

The Cellwise Minimum Covariance Determinant Estimator

Citation for published version (APA):

Raymaekers, J., & Rousseeuw, P. J. (2022). *The Cellwise Minimum Covariance Determinant Estimator*. Cornell University - arXiv. arXiv.org No. 2207.13493 <https://doi.org/10.48550/arXiv.2207.13493>

Document status and date:

Published: 27/07/2022

DOI:

[10.48550/arXiv.2207.13493](https://doi.org/10.48550/arXiv.2207.13493)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

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

[Link to publication](#)

General rights

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

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

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

www.umlib.nl/taverne-license

Take down policy

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

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

The Cellwise Minimum Covariance Determinant Estimator

Jakob Raymaekers

Department of Quantitative Economics, Maastricht University, The Netherlands

Peter J. Rousseeuw*

Section of Statistics and Data Science, University of Leuven, Belgium

July 27, 2022

Abstract

The usual Minimum Covariance Determinant (MCD) estimator of a covariance matrix is robust against casewise outliers. These are cases (that is, rows of the data matrix) that behave differently from the majority of cases, raising suspicion that they might belong to a different population. On the other hand, cellwise outliers are individual cells in the data matrix. When a row contains one or more outlying cells, the other cells in the same row still contain useful information that we wish to preserve. We propose a cellwise robust version of the MCD method, called cellMCD. Its main building blocks are observed likelihood and a sparsity penalty on the number of flagged cellwise outliers. It possesses good breakdown properties. We construct a fast algorithm for cellMCD based on concentration steps (C-steps) that always lower the objective. The method performs well in simulations with cellwise outliers, and has high finite-sample efficiency on clean data. It is illustrated on real data with visualizations of the results.

Keywords: Cellwise outliers, Covariance matrix, Likelihood, Missing values, Sparsity.

1 Motivation

Any practicing statistician or data scientist knows that real data sets often contain outliers. One definition of outliers says that they are cases that do not obey the fit suggested by the majority of the data, which raises suspicion that they may have been generated by

*This research was funded by projects of Internal Funds KU Leuven.

a different mechanism. Since cases typically correspond to rows of the data matrix, they are often called rowwise outliers. They may be the result of gross errors, but they can also be nuggets of valuable information. In either case, it is important to find them. In computer science this is called anomaly detection, and in some areas it is known as exception mining. In statistics several approaches were tried, such as testing for outliers and the computation of outlier diagnostics. In our experience the approach working best is that of robust statistics, which aims to fit the majority of the data first, and then flags outliers by their large deviation from that fit.

In this paper we focus on single-class multivariate numerical data without a response variable (although the results are relevant for classification and regression too). The goal is to robustly estimate the central location of the point cloud as well as its covariance matrix, and at the same time flag the outliers that may be present. The underlying model is that the data come from a multivariate Gaussian distribution, in which some data has been replaced by outliers that can be anywhere.

The Minimum Covariance Determinant (MCD) estimator introduced by Rousseeuw (1984, 1985) is highly robust to casewise outliers. Its definition is quite intuitive. Take an integer h that is at least half the sample size n . We then look for the subset containing h cases such that the determinant of its usual covariance matrix is as small as possible. The resulting robust location estimate is then the mean of that subset, and the robust covariance matrix is its covariance matrix multiplied by a consistency factor. One can show that the estimates are not overly affected when there are fewer than $n - h$ outlying cases. The MCD became computationally feasible with the algorithm of Rousseeuw and Van Driessen (1999), followed by even faster algorithms by Hubert et al. (2012) and De Ketelaere et al. (2020). Copt and Victoria-Feser (2004) computed the MCD for incomplete data. The MCD has also been generalized to high dimensions (Boudt et al., 2020), and to non-elliptical distributions using kernels (Schreurs et al., 2021). For a survey on the MCD and its applications see Hubert et al. (2018). The MCD is available in the procedure ROBUSTREG in SAS, in SAS/IML, in Matlab's PLS Toolbox, as the function `covMcd` in the R package *robustbase* (Maechler et al., 2022) on CRAN, and as the function `CovMcd` in the R package *rrcov* (Todorov, 2012). In Python one can use `MinCovDet` in *scikit-learn* (Pedregosa et al., 2011).

In recent times a different outlier paradigm has gained prominence, that of *cellwise*

outliers, first published by Alqallaf et al. (2009). It assumes that some individual cells (entries) of the data matrix deviate from what they should have been, perhaps due to gross measurement errors, whereas the remaining cells in the same row still contain useful information that we want to fit. The difference between the casewise and the cellwise paradigm is illustrated in Figure 1. In the left panel the outlying cases are shown as black rows. In the panel on the right the cellwise outliers correspond to fewer black squares in total, but together they contaminate over half of the cases, so the existing methods for casewise outliers may fail.

