Fast Computation of Robust Subspace Estimators

Holger Cevallos-Valdiviezo^{a,b}, Stefan Van Aelst^{c,*}

^a*ESPOL Polytechnic University, Escuela Superior Politécnica del Litoral, ESPOL, (Facultad de Ciencias Naturales y Matemáticas, FCNM), Campus Gustavo Galindo Km. 30.5 Va Perimetral, P.O. Box 09-01-5863, Guayaquil, Ecuador*

^b*Ghent University, Department of Applied Mathematics, Computer Science and Statistics, Krijgslaan 281 S9, 9000 Gent, Belgium*

^c*KU Leuven, Department of Mathematics, Celestijnenlaan 200B, 3001 Leuven, Belgium*

Abstract

Dimension reduction is often an important step in the analysis of high-dimensional data. PCA is a popular technique to find the best low-dimensional approximation of high-dimensional data. However, classical PCA is very sensitive to atypical data. Robust methods to estimate the low-dimensional subspace that best approximates the regular data have been proposed. However, for high-dimensional data these algorithms become computationally expensive. Alternative algorithms for the robust subspace estimators are proposed that are better suited to compute the solution for high-dimensional problems. The main ingredients of the new algorithms are twofold. First, the principal directions of the subspace are estimated directly by iterating the first order solutions corresponding to the estimators. Second, to reduce the computation time even further five robust deterministic values are proposed to initialize the algorithms instead of using random starting values. It is shown that the new algorithms yield robust solutions and the computation time is largely reduced, especially for high-dimensional data.

Keywords: Deterministic algorithm, High-dimensional data, Least trimmed squares, M-scale, Principal Component Analysis

1. Introduction

Principal component analysis (PCA) is a popular exploratory tool for multivariate data. In particular, PCA is extremely useful to find a low-dimensional representation of high-dimensional data that yields the best possible approximation to the original data. Classical PCA minimizes the squared euclidean distances between the original observations and their orthogonal projections

*Corresponding author: KU Leuven, Department of Mathematics, Celestijnenlaan 200B, 3001 Leuven, Belgium. Phone: +3216372383, fax: +3216327998.

Email addresses: holgceva@espol.edu.ec (Holger Cevallos-Valdiviezo), Stefan.VanAelst@kuleuven.be (Stefan Van Aelst)

onto the lower dimensional subspace. However, classical PCA is very sensitive to atypical data due to the use of quadratic loss. Therefore, several approaches to robustify PCA have been proposed.

The earliest and easiest approach to robust PCA consists of taking the eigenvectors and eigenvalues of a robust scatter estimator instead of the standard sample covariance matrix (see e.g. Campbell, 1980; Devlin et al., 1981; Naga & Antille, 1990; Croux & Haesbroeck, 2000; Salibián-Barrera et al., 2006). However, this approach cannot be used for high-dimensional data because calculating high-dimensional robust scatter matrices is computationally complex or even infeasible if the sample size is small compared to the dimension. Moreover, while the efficiency of robust scatter estimators increases with dimension this comes at the expense of a loss of robustness. Therefore, Locantore et al. (1999) introduced spherical PCA which uses the covariance matrix of the data projected onto the unit sphere and is fast to compute.

A second approach to robust PCA sequentially looks for univariate directions that maximize a robust estimator of scale and are orthogonal to each other. This robust projection pursuit (PP) approach has been studied by e.g. Li & Chen (1985) and Croux & Ruiz-Gazen (1996, 2005).

Instead of looking for one direction at a time as in PP, one can robustly estimate a lower-dimensional subspace directly (see e.g. Liu et al., 2003; Croux et al., 2003). ROBPCA (Hubert et al., 2005) seeks for a lower-dimensional subspace via a multiple-step procedure. Briefly, the first step aims to identify a subset of $n/2 \leq h < n$ observations based on the Stahel-Donoho outlyingness. Classical PCA is then applied on this h -subset and the optimal dimension of the subspace is determined. Next, the data points are projected onto the lower-dimensional subspace. Finally, the reweighted MCD estimator is computed on the projected data to obtain the estimates for the principal directions of the subspace.

Principal Component Pursuit method (PCP) aims to decompose the data matrix into a low-rank component and a gross outlier component (Candés et al., 2011). See e.g. Chiang et al. (2016) and Rahmani & Atia (2017) for related work. However, PCP may fail to detect outliers in the orthogonal complement of the subspace (cfr. Maronna et al., 2015; She et al., 2016). Therefore, She et al. (2016) introduced a robust orthogonal complement (ROCPA) approach to deal with orthogonal complement outliers. The authors model the projected data onto the orthogonal complement subspace with a mean term, a sparse outlier matrix and a noise term. The sparse outlier matrix identifies orthogonal outliers. More recently, Brahma et al. (2018) combined ideas from PCP and ROCPA to construct a reinforced robust PCP (RRPCP) method that considers outliers in both the observation space and in the orthogonal complement subspace. RRPCP decomposes the data matrix into a low-rank component, a sparse outlier matrix to represent outliers in the observation space, a noise term, a mean term and another sparse outlier matrix to represent outliers in the orthogonal complement space. The two latter components are represented in the orthogonal complement space and then transformed back to the observation space. However, according to their model formulation ROCPA and

RRPCP have to estimate the directions of the orthogonal complement space, which becomes computationally intensive for high-dimensional data.

Maronna (2005) proposed to robustly estimate the best lower-dimensional approximation by minimizing either an M-scale or a least trimmed squares (LTS) scale of the Euclidean distances corresponding to the observations. He also characterized the solutions by the orthogonal complement directions and showed that these directions correspond to the eigenvectors associated with the smallest eigenvalues of a weighted covariance matrix. Based on this characterization Maronna (2005) proposed an iterative algorithm to compute the robust subspace estimators. The robustness of these subspace estimators has been widely investigated empirically. See e.g. Maronna (2005); Serneels & Verdonck (2008) and Tharrault et al. (2008) for the M-scale estimator, and Maronna (2005); Engelen et al. (2005) and Croux et al. (2017) for the LTS-scale estimator. Moreover, Croux et al. (2017) also provides a thorough theoretical study of the properties of the estimator based on the LTS scale.

However, in case of a low-dimensional approximation for high-dimensional data Maronna’s orthogonal complement algorithm requires to decompose a high dimensional covariance matrix and needs a large number of its eigenvectors to characterize the solution. This makes computing the subspace estimators time consuming or even infeasible in high dimensions. Therefore, we propose an algorithm for the robust subspace estimators of Maronna (2005) that directly calculates principal directions of the low-dimensional subspace.

The main ingredients of our new algorithm are twofold. First, we use the first order conditions corresponding to the estimator to update the principal directions of the subspace iteratively. This approach only requires low-dimensional vector and matrix operations rather than manipulating high-dimensional covariance matrices. Second, instead of using random starting values, similarly to Hubert et al. (2012) we propose five robust deterministic values to initialize the algorithm. These starting values yield robust fits that are usually close to the sought after robust solution, so that convergence occurs quickly.

The remainder of the manuscript is organized as follows. Section 2 reviews the robust subspace estimators based on an M-scale or the LTS-scale. Our definition is equivalent to the definition in Maronna (2005) but characterizes the solution by the principal directions of the subspace which better serves our needs for the development of the new algorithm in Section 3. In Section 4 we introduce the robust deterministic values to initialize the algorithm. In Section 5 we compare the solutions calculated with the new algorithm to those obtained with the original algorithm by means of an extensive simulation study. We also include other robust subspace estimators in the simulations to compare their robustness properties. Since the deterministic starting values are not orthogonally equivariant, in Section 6 we empirically investigate the effect of orthogonal transformations of the data on our algorithm with deterministic starting values. We compare all methods in terms of their computation time in Section 7, while in Section 8 we extend the simulation study in Section 5 to higher-dimensions. Section 9 contains a real data illustration and Section 10 presents our final conclusions.

2. Robust subspace estimators

Consider a data matrix $\mathbf{X} = (\mathbf{x}_1 \dots \mathbf{x}_n)^T \in \mathbb{R}^{n \times p}$ which contains the measurements of p variables for n observations. The goal is to approximate the n observations $\mathbf{x}_i$ by points $\hat{\mathbf{x}}_i$ that lie in a q -dimensional subspace. That is, $\hat{\mathbf{x}}_i \equiv \hat{\mathbf{x}}_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}) = \mathbf{m} + \mathbf{B}_q \mathbf{a}_i$ for some $\mathbf{m} \in \mathbb{R}^p$, $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_n)^T \in \mathbb{R}^{n \times q}$ and orthogonal matrix $\mathbf{B}_q \in \mathbb{R}^{p \times q}$, i.e. $\mathbf{B}_q^T \mathbf{B}_q = \mathbf{I}_q$. Let $\mathbf{b}_1, \dots, \mathbf{b}_p$ denote the rows of $\mathbf{B}_q$. The Euclidean distance between $\mathbf{x}_i$ and its approximation $\hat{\mathbf{x}}_i$ is denoted by $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}) = d_i = \|\mathbf{r}_i\|$, where $\mathbf{r}_i = \mathbf{x}_i - \hat{\mathbf{x}}_i$.

