.

In our simulations, we have focused on configurations of outliers that are worst-case for affine-equivariant algorithms, and found that FastPCS behaves notably better than the other procedures we considered, often revealing outliers that would not have been identified by the other approaches. In practice, admittedly, contamination patterns will not always be as difficult as those we considered above and in many cases the different methods will, hopefully, concur. Nevertheless, given that in practice we do not know the configuration of the outliers, as data analysts, we prefer to carry our inferences while planing for the worst contingencies.

Also through simulations, we found that the performance of FastPCS is not affected much by the tail index of the majority of the data. For example, FastPCS is capable of finding the outliers both when the majority of the data is distributed multivariate normal, or is drawn from a heavy-tailed distribution, such as the multivariate Cauchy.

In this article we emphasized the practical aspects of PCS. A number of theoretical properties deserve further investigation. In particular, we suspect that PCS has maximum breakdown point. Arguments supporting this conjecture are that the maximum biases are low and that the procedure has the exact fit property.

References.

- [1] Deepayan, S. (2008). *Lattice: Multivariate Data Visualization with R*. Springer, New York.
- [2] Maronna R. A., Martin R. D. and Yohai V. J. (2006). *Robust Statistics: Theory and Methods*. Wiley, New York.
- [3] R Core Team (2012). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria.
- [4] Rocke D. M. and Woodruff D. L. (1996). Identification of Outliers in Multivariate Data. *Journal of the American Statistical Association*, 91, 1047–1061.
- [5] Rousseeuw, P.J. and Leroy, A.M. (1987). *Robust Regression and Outlier Detection*. Wiley, New York.
- [6] Rousseeuw, P. J. and van Zomeren B.C. (1990). Unmasking Multivariate Outliers and Leverage Points. *Journal of the American Statistical Association* , Vol. 85, No. 411, pp. 633-639.
- [7] Rousseeuw, P. J. (1994). Unconventional features of positive-breakdown estimators. *Statistics and Probability Letters*, 19, 417–431.
- [8] Rousseeuw P. J. and Van Driessen K. (1999). A Fast Algorithm for the Minimum Covariance Determinant Estimator. *Technometrics*, 41, 212–223.
- [9] Stahel W. (1981). Breakdown of Covariance Estimators. Research Report 31, Fachgrupp für Statistik, E.T.H. Zürich.
- [10] Stahel W. and Maechler M. (2009). *robustX: eXperimental eXtraneous eXtraordinary ... Functionality for Robust Statistics*. R package version 1.1-2.
- [11] Todorov V. and Filzmoser P. (2009). An Object-Oriented Framework for Robust Multivariate Analysis. *Journal of Statistical Software*, 32, 1–47.
- [12] Yeh, I. (2007). Modeling slump flow of concrete using second-order regressions and artificial neural networks. *Cement and Concrete Composites*, Vol.29, No. 6, 474–480.
- [13] Yohai, V.J. and Maronna, R.A. (1990). The Maximum Bias of Robust Covariances. *Communications in Statistics–Theory and Methods*, 19, 2925–2933.

E-MAIL: kaveh.vakili@wis.kuleuven.be
eric.schmitt@wis.kuleuven.be