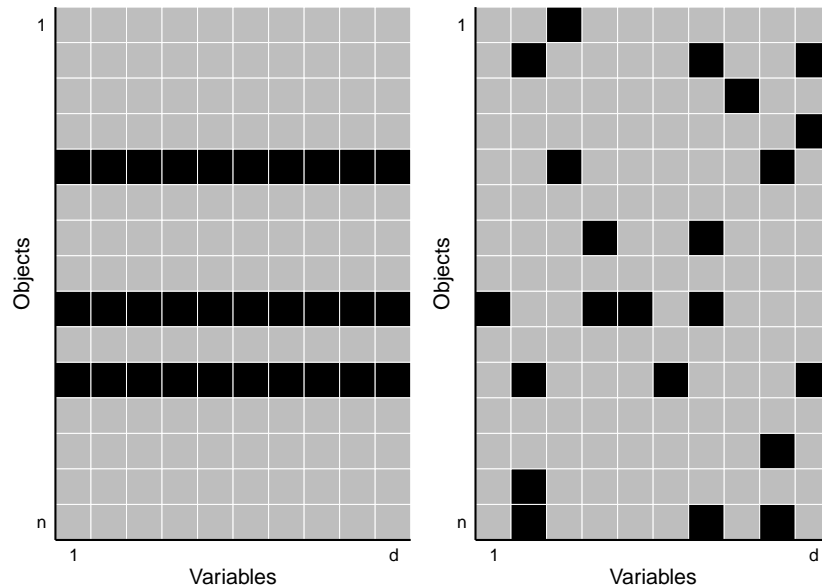


Figure 1: The casewise (left) and cellwise (right) outlier paradigms. (Black means outlying.)

In reality we do not know in advance *which* cells in the right panel of Figure 1 are outlying (black), unlike the simpler problem of incomplete data where we do know which cells are missing. When the variables have substantial correlations the cellwise outliers need not be marginally outlying, and then it can be quite hard to detect them. Van Aelst et al. (2011) proposed one of the first detection methods, based on an outlyingness measure of the Stahel-Donoho type. Rousseeuw and Van den Bossche (2018) predict the values of all cells and flag the observed cells that differ much from their prediction. Debruyne et al. (2019) consider casewise outliers and ask which variables contribute the most to their outlyingness. The O3 plot of Unwin (2019) visualizes cases that are outlying in fewer dimensions than the entire dataset.

There has been some work on estimating the underlying covariance matrix in the presence of cellwise outliers. One approach is to compute robust covariances between each pair of variables, and to assemble them in a matrix. To estimate these pairwise covariances, Öllerer and Croux (2015) and Croux and Öllerer (2016) use rank correlations. Tarr et al. (2016) instead use the robust pairwise correlation estimator of Gnanadesikan and Kettenring (1972) in combination with the robust scale estimator Q_n of Rousseeuw and Croux (1993). As the resulting matrix is not necessarily positive semidefinite (PSD), they then compute the nearest PSD matrix by the method of Higham (2002). Raymaekers and Rousseeuw (2021a) obtain a PSD covariance matrix by transforming (‘wrapping’) the original data variables. Another approach is the two-step generalized S-estimator (2SGS) of Agostinelli et al. (2015) and Leung et al. (2017). It starts with a filter to detect cellwise outliers. These cells are then set to missing, and the generalized S-estimator of Danilov et al. (2012) is run. Finally, the DI method of Raymaekers and Rousseeuw (2021b) alternates the detection of outlying cells with estimation of the covariance matrix.

Many cellwise robust methods were developed for specific settings, such as principal components (Hubert et al., 2019), discriminant analysis (Aerts and Wilms, 2017), clustering (García-Escudero et al., 2021), graphical models (Katayama et al., 2018), low-rank approximation (Maronna and Yohai, 2008), regression (see Öllerer et al. (2016) and Filzmoser et al. (2020)), variable selection (Su et al., 2021), and compositional data (Štefelová et al., 2021). Also, isolated outliers in functional data (Hubert et al., 2015) can be seen as cellwise outliers.

In the next section we introduce the cellwise MCD estimator, followed by its breakdown properties in section 3. Section 4 describes its algorithm. Some illustrations on real data are shown in section 5. The performance of the method is studied by simulation in section 6, and section 7 concludes with a discussion.

2 A cellwise MCD

We first note that the *casewise* MCD can be reformulated in terms of likelihood. The likelihood of a d -variate Gaussian distribution is

$$f(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} e^{-\text{MD}^2(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma})/2} \quad (1)$$

where $\boldsymbol{\mu}$ is a column vector, $\boldsymbol{\Sigma}$ is a positive definite matrix, and the Mahalanobis distance is $\text{MD}(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sqrt{(\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})}$. For a sample $\mathbf{x}_1, \dots, \mathbf{x}_n$ we put $L(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) := -2 \ln(f(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}))$ so the maximum likelihood estimator (MLE) of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ minimizes

$$\sum_{i=1}^n L(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{i=1}^n (\ln |\boldsymbol{\Sigma}| + d \ln(2\pi) + \text{MD}^2(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})) . \quad (2)$$

Let us now look for a subset $H \subset \{1, \dots, n\}$ with h elements which minimizes (2) where the sum is only over i in H . We can also write this with weights w_i that are 0 or 1 in the objective $\sum_{i=1}^n w_i L(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})$, so we minimize

$$\sum_{i=1}^n w_i (\ln |\boldsymbol{\Sigma}| + d \ln(2\pi) + \text{MD}^2(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})) \quad (3)$$

under the constraint that $\sum_{i=1}^n w_i = h$.