Maronna (2005) proposed to robustly estimate the optimal subspace by minimizing a robust scale estimator of the Euclidean distances $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})$. Note that if the nonrobust standard deviation is used, then the estimator minimizes the sum of squared Euclidean distances and thus the classical PCA solution is retrieved. Maronna's estimators are thus robust extensions of classical PCA dimension reduction.

Although the Euclidean distance between each observation $\mathbf{x}_i$ and its projection $\hat{\mathbf{x}}_i$ onto the q -dimensional subspace is measured in the $p - q$ dimensional orthogonal subspace in Maronna (2005), this is equivalent to our current formulation in the p -dimensional space.

Subspace S-estimator

The subspace S-estimator $(\hat{\mathbf{B}}_S, \hat{\mathbf{A}}_S, \hat{\mathbf{m}}_S)$ is obtained by minimizing an M-scale of the Euclidean distances $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})$. That is, $(\hat{\mathbf{B}}_S, \hat{\mathbf{A}}_S, \hat{\mathbf{m}}_S)$ is a solution of

$$\min_{\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}} \hat{\sigma}_M(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})), \quad (1)$$

over all $\mathbf{m} \in \mathbb{R}^p$, $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_n)^T \in \mathbb{R}^{n \times q}$ and orthogonal matrices $\mathbf{B}_q \in \mathbb{R}^{p \times q}$. For any $\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}) = (d_1, d_2, \dots, d_n)$, the corresponding M-scale $\hat{\sigma}_M(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$ is defined as the solution in s of

$$\frac{1}{n} \sum_{i=1}^n \rho\left(\frac{d_i}{s}\right) = b, \quad (2)$$

where ρ is an even function that is differentiable and non-decreasing on the positive real line with $\rho(0) = 0$ (see e.g. Maronna et al., 2006).

Similarly as in Boente & Salibian-Barrera (2015) first order conditions for the subspace S-estimator can be obtained by implicitly differentiating the M-scale in (2). The solutions of these first order conditions can be written as

$$\mathbf{a}_i = \left(\sum_{j=1}^p \mathbf{b}_j \mathbf{b}_j^T \right)^{-1} \left(\sum_{j=1}^p (x_{ij} - m_j) \mathbf{b}_j \right), \quad (3)$$

$$\mathbf{b}_j = \left(\sum_{i=1}^n w_i \mathbf{a}_i \mathbf{a}_i^T \right)^{-1} \left(\sum_{i=1}^n w_i (x_{ij} - m_j) \mathbf{a}_i \right), \quad (4)$$

and

$$m_j = \frac{\sum_{i=1}^n w_i (x_{ij} - \mathbf{a}_i^T \mathbf{b}_j)}{\sum_{i=1}^n w_i}, \quad (5)$$

with weights

$$w_i \equiv w(d_i) = \rho' \left(\frac{d_i}{\hat{\sigma}_M} \right) \frac{\hat{\sigma}_M}{d_i}, \quad (6)$$

for the rows of $\mathbf{A}_q$, $i = 1, \dots, n$, and the rows of $\mathbf{B}_q$ and components of $\mathbf{m}$, $j = 1, \dots, p$, respectively. Expressions (3)-(6) suggest an iterative reweighted least squares procedure. Initial values $\mathbf{B}_q^{(0)}$, $\mathbf{A}_q^{(0)}$ and $\mathbf{m}^{(0)}$ can be used to obtain starting weights $w_i^{(1)}$ from expression (6). Based on $\mathbf{B}_q^{(0)}$ and $\mathbf{m}^{(0)}$ one can obtain updated scores $\mathbf{A}_q^{(1)}$ from (3). Next, based on $\mathbf{A}_q^{(1)}$, $\mathbf{m}^{(0)}$ and weights $w_i^{(1)}$ one can obtain $\mathbf{B}_q^{(1)}$ from (4). Finally, based on $\mathbf{A}_q^{(1)}$, $\mathbf{B}_q^{(1)}$ and weights $w_i^{(1)}$ one can obtain $\mathbf{m}^{(1)}$ from (5). This is the first iteration. This procedure is repeated until we obtain a stable solution. Note that the resulting estimate of the matrix $\mathbf{B}_q$ is not necessarily orthogonal. Therefore, we orthogonalize this estimate at the end of the procedure and update the scores accordingly. Using the orthogonality of $\mathbf{B}_q$, it can be seen from (3) that $\mathbf{a}_i = \mathbf{B}_q^T (\mathbf{x}_i - \mathbf{m})$. Hence, once $\mathbf{B}_q$ and $\mathbf{m}$ are known, the corresponding scores $\mathbf{a}_i$ of the observations are easily obtained. By combining this result with (5) we can also obtain that $\mathbf{m} = \sum_{i=1}^n w_i \mathbf{x}_i / (\sum_{i=1}^n w_i)$. Note that if we put $w_i = 1$ for all observations, then the procedure reduces to an alternating least squares algorithm whose solution yields the classical PCA estimates (see also Gabriel & Zamir, 1979). Finally, by combining these expressions and using that $\hat{\mathbf{B}}_S^T \hat{\mathbf{B}}_S = \mathbf{I}_q$ it can also be derived that the subspace S-estimators $(\hat{\mathbf{B}}_S, \hat{\mathbf{m}}_S)$ satisfy the equation

$$\sum_{i=1}^n w_i (\mathbf{x}_i - \mathbf{m})(\mathbf{x}_i - \mathbf{m})^T \mathbf{B}_q = \mathbf{B}_q \Lambda. \quad (7)$$

where $\Lambda = \mathbf{B}^T \sum_{i=1}^n w_i (\mathbf{x}_i - \mathbf{m})(\mathbf{x}_i - \mathbf{m})^T \mathbf{B}$. From (7) and (1) it follows that the columns of $\hat{\mathbf{B}}_S$ correspond to the first q eigenvectors of the weighted covariance matrix

$$\mathbf{C}(\hat{\mathbf{m}}_S, \hat{\mathbf{B}}_S) = \frac{1}{n} \sum_{i=1}^n w_i (\mathbf{x}_i - \hat{\mathbf{m}}_S)(\mathbf{x}_i - \hat{\mathbf{m}}_S)^T, \quad (8)$$

which coincides with expression (9) of Maronna (2005).

Subspace LTS-estimator

The subspace LTS-estimator $(\hat{\mathbf{B}}_{\text{LTS}}, \hat{\mathbf{A}}_{\text{LTS}}, \hat{\mathbf{m}}_{\text{LTS}})$ is obtained by minimizing the LTS-scale of the Euclidean distances $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})$. That is, $(\hat{\mathbf{B}}_{\text{LTS}}, \hat{\mathbf{A}}_{\text{LTS}}, \hat{\mathbf{m}}_{\text{LTS}})$ is a solution of

$$\min_{\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}} \hat{\sigma}_{\text{LTS}}(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})), \quad (9)$$

over all $\mathbf{m} \in \mathbb{R}^p$, $\mathbf{A}_q \in \mathbb{R}^{n \times q}$ and orthogonal matrices $\mathbf{B}_q \in \mathbb{R}^{p \times q}$. For any $\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}) = (d_1, d_2, \dots, d_n)$, the corresponding LTS scale is defined as

$$\hat{\sigma}_{\text{LTS}}^2(\mathbf{d}) = \frac{1}{h} \sum_{i=1}^h d_{(i:n)}^2, \quad (10)$$

where $d_{(1:n)} \leq \dots \leq d_{(n:n)}$ are the ordered Euclidean distances and $h = n - \lfloor n\alpha \rfloor$ with $0 \leq \alpha < 0.5$. A fraction α of the observations is not taken into account when calculating the LTS estimator and thus α determines the robustness of the estimator.

Similarly as in Maronna (2005) it can be shown that the LTS solution satisfies the expressions (3)-(5) with weights now given by

$$w_i \equiv w(d_i) = \begin{cases} 1 & \text{if } d_i \leq d_{(h:n)}, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

With these weights the LTS estimator $(\hat{\mathbf{B}}_{\text{LTS}}, \hat{\mathbf{m}}_{\text{LTS}})$ also satisfies (7) and (8) which again coincides with expression (9) in Maronna (2005).