For the minimizing set of weights w_i we know from maximum likelihood that $\hat{\boldsymbol{\mu}}$ is the mean of the \mathbf{x}_i in H , so it is the weighted mean of all \mathbf{x}_i , and similarly

$$\hat{\boldsymbol{\Sigma}} = \frac{1}{h} \sum_{i=1}^n w_i (\mathbf{x}_i - \hat{\boldsymbol{\mu}})(\mathbf{x}_i - \hat{\boldsymbol{\mu}})^\top . \quad (4)$$

But then the third term of (3) becomes

$$\begin{aligned} \sum_{i=1}^n w_i (\mathbf{x}_i - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}) &= \sum_{i=1}^n \text{trace}(w_i (\mathbf{x}_i - \hat{\boldsymbol{\mu}})(\mathbf{x}_i - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1}) = \\ \text{trace}\left(\sum_{i=1}^n w_i (\mathbf{x}_i - \hat{\boldsymbol{\mu}})(\mathbf{x}_i - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1}\right) &= \text{trace}(h \hat{\boldsymbol{\Sigma}} \hat{\boldsymbol{\Sigma}}^{-1}) = hd \end{aligned}$$

which is constant, and so is the second term. Therefore minimizing (3) is equivalent to minimizing the determinant of (4), which is the definition of the casewise MCD.

In the context of incomplete data, Dempster et al. (1977) and others defined the *observed likelihood*. Let us denote the missingness pattern of the $n \times d$ data matrix \mathbf{X} by the $n \times d$ matrix \mathbf{W} with entries w_{ij} that are 0 for missing x_{ij} and 1 otherwise. Its rows \mathbf{w}_i take the place of the scalar weights w_i in (3). For the Gaussian model the observed likelihood of the i th observation (Little and Rubin, 2020) is given by:

$$f(\mathbf{x}_i^{(\mathbf{w}_i)}, \boldsymbol{\mu}^{(\mathbf{w}_i)}, \boldsymbol{\Sigma}^{(\mathbf{w}_i)}) := \frac{1}{(2\pi)^{d(\mathbf{w}_i)/2} |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}|^{1/2}} e^{-\text{MD}^2(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})/2} \quad (5)$$

in which

$$\text{MD}(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) := \sqrt{(\mathbf{x}^{(\mathbf{w}_i)} - \boldsymbol{\mu}^{(\mathbf{w}_i)})^\top (\boldsymbol{\Sigma}^{(\mathbf{w}_i)})^{-1} (\mathbf{x}_i^{(\mathbf{w}_i)} - \boldsymbol{\mu}^{(\mathbf{w}_i)})} \quad (6)$$

is called the *partial Mahalanobis distance* by Danilov et al. (2012). Here $\mathbf{x}_i^{(\mathbf{w}_i)}$ is the vector with only the entries for which $w_{ij} = 1$, and similarly for $\boldsymbol{\mu}^{(\mathbf{w}_i)}$. The matrix $\boldsymbol{\Sigma}^{(\mathbf{w}_i)}$ is the submatrix of $\boldsymbol{\Sigma}$ containing only the rows and columns of the variables j with $w_{ij} = 1$. Finally, $d(\mathbf{w}_i)$ is the number of entries in $\mathbf{x}_i^{(\mathbf{w}_i)}$, i.e. the number of non-missing entries of \mathbf{x}_i . By convention, a case \mathbf{x}_i consisting exclusively of NA's has $d(\mathbf{w}_i) = 0$, $\text{MD}(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = 0$ and $|\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| = 1$. Putting $L(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) := -2 \ln(f(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}))$ we see that maximizing the observed likelihood of the entire data set comes down to minimizing

$$\sum_{i=1}^n L(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{i=1}^n (\ln |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| + d(\mathbf{w}_i) \ln(2\pi) + \text{MD}^2(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})) . \quad (7)$$

This maximum likelihood estimate of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is typically computed by the EM algorithm (Dempster et al., 1977).

When constructing a cellwise MCD, h can no longer be the number of *cases* to be included. Instead, we use it for the number of *cells* to be included per column. We could minimize

$$\sum_{i=1}^n (\ln |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| + d(\mathbf{w}_i) \ln(2\pi) + \text{MD}^2(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})) \quad (8)$$

under the constraints $\xi_d(\boldsymbol{\Sigma}) \geq a$ and $\|\mathbf{W}_{\cdot j}\|_0 \geq h$ for all $j = 1, \dots, d$

over $(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{W})$. The first constraint says that the smallest eigenvector of $\boldsymbol{\Sigma}$ is at least as large as a number $a > 0$, where the eigenvalues of $\boldsymbol{\Sigma}$ are denoted as $\xi_1(\boldsymbol{\Sigma}) \geq \dots \geq \xi_d(\boldsymbol{\Sigma})$. This ensures that $\boldsymbol{\Sigma}$ is nonsingular, which is required to compute Mahalanobis distances. In the second constraint, $\|\mathbf{W}_{\cdot j}\|_0$ is the number of nonzero entries in the j -th column of \mathbf{W} . Note that we should not choose h too low. Whereas for the casewise MCD we can take h as low as $0.5n$, that would be ill-advised here because it could happen that two variables j and k do not overlap in the sense that $w_{ij}w_{ik} = 0$ for all i , making it impossible to estimate their covariance. We will impose that $h \geq 0.75n$ throughout.

However, we found that minimizing (8) tends to flag too many cells. This is because a value of h that is suitable for one variable may be too low for another, and we do not know ahead of time which variables have many outlying cells and which have few or none. To avoid flagging too many cells, we add a penalty counting the number of flagged cells in

each column. The objective function of the **cellwise MCD** (cellMCD) then becomes

$$\sum_{i=1}^n (\ln |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| + d(\mathbf{w}_i) \ln(2\pi) + \text{MD}^2(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})) + \sum_{j=1}^d \lambda_j \|\mathbf{1}_d - \mathbf{W}_{.j}\|_0 \quad (9)$$

under the constraints $\xi_d(\boldsymbol{\Sigma}) \geq a$ and $\|\mathbf{W}_{.j}\|_0 \geq h$ for all $j = 1, \dots, d$.

The notation $\|\mathbf{1}_d - \mathbf{W}_{.j}\|_0$ stands for the number of nonzero elements in this vector, so the number of zero weights in column j of \mathbf{W} , i.e. the number of flagged cells in column j of \mathbf{X} . In this way we reduce the number of flagged cells, while still keeping the previous constraint that $\|\mathbf{W}_{.j}\|_0 \geq h$. We'll discuss the choice of the λ_j later. Combining a sparsity penalty with a $\|\cdot\|_0$ constraint is not new, see the work of She et al. (2022) on casewise robust regression.

We think that cellMCD is the first cellwise robust method combining the fitting of the parameters $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and the flagging of outlying cells (\mathbf{W}) in one objective function. The constraint $\|\mathbf{W}_{.j}\|_0 \geq h$ for $j = 1, \dots, d$ says that we require at least h unflagged cells in each column. In order to avoid a singular covariance matrix, we obviously need $h > d$. Combining these inequalities we obtain $n > 4d/3$. But the curse of dimensionality implies that many spurious structures can be found in increasing dimensions, so we want a more comfortable ratio of cases per dimension. For the casewise MCD the rule of thumb is $n/d \geq 5$ (Rousseeuw and van Zomeren, 1990), and we will require that here too.

The cellMCD method defined by (9) is equivariant for permuting the cases, for shifting the data, and for multiplying the variables by nonzero constants. But unlike the casewise MCD it is not equivariant under general nonsingular linear transformations, or even orthogonal transformations. This is because cells are intimately tied to the coordinate system, and an orthogonal transformation changes the cells. This is an important difference between the casewise and cellwise approaches. For instance, consider the standard multivariate Gaussian model in dimension $d = 4$ with the suspicious point $(10, 0, 0, 0)$. By an orthogonal transformation of the data, this point can be moved to $(\sqrt{50}, \sqrt{50}, 0, 0)$ or to $(5, 5, 5, 5)$. The casewise MCD is equivariant to such transformations and will still flag the same case. But in the cellwise paradigm $(10, 0, 0, 0)$ has one outlying cell, $(\sqrt{50}, \sqrt{50}, 0, 0)$ has two, and $(5, 5, 5, 5)$ has four, so cellMCD will react differently, as it should.

3 Breakdown properties

Alqallaf et al. (2009) define the cellwise breakdown value of a location estimator. Here we will focus on finite-sample breakdown values in the sense of Donoho and Huber (1983) and Lopuhaä and Rousseeuw (1991). The *finite-sample cellwise breakdown value* of an estimator $\hat{\boldsymbol{\mu}}$ at a dataset \mathbf{X} is given by the smallest fraction of cells per column that need to be replaced to carry the estimate outside all bounds. Formally, let \mathbf{X} be a dataset of size n , and denote by \mathbf{X}^m any corrupted sample obtained by replacing at most m cells in each column of \mathbf{X} by arbitrary values. Then the finite-sample cellwise breakdown value of a location estimator $\hat{\boldsymbol{\mu}}$ at \mathbf{X} is given by

$$\varepsilon_n^*(\hat{\boldsymbol{\mu}}, \mathbf{X}) = \min \left\{ \frac{m}{n} : \sup_{\mathbf{X}^m} \|\hat{\boldsymbol{\mu}}(\mathbf{X}^m) - \hat{\boldsymbol{\mu}}(\mathbf{X})\| = \infty \right\}. \quad (10)$$

Analogously to the casewise setting, we can also define the *cellwise explosion breakdown value* of a covariance estimator $\hat{\boldsymbol{\Sigma}}$ as

$$\varepsilon_n^+(\hat{\boldsymbol{\Sigma}}, \mathbf{X}) = \min \left\{ \frac{m}{n} : \sup_{\mathbf{X}^m} \xi_1(\hat{\boldsymbol{\Sigma}}) = \infty \right\}. \quad (11)$$

Moreover, we define the *cellwise implosion breakdown value* of $\hat{\boldsymbol{\Sigma}}$ as

$$\varepsilon_n^-(\hat{\boldsymbol{\Sigma}}, \mathbf{X}) = \min \left\{ \frac{m}{n} : \inf_{\mathbf{X}^m} \xi_d(\hat{\boldsymbol{\Sigma}}) = 0 \right\}. \quad (12)$$

We will typically assume that the original data set \mathbf{X} is in *general position*, meaning that no more than d points lie in any $d - 1$ dimensional affine subspace. In particular, no three points lie on a line, no 4 points lie on a plane, and so on. When the data are drawn from a continuous distribution, it is in general position with probability 1.