3. The algorithm

Maronna (2005) characterizes a q -dimensional subspace by an equation $\mathbf{B}_{p-q}^T \mathbf{x} = \mathbf{a}$ with $\mathbf{a} \in \mathbb{R}^{p-q}$ and $\mathbf{B}_{p-q} \in \mathbb{R}^{p \times (p-q)}$ an orthogonal matrix. His algorithm can be summarized by the following steps. First, 50 random orthogonal matrices $\mathbf{B}_{p-q}$ are generated to initialize the algorithm. The corresponding optimal value of $\mathbf{a}$ that minimizes the robust scale can then be calculated easily. A few iterative improvement steps are then applied to each of these starting values. In each iteration the weights w_i of the observations corresponding to the current solution are updated. Based on the updated weights, a new solution is then obtained by calculating the smallest $p - q$ eigenvectors of the weighted covariance matrix in (8). After two iterations, the 10 best solutions are selected and these are iterated further until convergence (with a maximum of 10 iterations).

If p is large and the subspace dimension q is small, Maronna's algorithm requires the storage of a high-dimensional covariance matrix and the calculation of a large number (namely, $p - q$) of its eigenvectors which becomes computationally demanding. To improve the computation time and reduce the memory load, we instead propose to directly calculate the q basis directions of the subspace by iterating expressions (3)-(5). Extensive experiments have shown that iterating these expressions only a few times suffices to obtain close approximations to the first eigenvectors of the weighted covariance matrix in (8). Note that our iterations of these expressions only require operations with q -dimensional vectors and matrices and thus will be more suitable for high-dimensional settings.

Algorithm 1 contains a detailed description of the main part of our modified algorithm in pseudo-code. The algorithm requires initial values of $\mathbf{B}_q$ and $\mathbf{m}$ as input. It also depends on tuning parameters N_1 , N_2 , N_3 and tol . tol specifies

the precision with which the solution is calculated. The tuning parameters N_1 and N_2 play the same role as in Maronna's original algorithm. That is, for each initial orthogonal matrix $\mathbf{B}_q$ first N_1 iterations are performed to improve the corresponding estimates of the location $\mathbf{m}$ and the scores matrix $\mathbf{A}_q$ while keeping $\mathbf{B}_q$ fixed. In the next N_2 iterations, the estimates of all three quantities $\mathbf{B}_q$, $\mathbf{m}$ and $\mathbf{A}_q$ are updated. Since these updates are now calculated by iterating expressions (3)-(5), an additional tuning parameter N_3 specifies how often these expressions are iterated. Using a similar proof as in Maronna (2005), it can easily be shown that either the M-scale $\hat{\sigma}_M$ or LTS scale $\hat{\sigma}_{LTS}$ decreases in each iteration of our algorithm. Moreover, the performance of our algorithm is similar to Maronna's algorithm if we use the same settings for the tuning parameters and start from the same initial values for $\mathbf{B}_q$ and $\mathbf{m}$.

Algorithm 1 Algorithm for the subspace S and LTS estimators

Input: $\mathbf{X}$, $\mathbf{B}_q$ (with $\mathbf{B}_q^T \mathbf{B}_q = \mathbf{I}_q$), $\mathbf{m}$, $N_1 = 3$, $N_2 = 2$, $N_3 = 3$, $tol = 1 \times 10^{-6}$.

1. Set $it \leftarrow 1$.
 - a. Compute $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_n)^T$ with $\mathbf{a}_i = \mathbf{B}_q^T(\mathbf{x}_i - \mathbf{m})$, $i = 1, \dots, n$.
 - b. Compute distances $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})$, $i = 1, \dots, n$.
 - c. Compute $\hat{\sigma}_0 = \hat{\sigma}_M(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$ or $\hat{\sigma}_0 = \hat{\sigma}_{LTS}(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$.
2. Do until $it = N_1 + N_2$ or $\Delta \leq tol$:
 - a. Compute w_i from (6) or (11) and update $\mathbf{m} = \frac{\sum_{i=1}^n w_i \mathbf{x}_i}{\sum_{i=1}^n w_i}$.
 - b. If $it > N_1$:
 - (1) Set $iter \leftarrow 1$ and $s_0^2 \leftarrow \hat{\sigma}_0^2$.
 - (2) Do until $iter = N_3$ or $\hat{\Delta} \leq tol$:
 - i. Compute $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_n)^T$, $\mathbf{B}_q = (\mathbf{b}_1, \dots, \mathbf{b}_p)^T$ and $\mathbf{m}$ from (3)-(5).
 - ii. Compute distances $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})$, $i = 1, \dots, n$.
 - iii. Compute $s = \hat{\sigma}_M(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$ or $s = \hat{\sigma}_{LTS}(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$.
 - iv. Set $iter \leftarrow iter + 1$, $\hat{\Delta} \leftarrow 1 - s^2/s_0^2$ and $s_0^2 \leftarrow s^2$.
 - (3) End do.
 - c. Compute $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_n)^T$ using (3) and update distances $d_i(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m})$, $i = 1, \dots, n$.
 - d. Compute $\hat{\sigma} = \hat{\sigma}_M(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$ or $\hat{\sigma} = \hat{\sigma}_{LTS}(\mathbf{d}(\mathbf{B}_q, \mathbf{A}_q, \mathbf{m}))$.
 - e. Set $\Delta \leftarrow 1 - \hat{\sigma}^2/\hat{\sigma}_0^2$ and $\hat{\sigma}_0^2 \leftarrow \hat{\sigma}^2$.
 - f. Set $it \leftarrow it + 1$.
3. End do.

Output: $\hat{\mathbf{B}}_q$, $\hat{\mathbf{A}}_q$, $\hat{\mathbf{m}}$.

For the subspace S-estimator we have used the popular Tukey biweight loss function $\rho(y) = \min(3y^2 - 3y^4 + y^6, 1)$ to calculate the M-scales. Two standard choices for the tuning parameters are $b = 0.5$ which yields the maximal breakdown point (BDP) of 50% and $b = 0.2426$ which yields a better compromise between efficiency and robustness (breakdown point $\approx 25\%$). For the subspace LTS estimator the trimming fraction α determines the breakdown point. The

most common choices are $\alpha = 0.5$ (50% BDP) and $\alpha = 0.25$ (25% BDP). See Maronna (2005) for more details on the breakdown point.

4. Starting values

It is well-known that the objective functions in (1) and (9) are nonconvex and thus may have several local minima. Not all of these minima correspond to robust solutions. A standard approach is to aim for the global minimum in (1) or (9) by generating many starting values. The solution that corresponds to the smallest value of the objective function that is reached by iterating these starting values is then the approximation for the global optimum (see e.g. Rousseeuw & Van Driessen, 1999; Salibián-Barrera & Yohai, 2006). In Maronna’s algorithm 50 random initial orthogonal matrices $\mathbf{B}_q$ are generated while the initial estimate for the location $\mathbf{m}$ is the coordinatewise median of the data. Each of these starting values is iterated $N_1 = 3$ times with $\mathbf{B}_q$ fixed, followed by $N_2 = 2$ iterations to improve the estimates of $\mathbf{B}_q$, $\mathbf{A}$ and $\mathbf{m}$ together. Then, the 10 best solutions (with the lowest scale) are selected and these are iterated further until convergence ($tol = 0.001$) with a maximum of $N'_2 = 10$ iterations. Extensive experiments by Maronna (2005) revealed that the algorithm shows good performance with these choices for the tuning parameters. Following Maronna (2005) we denote the solutions obtained with his algorithm by S-M when the M-scale is minimized and by S-L when the LTS scale is minimized.

Empirical comparisons have confirmed that if we use the same starting values for our new algorithm as well as the same settings for the tuning parameters, then we obtain a similar performance as Maronna’s algorithm. Note that our algorithm requires the additional tuning parameter N_3 which determines the number of iterations to calculate the updated estimates by iterating expressions (3)-(5). Extensive experiments showed that it suffices to use $N_3 = 3$ iterations to obtain good approximations that yield stable results. Note that while Maronna’s algorithm uses the fast-to-compute coordinatewise median as starting value for $\mathbf{m}$, we prefer to use as default the orthogonal equivariant spatial median which can also be computed efficiently (Vardi & Zhang, 2000). Moreover, we generate random orthogonal matrices $\mathbf{B}_q$ as starting values by the method of Stewart (1980) which consists of orthogonalizing a matrix of normal random numbers while Maronna (2005) instead orthogonalizes a matrix of uniform random numbers. We denote the solutions of our algorithm with random orthogonal starting values by RsubS for the subspace S-estimator and by RsubLTS for the subspace LTS-estimator.