Although the breakdown value definitions above look similar to those for casewise contamination, there is an important difference concerning implosion. The casewise implosion breakdown value of the classical covariance matrix \mathbf{Cov} at a dataset in general position is very high, in fact it is $(n - d)/n$ which goes to 1 for increasing sample size n . This is because whenever $d + 1$ of the original data points are kept, \mathbf{Cov} remains nonsingular. In stark contrast, its *cellwise* implosion breakdown value is quite low:

$$\varepsilon_n^-(\mathbf{Cov}, \mathbf{X}) = \left\lceil \frac{n - d}{d} \right\rceil / n \leq \frac{1}{d}. \quad (13)$$

To see why, let us pick d points of \mathbf{X} which lie on a hyperplane that is not parallel to any coordinate axis, which is possible due to general position. In the remaining $n - d$ rows we

can then replace a single cell such that all of the resulting points lie on the same hyperplane, so \mathbf{Cov} becomes singular. We can do this by replacing no more than $\lceil (n-d)/d \rceil$ cells in each variable, which is a fraction $\lceil (n-d)/d \rceil/n$ of its n cells.

The fact that we only need to replace a fraction $1/d$ of cells per variable to make the entire dataset coplanar is also bad news for most other covariance matrix estimators. This makes it useful to rule out implosion by a constraint like $\xi_d(\widehat{\Sigma}) \geq a$ in (9). (This constraint will also guarantee that the algorithm converges.) Such a constraint would be at odds with affine equivariance, but here we are in the cellwise setting.

Due to the constraint which is built into the definition (9) of cellMCD, its cellwise implosion breakdown value is 1. We also want to know the breakdown value of its location estimate $\widehat{\boldsymbol{\mu}}$ and the explosion breakdown value of $\widehat{\Sigma}$. These naturally depend on the choice of h . The result below would hold even for h as low as $\lfloor \frac{n}{2} \rfloor + 1$, but as explained before this could lead to some poorly defined covariances and numerical instability, so we stick with our earlier recommendation of $h \geq 0.75n$ (and in fact $h = 0.75n$ is the default in our implementation).

Proposition 1. *If the dataset \mathbf{X} is in general position and $h \geq 0.75n$, the cellMCD estimators $\widehat{\boldsymbol{\mu}}$ and $\widehat{\Sigma}$ satisfy the properties*

$$(a) \ \varepsilon_n^-(\widehat{\Sigma}, \mathbf{X}) = 1$$

$$(b) \ \varepsilon_n^+(\widehat{\Sigma}, \mathbf{X}) \geq (n - h + 1)/n$$

$$(c) \ \varepsilon_n^*(\widehat{\boldsymbol{\mu}}, \mathbf{X}) \geq (n - h + 1)/n$$

(d) *The lower bound $(n - h + 1)/n$ is sharp.*

Proposition 1 shows that cellMCD is highly robust. The proof is in section A.1 of the Supplementary Material.

4 Algorithm

In the algorithm we will need the following result about decomposing the Mahalanobis distance and the likelihood.

Proposition 2. *Let us split the d -variate case \mathbf{x} into two nonempty blocks, and split $\boldsymbol{\mu}$ and the $d \times d$ positive definite matrix $\boldsymbol{\Sigma}$ accordingly, like*

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{bmatrix} \quad \boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{bmatrix} \quad \boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{bmatrix}.$$

Then $MD^2(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = (\mathbf{x}_i - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu})$ and $L(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = -2 \ln(f(\mathbf{x}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}))$ satisfy

$$MD^2(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = MD^2(\mathbf{x}_1, \hat{\mathbf{x}}_1, \mathbf{C}_1) + MD^2(\mathbf{x}_2, \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{22}) \quad (14)$$

$$L(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = L(\mathbf{x}_1, \hat{\mathbf{x}}_1, \mathbf{C}_1) + L(\mathbf{x}_2, \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{22}) \quad (15)$$

for $\hat{\mathbf{x}}_1 = \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} \mathbf{x}_2$ and $\mathbf{C}_1 = \boldsymbol{\Sigma}_{11} - \boldsymbol{\Sigma}_{12} \boldsymbol{\Sigma}_{22}^{-1} \boldsymbol{\Sigma}_{21}$.

The proof can be found in section A.1 in the Supplementary Material.

The proposition can be interpreted as follows. Take a case \mathbf{x}_i with some but not all cells missing, and for simplicity assume that its missing components come first. Then put $\mathbf{x}_1 = \mathbf{x}_i^{(1-w_i)}$ and \mathbf{x}_2 the remainder. If $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ are the true underlying parameters, $\hat{\mathbf{x}}_1$ is the conditional expectation $E[\mathbf{X}_1 | \mathbf{X}_2 = \mathbf{x}_2]$ and \mathbf{C}_1 is the conditional covariance matrix $\text{Cov}[\mathbf{X}_1 | \mathbf{X}_2 = \mathbf{x}_2]$. The additivity in (14) and (15) justifies the use of the partial Mahalanobis distances and the observed likelihood in our setting. Moreover, the fact that the difference of two ‘nested’ MD^2 is again an MD^2 and hence non-negative implies that the MD^2 is monotone for nested sets of variables. In particular, if \mathbf{x} is observed fully we can write

$$\begin{aligned} MD^2(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) &= \frac{r^2(x_1 | x_2, \dots, x_d)}{s^2(X_1 | x_2, \dots, x_d)} + \frac{r^2(x_2 | x_3, \dots, x_d)}{s^2(X_2 | x_3, \dots, x_d)} + \dots + \frac{r^2(x_{d-1} | x_d)}{s^2(X_{d-1} | x_d)} + \frac{(x_d - \mu_d)^2}{\Sigma_{dd}} \end{aligned} \quad (16)$$

where each time s^2 is the matrix \mathbf{C}_1 (which is a scalar here) and the residuals are $r(x_1 | x_2, \dots, x_d) = x_1 - \hat{x}_1(x_2, \dots, x_d)$ and so on. Note that (16) holds for any order of the d variables. However, in each order the relative contribution of variable j to the total $MD^2(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma})$ may be different. For the likelihood we obtain similarly

$$\begin{aligned} L(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) &= L(x_1, \mu_1, C_{1|2,\dots,d}) + L(x_2, \mu_2, C_{2|3,\dots,d}) + \dots \\ &\quad + L(x_{d-1}, \mu_{d-1}, C_{d-1|d}) + L(x_d, \mu_d, \Sigma_{dd}) \end{aligned} \quad (17)$$

in which the terms do not need to be positive.