While our adaptation of Maronna’s algorithm is indeed faster for high-dimensional data (see results in Section 7), the computation time remains considerable because a sufficient number of random starting values is needed to obtain a robust solution. While the default setting is to use a fixed number of 50 random orthogonal matrices as starting values, it is clear that the search space increases dramatically with increasing dimension. Hence, it can be expected that for high-dimensional data (many) more random starting values may

be needed to find a stable robust solution. Many algorithms for robust estimators encounter the same issue. To address this issue for the calculation of the minimum covariance determinant (MCD) estimator of multivariate location and scatter in high dimensions, Hubert et al. (2012) introduced a deterministic MCD algorithm (see also Hubert et al., 2015). The main idea is to replace the random starting values by a few well-chosen robust starting values. These robust starting values should be fast to compute while at the same time they are expected to lie close to a robust minimum of the objective function, which in our case is given by (1) or (9). Hence, instead of exploring the whole parameter space to find the optimum, a few robust starting values should point us to that part of the parameter space where robust solutions can be found. Since convergence from the robust starting values to their closest local minimum is generally fast as well, this results in an algorithm with a much lower computation time which allows us to handle problems in higher dimensions.

For the deterministic version of our algorithm, we could use as starting values for $\mathbf{B}_q$ the q largest eigenvectors corresponding to the robust starting values for the p -dimensional scatter matrix proposed by Hubert et al. (2012). However, we want to avoid having to calculate p -dimensional scatter matrices which may be unstable for large p and consumes a lot of memory. Therefore, inspired by Hubert et al. (2012) we propose five robust starting values for $\mathbf{B}_q$ which are obtained by the following procedure:

Step 1. Robustly standardize each variable X_j by subtracting its median and dividing by its Q_n scale estimate (Rousseeuw & Croux, 1993). Let $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_n)^T$ denote the standardized data matrix and $Z_1, \dots, Z_p$ the standardized variables. Consider the following transformations of the data:

- 1) Hyperbolic tangent (sigmoid) transformation: let $H_j = \tanh(Z_j)$, $j = 1, \dots, p$.
Set $\mathbf{U}'_1 = (H_1, \dots, H_p)$.
- 2) Rank transformation: let R_j be the vector of ranks of column X_j , $j = 1, \dots, p$.
Set $\mathbf{U}_2 = (R_1, \dots, R_p)$.
- 3) Normal scores: compute normal scores from the ranks R_j , i.e. $T_j = \Phi^{-1}[(R_j - 1/3)/(n + 1/3)]$, where $\Phi(\cdot)$ is the normal cumulative distribution function.
Set $\mathbf{U}_3 = (T_1, \dots, T_p)$.
- 4) Spatial signs: let $\mathbf{s}_i = \mathbf{z}_i / \|\mathbf{z}_i\|$, $i = 1, \dots, n$.
Set $\mathbf{U}'_4 = (\mathbf{s}_1, \dots, \mathbf{s}_n)^T$.

Robustly standardize the columns of $\mathbf{U}'_1$ and $\mathbf{U}'_4$ obtained above by subtracting their median and dividing by their Q_n scale estimate. Denote these standardized transformed data matrices by $\mathbf{U}_1$ and $\mathbf{U}_4$ and take $\mathbf{U}_5 = \mathbf{Z}$.

Step 2. Calculate the classical q -dimensional PCA subspace using the alternating least squares algorithm (see Algorithm 2 below) for $\mathbf{Y} = \mathbf{U}_k$. Let $\tilde{\mathbf{B}}_q^{(k)}$ be the resulting estimate of the q PC directions for $k = 1, \dots, 5$.

Step 3. For each of the score matrices $\tilde{\mathbf{A}}_q^{(k)} = \mathbf{Z}\tilde{\mathbf{B}}_q^{(k)}$, $k = 1, \dots, 5$, select the $l = \lceil n/2 \rceil$ rows with smallest Euclidean norm. Let $\mathcal{I}_k \subset \{1, \dots, n\}$ denote the indices of the selected rows.

Step 4. Set $\tilde{\mathbf{X}}^{(k)} = \mathbf{X}_{\mathcal{I}_k}$ for $k = 1, \dots, 5$. For each of these five reduced data matrices, calculate the classical q -dimensional PCA subspace by applying Algorithm 2 on $\mathbf{Y} = \tilde{\mathbf{X}}^{(k)} - \tilde{\mathbf{m}}^{(k)}$ where $\tilde{\mathbf{m}}^{(k)} = \frac{1}{l} \sum_{i=1}^l \tilde{\mathbf{x}}_i^{(k)}$. The resulting estimates of $\mathbf{B}_q$ together with $\tilde{\mathbf{m}}^{(k)}$ are the starting values for Algorithm 1.

Algorithm 2 PCA subspace algorithm

Input: $\mathbf{Y}, \mathbf{B}_q$ (with $\mathbf{B}_q^T \mathbf{B}_q = \mathbf{I}_q$), $N_3 = 3$, $tol = 1 \times 10^{-6}$.

1. Compute $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_l)^T$ with $\mathbf{a}_i = \mathbf{B}_q^T \mathbf{y}_i$.
2. Compute distances $d_i(\mathbf{B}_q, \mathbf{A}_q)$, $i = 1, \dots, l$.
3. Set $iter \leftarrow 1$ and compute $\hat{s}_0^2 = \frac{1}{l} \sum_{i=1}^l d_i^2$.
4. Do until $iter = N_3$ or $\tilde{\Delta} \leq tol$:
 - i. Compute $\mathbf{A}_q = (\mathbf{a}_1, \dots, \mathbf{a}_l)^T$ and $\mathbf{B}_q = (\mathbf{b}_1, \dots, \mathbf{b}_p)^T$ from (3)-(4) with weights $w_i = 1$, $i = 1, \dots, l$.
 - ii. Compute distances $d_i(\mathbf{B}_q, \mathbf{A}_q)$, $i = 1, \dots, l$.
 - iii. Compute $\hat{s}^2 = \frac{1}{l} \sum_{i=1}^l d_i^2$.
 - iv. Set $iter \leftarrow iter + 1$, $\tilde{\Delta} \leftarrow 1 - \hat{s}^2 / \hat{s}_0^2$ and $\hat{s}_0^2 \leftarrow \hat{s}^2$.
5. End do.

Output: $\hat{\mathbf{B}}_q, \hat{\mathbf{A}}_q$.

The alternating least squares algorithm for classical PCA, summarized in Algorithm 2, calculates the q -dimensional PCA subspace without performing singular value decomposition of a high-dimensional matrix, and has the advantage that it only requires operations with q -dimensional vectors and matrices in the iterative updating of the estimates in step 4.i. Note that for simplicity Algorithm 2 uses a centered data matrix as input data. Similarly as for Algorithm 1, we found that $N_3 = 3$ iterations in step 4 of Algorithm 2 were enough to obtain close approximations to the first eigenvectors of the classical covariance matrix. The q -dimensional canonical basis or a random orthogonal matrix can be used to initialize $\mathbf{B}_q$. In fact, we found that Algorithm 2 converges to the global solution independent of the initial value that is chosen for $\mathbf{B}_q$.

The procedure to obtain deterministic starting values can be summarized in two parts. The first part (first three steps) aims to identify subsets of the data with only regular observations. Transformations of the standardized data are

used in Step 1 to reduce the effect of outliers. In Step 2 naive robust estimates of the subspace are calculated from the transformed data which are used to project the data in Step 3 before obtaining clean subsets. This projection makes the algorithm better suitable for correlated data. In Section 6 we investigate by means of a simulation study the effect of orthogonal transformations of the data on the estimates obtained with these deterministic algorithms. The second part (last step) calculates the classical PCA subspace corresponding to each of the obtained subsets of the data matrix to obtain promising starting values for Algorithm 1. Similarly as for orthogonal starting values, each of these five starting values is iterated $N_2 = 2$ times to improve the estimates of $\mathbf{B}_q$, $\mathbf{A}$ and $\mathbf{m}$. Then, the best solution (with the lowest scale) is selected and iterated further until convergence ($tol = 1 \times 10^{-6}$) with a maximum of $N'_2 = 10$ iterations. We denote our algorithm with deterministic starting values by DsubS when the M-scale is minimized and by DsubLTS when the LTS scale is minimized.

5. Performance comparison

To compare our new algorithms to Maronna's original algorithms, we repeat the simulations of Maronna (2005). We use our algorithms with random orthogonal matrices as starting values as well as with deterministic starting values. For the simulations we choose the tuning parameters for the M-scale and LTS scale that yield the maximal breakdown point.

We also consider some other computationally efficient methods for robust subspace estimation. These are robust projection pursuit (PP) (Croux & Ruiz-Gazen, 1996), Spherical PCA (SPC) (Locantore et al., 1999) and ROBPCA (Hubert et al., 2005). For robust PP we use the algorithm of Croux & Ruiz-Gazen (2005) and maximize the LTS scale with $\alpha = 0.5$. Note that PP based on the M-scale was already considered in Maronna (2005). For ROBPCA we include the reweighting step as in Engelen et al. (2005). The number of random directions through two data points to compute the measure of outlyingness was fixed to 250 in all experiments as in Hubert et al. (2005). The parameter α in ROBPCA was set equal to $\alpha = 0.5$ for maximal robustness.