If in definition (9) of cellMCD we set $\lambda_j = 0$ and use casewise weights, i.e. rowwise constant w_{ij} , we recover the original casewise MCD. The crucial ingredient for computing

the latter is the concentration step (C-step) of Rousseeuw and Van Driessen (1999). After each C-step the new objective value is less than or equal to the old objective value, so iterating C-steps always converges. We will now construct a C-step for cellMCD with the same properties. Let us denote the current solution of cellMCD by $\hat{\boldsymbol{\mu}}^{(k)}$, $\hat{\boldsymbol{\Sigma}}^{(k)}$, and $\mathbf{W}^{(k)}$. Then the C-step proceeds as follows.

Part (a) of the C-step. In this part we update the matrix \mathbf{W} in (9) while keeping $\hat{\boldsymbol{\mu}}^{(k)}$ and $\hat{\boldsymbol{\Sigma}}^{(k)}$ unchanged. We start the new pattern $\widetilde{\mathbf{W}}$ as $\widetilde{\mathbf{W}} = \mathbf{W}^{(k)}$, and then we modify $\widetilde{\mathbf{W}}$ column by column, by cycling over the variables $j = 1, \dots, d$. The fact that this job can be done by column is advantageous for maintaining the constraint. Assume we are working on column j of $\widetilde{\mathbf{W}}$, possibly after having modified other columns of $\widetilde{\mathbf{W}}$ already. The current pattern of variable j is $\widetilde{\mathbf{W}}_{.j}$ and we want to obtain a new pattern for column j to reduce the objective while leaving the other columns of $\widetilde{\mathbf{W}}$ unchanged. Note that we can write the objective (9) as $\sum_{i=1}^n \widetilde{L}(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\lambda})$ where

$$\widetilde{L}(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}) = \ln |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| + d(\mathbf{w}_i) \ln(2\pi) + \text{MD}^2(\mathbf{x}_i, \mathbf{w}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma}) + \sum_{j=1}^d \lambda_j |1 - w_{ij}|$$

with $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_d)$. For each $i = 1, \dots, n$ we compute the difference in the total objective (9) between putting $\widetilde{w}_{ij} = 1$ and putting $\widetilde{w}_{ij} = 0$, which is

$$\begin{aligned} \Delta_{ij} &= \widetilde{L}(\mathbf{x}_i, \widetilde{w}_{ij} = 1, \hat{\boldsymbol{\mu}}^{(k)}, \hat{\boldsymbol{\Sigma}}^{(k)}, \boldsymbol{\lambda}) - \widetilde{L}(\mathbf{x}_i, \widetilde{w}_{ij} = 0, \hat{\boldsymbol{\mu}}^{(k)}, \hat{\boldsymbol{\Sigma}}^{(k)}, \boldsymbol{\lambda}) \\ &= \ln |\boldsymbol{\Sigma}^{(\widetilde{w}_{ij}=1)}| - \ln |\boldsymbol{\Sigma}^{(\widetilde{w}_{ij}=0)}| + \ln(2\pi) + \text{MD}^2(x_{ij}, \hat{x}_{ij}, C_{ij}) - \lambda_j \\ &= \ln(C_{ij}) + \ln(2\pi) + (x_{ij} - \hat{x}_{ij})^2 / C_{ij} - \lambda_j \end{aligned} \quad (18)$$

where the second and third equalities use Proposition 2 in which \hat{x}_{ij} and C_{ij} are now scalars. Note that $\hat{x}_{ij} = \hat{\mu}_j^{(k)} + \hat{\boldsymbol{\Sigma}}_{j,o}^{(k)} (\hat{\boldsymbol{\Sigma}}_{o,o}^{(k)})^{-1} (\hat{\mathbf{x}}_{i,o} - \hat{\boldsymbol{\mu}}_o^{(k)})$ is the conditional expectation of the cell X_{ij} conditional on the observed (subscript ‘o’) cells in row i , i.e. those with $\widetilde{w}_i = 1$, taking into account any earlier modifications to $\widetilde{\mathbf{W}}$. Analogously, $C_{ij} = \hat{\boldsymbol{\Sigma}}_{j,j}^{(k)} - \hat{\boldsymbol{\Sigma}}_{j,o}^{(k)} (\hat{\boldsymbol{\Sigma}}_{o,o}^{(k)})^{-1} \hat{\boldsymbol{\Sigma}}_{o,j}^{(k)}$ is the conditional variance of X_{ij} . So we assign

$$\widetilde{w}_{ij} = \begin{cases} 1 & \text{if } \ln(C_{ij}) + \ln(2\pi) + (x_{ij} - \hat{x}_{ij})^2 / C_{ij} \leq \lambda_j \\ 0 & \text{otherwise.} \end{cases} \quad (19)$$

If fewer than h of the Δ_{ij} are negative, we set $\widetilde{w}_{ij} = 1$ for the i with the h smallest Δ_{ij} and 0 otherwise. After cycling through all columns of $\widetilde{\mathbf{W}}$ we set $\mathbf{W}^{(k+1)} = \widetilde{\mathbf{W}}$.