Following Maronna (2005) we generated $M = 200$ samples of size $n = 100$ and dimension $p = 10$. The regular observations are generated according to $N(\mathbf{0}, \mathbf{\Sigma})$ with $\mathbf{\Sigma} = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_p)$ where the following two designs are considered for the diagonal elements.

- a) An abrupt increase of the eigenvalues: $\lambda_j = 1 + 0.1j$ for $1 \leq j \leq (p - q)$ and $\lambda_j = 20(1 + 0.5(j - p + q))$ for $(p - q + 1) \leq j \leq p$.
- b) A smooth increase of the eigenvalues: $\lambda_j = 2^{j-1}$ for $1 \leq j \leq p$.

A fraction $\epsilon = 0\%$, 10% or 20% of outlying observations is generated from $N(k\mathbf{x}_0, 0.25\mathbf{\Sigma})$, where $\mathbf{x}_0$ is a vector of length p with $x_{0j} = 1$ for $j \leq (p - q)$ and 0 otherwise. The value of k runs between 0 and 20 with steps of 0.5. All methods are applied for $q = 2$, so the best two-dimensional linear subspace approximation is estimated.

To compare the methods we measure their prediction performance as in Maronna (2005). Hence, we measure the proportion of variance in independent regular data that remains unexplained by their approximation according to the estimated subspace. More formally, let $\mathbf{x}$ be a $N(\mathbf{0}, \Sigma)$ vector independent of the random sample used to obtain $\hat{\mathbf{B}}_q$. Then, the variability of $\mathbf{x}$ around the subspace generated by $\hat{\mathbf{B}}_q$ equals

$$E\|\mathbf{x} - \hat{\mathbf{B}}_q \hat{\mathbf{B}}_q^T \mathbf{x}\|^2 = \text{tr}[\Sigma] - \text{tr}[\hat{\mathbf{B}}_q^T \Sigma \hat{\mathbf{B}}_q],$$

and the prediction proportion of unexplained variance is given by

$$u_q^{\text{pred}} = \frac{E\|\mathbf{x} - \hat{\mathbf{B}}_q \hat{\mathbf{B}}_q^T \mathbf{x}\|^2}{\text{tr}[\Sigma]} = 1 - \frac{\text{tr}[\hat{\mathbf{B}}_q^T \Sigma \hat{\mathbf{B}}_q]}{\text{tr}[\Sigma]}. \quad (12)$$

Since Maronna's algorithms characterize the subspace by an estimate $\hat{\mathbf{B}}_{p-q} \in \mathbb{R}^{p \times (p-q)}$ of its orthogonal complement, the corresponding prediction proportion of unexplained variability in this case becomes:

$$u_q^{\text{pred}} = \frac{\text{tr}[\hat{\mathbf{B}}_{p-q}^T \Sigma \hat{\mathbf{B}}_{p-q}]}{\text{tr}[\Sigma]}. \quad (13)$$

Comparing this prediction error with the optimal value, given by

$$u_q^{\text{opt}} = \frac{\sum_{j=1}^{p-q} \lambda_j}{\sum_{j=1}^p \lambda_j}, \quad (14)$$

yields the relative prediction error

$$e_{\text{pred}} = \frac{u_q^{\text{pred}}}{u_q^{\text{opt}}} - 1. \quad (15)$$

Table 1 shows the mean of the relative prediction errors over $M = 200$ samples. The values of k included in Table 1 are those values at which some estimator attains its maximal error. Standard errors are not shown because they were small (below 0.08) in all cases. For $\epsilon = 20\%$ of contamination, Figures 1 and 2 show the mean relative prediction errors as a function of k . From $k = 10$ onwards these prediction errors stabilize so we only show results up to $k = 10$. These results show that the estimates obtained by our new RsubS and RsubLTS algorithms are very similar to the results of Maronna's S-M and S-L algorithms, as expected. Moreover, the DsubS and DsubLTS algorithms often succeed in obtaining more robust estimates than their counterparts based on random starts. The advantage of using robust starting values is most pronounced when the outliers are at a small to moderate distance of the majority ($k \leq 6$) as can be seen from Figures 1 and 2. Looking at the robust competitors, we can see that PPLTS turns out to be inefficient for $\epsilon = 0\%$, similarly as for the PP methods considered by Maronna (2005). Moreover, also for contaminated data PPLTS

Table 1: Mean of relative prediction errors: $\bar{e}_{\text{pred}}$.

Des.	ϵ	k	SPC	S-M	RsubS	DsubS	PPLTS	S-L	RsubLTS	DsubLTS	ROBPCA
a)	0%		0.04	0.02	0.02	0.02	0.32	0.07	0.07	0.06	0.03
		10% 1	0.04	0.03	0.03	0.03	0.41	0.10	0.11	0.08	0.03
		3	0.05	0.03	0.03	0.03	0.72	0.07	0.07	0.06	0.03
		6	0.05	0.03	0.03	0.03	0.93	0.07	0.07	0.06	0.03
	20%	1	0.06	0.03	0.03	0.03	0.48	0.28	0.28	0.15	0.04
		1.5	0.09	0.03	0.03	0.03	0.78	1.18	0.96	0.09	0.05
		2	0.15	0.03	0.03	0.03	1.34	1.31	0.84	0.06	0.04
		3	0.44	0.17	0.03	0.03	1.86	0.11	0.11	0.06	0.03
		3.5	0.57	0.11	0.03	0.03	1.92	0.09	0.08	0.06	0.03
		4.5	0.77	0.03	0.03	0.03	1.80	0.07	0.07	0.06	0.03
			0.05	0.04	0.04	0.04	0.27	0.14	0.14	0.12	0.05
b)	10%	1	0.09	0.07	0.07	0.06	0.34	0.23	0.22	0.21	0.10
		1.5	0.13	0.09	0.08	0.07	0.40	0.24	0.20	0.16	0.14
		2	0.16	0.10	0.09	0.08	0.46	0.15	0.16	0.14	0.11
		4	0.19	0.05	0.05	0.05	0.49	0.11	0.12	0.12	0.06
		5	0.20	0.05	0.05	0.05	0.45	0.11	0.11	0.12	0.06
	20%	1.5	0.46	0.45	0.44	0.35	0.67	0.71	0.70	0.40	0.42
		2	0.55	0.67	0.66	0.45	0.78	0.75	0.74	0.28	0.33
		3	0.60	0.71	0.71	0.31	0.73	0.49	0.44	0.12	0.16
		3.5	0.61	0.71	0.70	0.24	0.68	0.24	0.19	0.11	0.11
		5	0.62	0.69	0.69	0.07	0.57	0.11	0.11	0.11	0.07

Figure 1: Mean relative unexplained variance ($\bar{e}_{\text{pred}}$) corresponding to the estimated 2 dimensional subspace as a function of k for eigenvalue configuration a) and $\epsilon = 20\%$.

gives worse results in these settings. Spherical PCA is generally better than PPLTS but its performance decreases considerably for $\epsilon = 20\%$. ROBPCA performs well in these settings, which is in accordance with findings in Hubert et al. (2005); Engelen et al. (2005); Croux et al. (2017). Overall, ROBPCA and DsubS show the best performance in these simulations.

Figure 2: Mean relative unexplained variance ($\bar{e}_{\text{pred}}$) corresponding to the estimated 2 dimensional subspace as a function of k for eigenvalue configuration b) and $\epsilon = 20\%$.

6. Orthogonal equivariance

In the previous section we showed that the RsubS/RsubLTS algorithms yield similar estimates as the original S-M/S-L algorithms while the DsubS/DsubLTS often performed even better. Since orthogonal equivariance is important in the context of PCA, we now investigate the effect of orthogonal transformations on our deterministic algorithms by means of a simulation study. Similarly to Maronna & Zamar (2002), let $\mathbf{X} \in \mathbb{R}^{n \times p}$ denote the data matrix and let $\mathbf{P} \in \mathbb{R}^{p \times p}$ be a random orthogonal matrix, i.e. $\mathbf{P}^T \mathbf{P} = \mathbf{P} \mathbf{P}^T = \mathbf{I}_p$. Let $\hat{\mathbf{B}}_q(\mathbf{X})$ denote an estimate of the first q principal component directions for the data matrix $\mathbf{X}$. If the estimator is orthogonal equivariant, then it must hold that:

$$\hat{\mathbf{B}}_q(\mathbf{X} \mathbf{P}^T) = \mathbf{P} \hat{\mathbf{B}}_q(\mathbf{X}).$$

To investigate orthogonal equivariance, we therefore compare $\hat{\mathbf{B}}_q(\mathbf{X})$ to

$$\hat{\mathbf{B}}'_q(\mathbf{X}) = \mathbf{P}^T \hat{\mathbf{B}}_q(\mathbf{X} \mathbf{P}^T).$$

To measure the effect of the transformation we calculate the standardized last principal angle between the subspaces generated by the columns of $\hat{\mathbf{B}}_q$ and $\hat{\mathbf{B}}'_q(\mathbf{X})$, using the algorithm of Björck & Golub (1973). We divide this angle by $\pi/2$ to obtain values between 0 and 1. Values closer to 0 indicate that the subspaces are more comparable.

We generated data matrices $\mathbf{X}$ according to designs a) and b) in Section 5 with $\epsilon = 20\%$ and estimate the best $q = 2$ linear dimensional subspace. We generated $M = 200$ replicates and for each data replicate we generated a random orthogonal matrix $\mathbf{P}$ using the method of Stewart (1980). Figure 3 displays the mean standardized angles as a function of k for both designs. We can see that for both DsubS and DsubLTS the effect of the transformation is

Figure 3: Mean standardized angle corresponding to the estimated 2 dimensional subspace as a function of k for eigenvalue configuration a) in the left panel and configuration b) in the right panel. In both cases $\epsilon = 20\%$ of contamination was introduced.

largest when the outliers are at a small to moderate distance of the regular data. In these settings the DsubS algorithm is also clearly less affected by orthogonal transformations than DsubLTS and the effect quickly disappears when the outlier distance increases.

7. Computation time

We now compare the computation times of our algorithms to those of the original algorithms of Maronna and the other methods considered in Section 5. For this purpose we generalize the data generating model in design a). More specifically, we generate clean data from the p -dimensional normal distribution with center zero and diagonal covariance matrix. The eigenvalues on the diagonal are set equal to

$$\lambda_j = \begin{cases} 1 + c_{p,q} j & \text{for } 1 \leq j \leq (p - q), \\ 20(1 + 0.5(j - p + q)) & \text{for } (p - q + 1) \leq j \leq p. \end{cases} \quad (16)$$

The constant $c_{p,q}$ is chosen such that the q main directions explain about 80% of the total variability. We consider sample sizes $n = 1000$ or 5000 in dimension $p = 10, 500, 1000$ with $q = 2, 5$. The data contain $\epsilon = 20\%$ of contamination, generated as explained in Section 5 with $k = 15$.

All algorithms are implemented in R (R Core Team, 2016). For the S-M, S-L, RsubS, RsubLTS, DsubS, DsubLTS and PPLTS algorithms we used our own implementation. For S-M and S-L we use the R function `eigen()` to calculate the smallest $p - q$ eigenvectors of the weighted covariance matrix in (8). This function uses LAPACK (Anderson et al., 1999) routines which are written in FORTRAN. Alternatives such as the R function `svd()` for instance could be

used, but the speed difference is generally small. For spherical PCA we use the function `PcaLocantore()` and for ROBPCA we use the function `PcaHubert()`, both in R package `rrcov` (Todorov & Filzmoser, 2009). `PcaLocantore()` also uses `svd()` or `eigen()` to compute classical PCA on the projected data onto the unit sphere while `PcaHubert()` uses a FORTRAN implementation to compute the MCD estimator of location and scatter on the projected data and `eigen()` to compute its decomposition. To allow a fair comparison of computation times between these algorithms and our new algorithms, for our algorithms we implemented the iterative updating of the estimates based on expressions (3)-(5) in C++, using R package `RcppArmadillo` (Eddelbuettel & Sanderson, 2014). Note that the implementations of all methods in the comparison use plain R, except for the calculation of covariance matrices and/or its decomposition. S-M, S-L, spherical PCA and ROBPCA use FORTRAN implementations for this purpose while our algorithms perform the iterative updates in C++. The algorithms were compared on a single Intel i7 CPU (3.4GHz) machine running Windows 7.

Table 2 shows the computation time of the algorithms in seconds, averaged over $M = 50$ replications. Not surprisingly, we see that spherical PCA is the

Table 2: Computation times in seconds

		$n = 1000$			$n = 5000$		
	p	10	500	1000	10	500	1000
$q = 2$	SPC	0.04	1.34	4.19	0.16	5.67	21.59
	S-M	2.47	121.55	567.92	39.78	1178.71	3195.05
	RsubS	0.61	25.32	154.42	3.06	57.18	250.19
	DsubS	0.19	3.33	6.70	1.00	16.46	33.87
	PPLTS	5.75	25.91	52.97	142.07	703.32	1372.31
	S-L	0.60	95.46	502.53	2.66	288.50	1383.55
	RsubLTS	0.90	24.70	152.38	4.90	52.32	243.24
	DsubLTS	0.27	3.01	6.02	1.31	15.04	31.29
	ROBPCA	1.02	3.08	6.67	4.47	13.16	40.37
$q = 5$	SPC	0.04	1.34	4.21	0.18	5.74	22.01
	S-M	2.33	119.93	569.61	35.00	1170.10	3206.02
	RsubS	0.94	28.70	161.39	4.63	74.42	281.46
	DsubS	0.30	4.53	9.08	1.53	21.88	44.95
	PPLTS	14.61	64.40	122.36	356.52	1774.19	3385.82
	S-L	0.58	95.11	507.15	2.68	300.09	1369.12
	RsubLTS	1.70	28.80	161.07	9.25	75.68	282.14
	DsubLTS	0.45	4.12	8.29	2.25	20.35	41.36
	ROBPCA	1.04	3.07	6.65	4.65	13.55	40.78

fastest to compute because it does not require any iterative process (see also Maronna, 2005; Wilcox, 2008). For $n = 1000$ and $q = 2$ projection-pursuit is relatively fast to compute regardless of the dimension of the data, as can be seen from Table 2. However, the computation time grows quickly with increasing subspace dimension (case $q = 5$) and/or sample size (case $n = 5000$). The computation time of Maronna's algorithms is reasonable for low dimensional

data, but increases quickly when the dimension grows. This result was expected because these algorithms need to compute a large number $(p - q)$ of eigenvectors of a high-dimensional covariance matrix which is very time-consuming. Clearly, our implementations of the new algorithms with random starting values are faster than Maronna's original algorithms, but computation time still increases quickly with growing dimension. The new algorithms with deterministic starting values are even faster and their computation time increases much slower when the dimension grows.

When using 250 random directions, the computation time of ROBPCA is comparable to our deterministic algorithms for the different sample sizes and subspace dimensions considered. Increasing the number of directions considered to compute the measure of outlyingness makes ROBPCA much slower (results not shown here), especially for the case $n = 5000$ and $q = 5$. Thus, computation time of ROBPCA critically depends on the number of directions considered. These results indicate that the algorithms with deterministic starting values do not only show good performance but are also computationally attractive. They make it possible to compute robust subspace estimates for a large scale of problems.

8. Performance for high dimensional data

In Sections 5 and 7 we have shown that DsubS and DsubLTS show good performance for low-dimensional data and remain computationally attractive for high-dimensional data. We now investigate whether their good performance also carries over to the high-dimensional setting. To this end we consider the model of the previous section with eigenvalues according to (16) for $q = 2$ or $q = 5$, and $\epsilon = 20\%$ of contamination generated as in Section 5. We let k range between 0 and 20 as before and generated data of size $n = 1000$ in dimension $p = 1000$ or $p = 10000$.

For $p = 1000$ we did not include the S-L and S-M algorithms in the comparison because their performance was similar to the RsubS and RsubLTS algorithms, but computationally more demanding. Moreover, we did not consider the RsubS and RsubLTS algorithms in the case $p = 10000$ because these algorithms require too much computation time in these settings. SPC implemented with the R function `PcaLocantore()` also required excessive computation time for the case $p = 10000$. Therefore, we wrote an alternative SPC R function that computes classical PCA on the projected data using Algorithm 2.

Figures 4 and 5 show the mean relative prediction errors, averaged over $M = 200$ samples, when $p = 1000$ and $p = 10000$, respectively. Relative prediction errors again stabilized from $k = 10$ onwards, so we only present results up to $k = 10$. From Figure 4 we can see that in higher dimensions DsubS and DsubLTS behave similar as for low-dimensional data and clearly perform better than SPC and PPLTS. Figure 5 shows that DsubS and DsubLTS keep their excellent performance when the dimension is increased further to $p = 10000$. The performance of ROBPCA is mostly similar to our deterministic algorithms in this case, but it has slightly more difficulty to estimate the best

$q = 5$ dimensional subspace when the outliers are at a small distance ($k = 1.5$) from the regular data. Similar conclusions were obtained for sample size $n = 100$ so these results are not shown.