Part (b) of the C-step. Keeping the new pattern $\mathbf{W}^{(k+1)}$ fixed we now want to update $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$. As $\mathbf{W}^{(k+1)}$ is fixed the penalty term in (9) does not enter the minimization, so we are in the situation of the objective (7) for incomplete data, where the EM algorithm can be used. We first carry out one E-step which computes conditional means and products for the data entries with $\mathbf{W}_{ij}^{(k+1)} = 0$, for all rows. Next, we carry out an M-step, followed by imposing the constraint $\xi_d \geq a$ by truncating the eigenvalues of $\hat{\boldsymbol{\Sigma}}$ from below at a . The C-step ends by reporting $\mathbf{W}^{(k+1)}$, $\hat{\boldsymbol{\mu}}^{(k+1)}$ and $\hat{\boldsymbol{\Sigma}}^{(k+1)}$.

Proposition 3. *Each C-step turns a triplet $(\hat{\boldsymbol{\mu}}^{(k)}, \hat{\boldsymbol{\Sigma}}^{(k)}, \mathbf{W}^{(k)})$ satisfying the constraints in (9) into a new triplet $(\hat{\boldsymbol{\mu}}^{(k+1)}, \hat{\boldsymbol{\Sigma}}^{(k+1)}, \mathbf{W}^{(k+1)})$ which satisfies the same constraints and whose objective (9) is less than or equal to before.*

For the proof see section A.1 in the Supplementary Material. Many variations of the C-step are possible, such as updating only one column of $\widetilde{\mathbf{W}}$ at a time in part (a), or cycling through all columns of $\widetilde{\mathbf{W}}$ more than once. We could also run more than one EM-step in part (b), etc. But in our experiments such variations were not faster than the current version, which is quick.

The algorithm iterates C-steps, and converges because the objective decreases in each C-step (when it remains the same the algorithm is done) and there is a finite lower bound on the objective (9). To see the latter, first consider a fixed matrix \mathbf{W} . Then the first term satisfies $\ln(|\boldsymbol{\Sigma}^{(w_i)}|) = \ln(\prod_j \xi_j(\boldsymbol{\Sigma}^{(w_i)})) = \sum_j \ln(\xi_j(\boldsymbol{\Sigma}^{(w_i)})) \geq \|w_i\|_0 \ln(\xi_d(\boldsymbol{\Sigma})) \geq \|w_i\|_0 \ln(a)$ which is finite, and all the other terms are bounded below by zero. The overall lower bound is the minimum of such lower bounds over the finite number of possible matrices \mathbf{W} that satisfy the constraint, so it is finite.

Note that cellMCD can still be used when the data contains missing cells, indicated by u_{ij} which are 0 for missing cells and 1 elsewhere. In that situation we first have to remove variables with more than $n - h$ missing values. In the C-step it then suffices to force $w_{ij} = 0$ whenever $u_{ij} = 0$.

In order to start our C-steps we need an initial estimator. In our experiments we found that the DDCW estimator of Raymaekers and Rousseeuw (2021b) gives good results and is very fast. It is a combination of the DetectDeviatingCells (DDC) method of Rousseeuw and Van den Bossche (2018) and the fast correlation method in (Raymaekers and Rousseeuw, 2021a). DDCW is described in section A.2 of the Supplementary Material.

The only remaining question is how to select the λ_j but this is quite simple, we do not need cross-validation or an information criterion. In (18) the term $(x_{ij} - \hat{x}_{ij})^2/C_{ij}$ is the square of the residual $x_{ij} - \hat{x}_{ij}$ standardized robustly. For inlying cells this should be below a cutoff, for which we take the chi-squared quantile $\chi_{1,p}^2$ with one degree of freedom and probability p . The term $\ln(C_{ij})$ is close to the average $\sum_{i=1}^n \ln(C_{ij}^*)/n$ where the C_{ij}^* are from the initial estimator. So we propose to set each λ_j equal to

$$\lambda_j = \chi_{1,p}^2 + \ln(2\pi) + \frac{1}{n} \sum_{i=1}^n \ln(C_{ij}^*). \quad (20)$$

Therefore we only have to choose a single cutoff probability p to generate all λ_j automatically. From simulations and examples we found that $p = 0.99$ was a good choice overall, so it is set as the default.

The algorithm has been implemented as the R function `cellMCD()`. It starts by checking the data for non-numerical variables, cases with too many NA's and so on. Next, it robustly standardizes the variables, and then computes the initial estimator followed by C-steps until convergence. The constraint $\xi_d(\hat{\Sigma}) \geq a$ is applied to the standardized data, with default $a = 10^{-4}$. The function also reports the number of flagged cells in each variable. All the plots in the next section were made by the companion function `plot_cellMCD()`. Both functions will be incorporated in the R package *cellWise* on CRAN.