Figure 4: Mean relative prediction errors $\bar{e}_{pred}$ for the case $n = 1000$ and $p = 1000$ with $q = 2$ subspace estimation (left panel) and $q = 5$ subspace estimation (right panel). In both cases the level of contamination is $\epsilon = 20\%$.

Figure 5: Mean relative prediction errors $\bar{e}_{pred}$ for the case $n = 1000$ and $p = 10000$ with $q = 2$ subspace estimation (left panel) and $q = 5$ subspace estimation (right panel). In both cases the level of contamination is $\epsilon = 20\%$.

9. Real data example

We illustrate our deterministic algorithms on real image data that have been vectorized to form a high-dimensional dataset with $p \gg n$. We use the Extended Yale Face Database B (Georghiades et al., 2001; Lee et al., 2005) which contains aligned grayscale face images of 38 subjects under the same frontal pose and 64 different illumination conditions. This database contains cropped face images which have resolution 192×168 ($=32256$) pixels each. For the analysis we randomly sampled 11 subjects and selected for each subject the 6 images with the highest illumination contrasts, so that the face characteristics are clearly identifiable. In particular, we used the images with the following codes for light conditions: "P00A+000E+00.pgm", "P00A+010E-20.pgm", "P00A-010E-20.pgm", "P00A+000E-20.pgm", "P00A-005E-10.pgm" and "P00A+005E-10.pgm". Similarly to Rahmani & Atia (2017), next to these 66 face images we also sampled 9 non-face images from the Caltech 101 database (Fei-Fei et al., 2007) which constitute potential outliers. The Caltech 101 database contains 9144 color and grayscale images from 102 object categories that includes vehicles, plants, animals and cartoon characters. We only considered grayscale images with a height of at least 192 pixels and a width of at least 168 pixels. We randomly sampled 9 object images from this subset and cropped the sampled images to 192 pixel height and 168 pixel width when necessary. Combining the vectorized face and object images resulted in 75 observations in $p = 192 \times 168 = 32256$ dimensions. The last 9 observations (rows 67-75 in the data matrix) correspond to the object images. Hence, 12% of the observations in the dataset are contamination. Figure 6 displays a random subset of the face images while Figure 7 displays the 9 object images in the data matrix. The goal of our analysis is to recognize the faces and thus to flag the object images as outliers. We compare the results obtained by DsubLTS and DsubS to those of SPC, PPLTS and ROBPCA. We estimate the $q = 2$ dimensional subspace which explains about 80% of the robustly estimated total variability.

Figure 6: Random subset of the face images in the image data example.

To identify outliers, we use the diagnostic plot introduced by Hubert et al. (2005). It plots the (robust) orthogonal distances between the observations

Figure 7: Object images in the image data example.

and the estimated subspace versus (robust) score distances of the projected observations in the subspace with respect to their center. To identify outlying observations cutoff values for both the robust orthogonal distances and the robust score distances were proposed by Hubert et al. (2005). To compute the robust score distances robust estimates for the variances according to the basis directions within the estimated subspace are required. Since our algorithms do not yield such estimates, we estimate these variances robustly using univariate LTS or M-scales of the scores corresponding to the estimated basis directions. Also SPC does not yield such variance estimates, so we computed the MCD estimator on the scores to compute robust distances in the PCA subspace, similarly to Hubert et al. (2005). Note that as in Section 8, we used our SPC implementation in R based on Algorithm 2 for these high-dimensional data.

Figure 8 shows the diagnostic plots of the solutions obtained by DsubLTS, DsubS, SPC, PPLTS and ROBPCA. All robust methods are able to flag the 9 object images as outliers with respect to the estimated 2-dimensional subspace. Object images are flagged as either orthogonal outliers or bad leverage points by the different methods. Note that SPC also detects four faces as outliers, but these false positives are merely borderline cases. This example illustrates that also in very high-dimensional settings our deterministic algorithms are able to robustly estimate a low-dimensional subspace that best represents the regular observations and allows to identify outliers.

10. Conclusions

We proposed new algorithms for the robust subspace estimation methods proposed by Maronna (2005). These algorithms directly estimate principal directions of the subspace, which makes them more suitable for high-dimensional problems. For the starting values of the algorithm we considered random orthogonal matrices, as well as five deterministic starting values. These well-chosen deterministic starting values can be computed fast in high-dimensional settings

because they avoid the need to calculate a high-dimensional scatter matrix and only use simple manipulations of the data. Our experiments show that the deterministic algorithms yield results that are as good as or better than the results for the algorithms with random starting values, while having a much lower computation time. Moreover, the excellent performance of our deterministic algorithms carries over to high-dimensional settings. Maronna’s algorithms and our algorithms with random orthogonal matrices attempt to find the global minimum of their objective function by using a sufficient number of starting values. On the other hand, our algorithms with deterministic starting values aim to find this minimum by using a few well-chosen robust but rough initial estimates of the subspace. While this may decrease the probability of obtaining the global minimum, it does often lead to a robust local minimum. This explains why these deterministic algorithms show a better performance.

For the deterministic algorithms orthogonal equivariance is not guaranteed anymore. However, our experiments have shown that in most cases there is not a serious loss of equivariance. Note that the computation time of our deterministic algorithms can be reduced further on multi-core machines by calculating each of the five starting solutions in parallel on different cores. Finally, the subspace S-estimator usually shows a better compromise between robustness and efficiency than the subspace LTS-estimator. Implementations of our algorithms in R (R Core Team, 2016) are available from the website <http://wis.kuleuven.be/stat/robust/software>.

Acknowledgments

The research by Van Aelst was supported by the Internal Funds KU Leuven under Grant C16-15-068 and COST Action IC1408 CRoNoS. Their support is gratefully acknowledged.

Anderson, E., Bai, Z., Bischof, C., Blackford, L. S., Demmel, J., Dongarra, J. J., Du Croz, J., Hammarling, S., Greenbaum, A., McKenney, A., & Sorensen, D. (1999). *LAPACK Users’ Guide (Third Ed.)*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics. URL: http://www.netlib.org/lapack/lug/lapack_lug.html.

Björck, Å., & Golub, G. H. (1973). Numerical methods for computing angles between linear subspaces. *Mathematics of Computation*, 27, 579–594. URL: <http://www.jstor.org/stable/2005662>.

Boente, G., & Salibian-Barrera, M. (2015). S-estimators for functional principal component analysis. *Journal of the American Statistical Association*, 110, 1100–1111. URL: <http://dx.doi.org/10.1080/01621459.2014.946991>. doi:10.1080/01621459.2014.946991. arXiv:<http://dx.doi.org/10.1080/01621459.2014.946991>.

Brahma, P. P., She, Y., Li, S., Li, J., & Wu, D. (2018). Reinforced robust principal component pursuit. *IEEE Transactions on Neural Networks and Learning Systems*, PP, 1–14. doi:10.1109/TNNLS.2017.2671849.

- Campbell, N. A. (1980). Robust procedures in multivariate analysis i: Robust covariance estimation. *Applied Statistics*, 29, 231–237.
- Candés, E. J., Li, X., Ma, Y., & Wright, J. (2011). Robust principal component analysis? *J. ACM*, 58, 11:1–11:37. URL: <http://doi.acm.org/10.1145/1970392.1970395>. doi:10.1145/1970392.1970395.
- Chiang, K.-Y., Hsieh, C.-J., & Dhillon, I. S. (2016). Robust principal component analysis with side information. In *International Conference on Machine Learning (ICML)*.
- Croux, C., Filzmoser, P., Pison, G., & Rousseeuw, P. J. (2003). Fitting multiplicative models by robust alternating regressions. *Statistics and Computing*, 13, 23–36. URL: <http://dx.doi.org/10.1023/A:1021979409012>. doi:10.1023/A:1021979409012.
- Croux, C., García-Escudero, L. A., Gordaliza, A., Ruwet, C., & Martín, R. S. (2017). Robust principal component analysis based on trimming around affine subspaces. *Statistica Sinica*, 27, 1437–1459.
- Croux, C., & Haesbroeck, G. (2000). Principal component analysis based on robust estimators of the covariance or correlation matrix: influence functions and efficiencies. *Biometrika*, 87, 603–618. URL: <http://biomet.oxfordjournals.org/content/87/3/603.short>. doi:10.1093/biomet/87.3.603.
- Croux, C., & Ruiz-Gazen, A. (1996). A fast algorithm for robust principal components based on projection pursuit. In A. Prat (Ed.), *COMP-STAT: Proceedings in Computational Statistics 12th Symposium held in Barcelona, Spain, 1996* (pp. 211–216). Heidelberg: Physica-Verlag HD. URL: http://dx.doi.org/10.1007/978-3-642-46992-3_22. doi:10.1007/978-3-642-46992-3_22.
- Croux, C., & Ruiz-Gazen, A. (2005). High breakdown estimators for principal components: The projection-pursuit approach revisited. *Journal of Multivariate Analysis*, 95, 206–226. doi:10.1016/j.jmva.2004.08.002.
- Devlin, S. J., Gnanadesikan, R., & Kettenring, J. R. (1981). Robust Estimation of Dispersion Matrices and Principal Components. *Journal of the American Statistical Association*, 76, 354–362. URL: <http://www.jstor.org/stable/2287836>. doi:10.2307/2287836.
- Eddelbuettel, D., & Sanderson, C. (2014). Rcpparmadillo: Accelerating r with high-performance c++ linear algebra. *Computational Statistics and Data Analysis*, 71, 1054–1063. URL: <http://dx.doi.org/10.1016/j.csda.2013.02.005>.
- Engelen, S., Hubert, M., & Vanden Branden, K. (2005). A Comparison of Three Procedures for Robust PCA in High Dimensions. *Austrian Journal of Statistics*, 34, 117–126.

- Fei-Fei, L., Fergus, R., & Perona, P. (2007). Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. *Comput. Vis. Image Underst.*, *106*, 59–70. URL: <http://dx.doi.org/10.1016/j.cviu.2005.09.012>. doi:10.1016/j.cviu.2005.09.012.
- Gabriel, K. R., & Zamir, S. (1979). Lower rank approximation of matrices by least squares with any choice of weights. *Technometrics*, *21*, 489–498. URL: <http://www.jstor.org/stable/1268288>.
- Georghiades, A. S., Belhumeur, P. N., & Kriegman, D. J. (2001). From few to many: illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *23*, 643–660. doi:10.1109/34.927464.
- Hubert, M., Rousseeuw, P., Vanpaemel, D., & Verdonck, T. (2015). The dets and detmm estimators for multivariate location and scatter. *Computational Statistics & Data Analysis*, *81*, 64 – 75. URL: <http://www.sciencedirect.com/science/article/pii/S0167947314002175>. doi:<https://doi.org/10.1016/j.csda.2014.07.013>.
- Hubert, M., Rousseeuw, P. J., & Vanden Branden, K. (2005). ROBPCA: A New Approach to Robust Principal Component Analysis. *Technometrics*, *47*, 64–79. doi:10.1198/004017004000000563.
- Hubert, M., Rousseeuw, P. J., & Verdonck, T. (2012). A deterministic algorithm for robust location and scatter. *Journal of Computational and Graphical Statistics*, *21*, 618–637. URL: <http://dx.doi.org/10.1080/10618600.2012.672100>
~~delimiter"026E30F\$nh~~[http://www.tandfonline.com/doi/abs/10.1080/10618600.2012.672100#.UdeZCOA-nCM\\$](http://www.tandfonline.com/doi/abs/10.1080/10618600.2012.672100#.UdeZCOA-nCM$)
~~delimiter"026E30F\$nh~~<http://www.tandfonline.com/doi/abs/10.1080/10618600.2012.672100>.
doi:10.1080/10618600.2012.672100.
- Lee, K.-C., Ho, J., & Kriegman, D. J. (2005). Acquiring linear subspaces for face recognition under variable lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *27*, 684–698. doi:10.1109/TPAMI.2005.92.
- Li, G., & Chen, Z. (1985). Projection-pursuit approach to robust dispersion matrices and principal components: primary theory and Monte Carlo. *Journal of the American Statistical Association*, *80*, 759–766. URL: <http://www.tandfonline.com/doi/abs/10.1080/01621459.1985.10478181>
~~delimiter"026E30F\$npapers3:~~
~~//publication/uuid/E7A6D092-1420-4117-A47C-C6F759DA3121.~~
doi:10.1080/01621459.1985.10478181.
- Liu, L., Hawkins, D. M., Ghosh, S., & Young, S. S. (2003). Robust singular value decomposition analysis of microarray data. *Proceedings of the National Academy of Sciences of the United States of America*, *100*,

- 13167–72. URL: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=263735&tool=pmcentrez&rendertype=abstract>. doi:10.1073/pnas.1733249100.
- Locantore, N., Marron, J. S., Simpson, D. G., Tripoli, N., Zhang, J. T., & Cohen, K. L. (1999). Robust principal component analysis for functional data. *Test*, 8, 1–73. URL: <http://dx.doi.org/10.1007/BF02595862>. doi:10.1007/BF02595862.
- Maronna, R. A. (2005). Principal Components and Orthogonal Regression Based on Robust Scales. *Technometrics*, 47, 264–273. URL: <http://pubs.amstat.org/doi/abs/10.1198/004017005000000166>. doi:10.1198/004017005000000166.
- Maronna, R. A., Martin, R. D., & Yohai, V. J. (2006). *Robust Statistics: Theory and Methods*. Wiley.
- Maronna, R. A., Méndez, F., & Yohai, V. J. (2015). Robust nonlinear principal components. *Statistics and Computing*, 25, 439–448. URL: <http://dx.doi.org/10.1007/s11222-013-9442-0>. doi:10.1007/s11222-013-9442-0.
- Maronna, R. A., & Zamar, R. H. (2002). Robust estimates of location and dispersion for high-dimensional datasets. *Technometrics*, 44, 307–317. URL: <http://www.jstor.org/stable/1271538>.
- Naga, R. A., & Antille, G. (1990). Stability of robust and non-robust principal components analysis. *Comput. Stat. Data Anal.*, 10, 169–174. URL: [http://dx.doi.org/10.1016/0167-9473\(90\)90062-M](http://dx.doi.org/10.1016/0167-9473(90)90062-M). doi:10.1016/0167-9473(90)90062-M.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing Vienna, Austria. URL: <https://www.R-project.org/>.
- Rahmani, M., & Atia, G. K. (2017). Coherence pursuit: Fast, simple, and robust principal component analysis. *IEEE Transactions on Signal Processing*, 65, 6260–6275.
- Rousseeuw, P. J., & Croux, C. (1993). Alternatives to the Median Absolute Deviation. *Journal of the American Statistical Association*, 88, 1273–1283. URL: <http://www.jstor.org/stable/2291267>. doi:10.1080/01621459.1993.10476408.
- Rousseeuw, P. J., & Van Driessen, K. (1999). A fast algorithm for the minimum covariance determinant estimator. *Technometrics*, 41, 212–223. URL: <http://dx.doi.org/10.2307/1270566>. doi:10.2307/1270566.
- Salibián-Barrera, M., Aelst, S. V., & Willems, G. (2006). Principal components analysis based on multivariate MM estimators

- with fast and robust bootstrap. *Journal of the American Statistical Association*, 101, 1198–1211. URL: <http://dx.doi.org/10.1198/016214506000000096>. doi:10.1198/016214506000000096. arXiv:<http://dx.doi.org/10.1198/016214506000000096>.
- Salibián-Barrera, M., & Yohai, V. J. (2006). A Fast Algorithm for S-Regression Estimates. *Journal of Computational and Graphical Statistics*, 15, 414–427. doi:10.1198/106186006X113629.
- Serneels, S., & Verdonck, T. (2008). Principal component analysis for data containing outliers and missing elements. *Computational Statistics & Data Analysis*, 52, 1712 – 1727. URL: <http://www.sciencedirect.com/science/article/pii/S0167947307002241>. doi:<https://doi.org/10.1016/j.csda.2007.05.024>.
- She, Y., Li, S., & Wu, D. (2016). Robust Orthogonal Complement Principal Component Analysis. *Journal of the American Statistical Association*, 111, 763–771. URL: <http://dx.doi.org/10.1080/01621459.2015.1042107>. doi:10.1080/01621459.2015.1042107.
- Stewart, G. W. (1980). The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM Journal on Numerical Analysis*, 17, 403–409.
- Tharrault, Y., Mourot, G., & Ragot, J. (2008). Fault detection and isolation with robust principal component analysis. In *2008 16th Mediterranean Conference on Control and Automation* (pp. 59–64). doi:10.1109/MED.2008.4602224.
- Todorov, V., & Filzmoser, P. (2009). An object-oriented framework for robust multivariate analysis. *Journal of Statistical Software*, 32, 1–47. URL: <http://www.jstatsoft.org/v32/i03/>.
- Vardi, Y., & Zhang, C.-H. (2000). The multivariate l1-median and associated data depth. *Proceedings of the National Academy of Sciences*, 97, 1423–1426. URL: <http://www.pnas.org/content/97/4/1423.abstract>. doi:10.1073/pnas.97.4.1423. arXiv:<http://www.pnas.org/content/97/4/1423.full.pdf>.
- Wilcox, R. R. (2008). Robust principal components: A generalized variance perspective. *Behavior Research Methods*, 40, 102–108. doi:10.3758/BRM.40.1.102.