5 Illustration on real data

We will illustrate `cellMCD` on the cars data obtained from the Top Gear website by Alfons (2016), focusing on the 11 numerical variables `price`, `displacement`, `horsepower`, `torque`, `acceleration time`, `top speed`, `miles per gallon`, `weight`, `length`, `width`, and `height`. This dataset is popular because both the variables and the cases (the cars) can easily be interpreted. After removing two cars with mostly NA's we have $n = 295$. We also replaced the highly right-skewed variables `price`, `displacement`, `horsepower`, `torque`, and `top speed` by their logarithms. On these data we ran `cellMCD` in its default version.

To visualize the results, we first look by variable. Consider variable j , say `horsepower`. Its i -th cell has observed value x_{ij} as well as its prediction \hat{x}_{ij} obtained from the *unflagged* cells in the same row i , as in (18). In (18) we also see the conditional variance C_{ij} of this




Figure 2: Top Gear data: (left) index plot of the standardized residual of `horsepower`; (right) standardized residual of `length` versus observed `length`.

cell. It is then natural to plot the *standardized cellwise residual*

$$\text{stdres}_{ij} = \frac{x_{ij} - \hat{x}_{ij}}{\sqrt{C_{ij}}} \quad (21)$$

which is NA when x_{ij} is missing. The left panel of Figure 2 shows the standardized residuals of the variable `horsepower` versus the index (case number) i . The function `plot.cellMCD()` shows the flagged points, i.e. those with $w_{ij} = 0$, in red. It also draws a horizontal tolerance band given by $\pm c$ where $c = \sqrt{\chi_{1,0.99}^2} \approx 2.57$. Here, some residuals stick out below the tolerance band. The Renault Twizy and Citroen DS3 are energy savers, whereas the Caterham is a super lightweight fun car. The most extreme outlier is the Chevrolet Volt with a standardized residual around -8 . Top Gear lists this car's power as 86 hp, which `cellMCD` says is very low compared to what would be expected from the other 10 characteristics of this car. Looking it up revealed that the Volt actually has 149 hp. As far as we know this data error was not detected before.

The right panel of Figure 2 plots the standardized residuals of the variable `length` versus the observed `length` itself. The vertical lines are at $T \pm cS$ where T and S are robust univariate location and scale estimates of `length`, obtained from the function `estLocScale()` in the R package *cellWise*. The points to the left and right of such a vertical tolerance band are marginally outlying, i.e. their `length` stands out by itself without regard to the other

variables. In the bottom left region of the plot we see five cars that are marginal outliers to the left and at the same time have outlying negative residuals, so they are short in absolute terms, as well as relative to what would be expected from their other characteristics. The Smart fortwo, Renault Twizy and Toyota IQ are indeed tiny. The Aston Martin Cygnet is merely a rebadged Toyota IQ, which enabled Aston Martin to comply with the 2012 European Union-imposed fleet average emissions regulations.

However, not all cellwise outliers are marginal outliers. In the middle bottom part of the plot we marked three cars whose length is not unusual by itself, but that are short relative to what would be expected based on their other 10 variables. They are sports cars, often built small to achieve high speeds. Note that there can also be points that lie inside the horizontal band but (slightly) outside the vertical band. They correspond to cells that look a bit unusual in the variable j , but whose observed value x_{ij} is not that far from the predicted \hat{x}_{ij} based on its other variables.




Figure 3: Top Gear data: (left) standardized residual of `weight` versus its prediction; (right) observed `top speed` versus its prediction.

The left panel of Figure 3 plots the standardized residual of each car’s `weight` versus its prediction. Since all the points lie within the vertical tolerance band, no predictions are outlying. But we do see some outlying residuals, most of which can easily be explained. The Bentley is a heavy luxury car, and the Mercedes-Benz G an all-terrain vehicle. Below

the horizontal tolerance band we see four lightweight sports cars. What remains is the Peugeot 107 which is small but not sporty at all. Top Gear reports its weight as 210 kg, which seems much too light for a car. Based on its other characteristics, cellMCD predicts its weight as 757 kg with a standard error of 89.5 kg. Looking up this car, its actual weight turns out to be 800 kg, so the value in the Top Gear dataset was mistaken.

The right panel of Figure 3 shows the observed value of **top speed** versus its prediction. Below the superimposed $y = x$ line we find some electric cars (BMW i3, Vauxhall Ampera) and some small cars (Smart fortwo and Renault Zoe). The one standing out most is the Renault Twizy, a tiny electric one-seater vehicle. Above the line we see some extremely fast sports cars. Also note that some points appear to lie on a horizontal line. Top Gear reports their top speed as 155 mph, corresponding to 250 km/hour. Many of these cars were produced by Audi, BMW and Mercedes with a built-in 250 km/hour speed limiter.

The four plot types in Figures 2 and 3 all focus on a single variable. It can also be instructive to look at a pair of variables, say j and k . The left panel of Figure 4 plots the variable **miles/gallon** versus **torque**. The points for which $w_{ij} = 0$ or $w_{ik} = 0$ or both are automatically plotted in red. The figure also contains an ellipse, given by

$$\begin{bmatrix} x - \hat{\mu}_j & y - \hat{\mu}_k \end{bmatrix} \begin{bmatrix} \hat{\Sigma}_{jj} & \hat{\Sigma}_{jk} \\ \hat{\Sigma}_{kj} & \hat{\Sigma}_{kk} \end{bmatrix}^{-1} \begin{bmatrix} x - \hat{\mu}_j \\ y - \hat{\mu}_k \end{bmatrix} = q \quad (22)$$

where q is the 0.99 quantile of the χ_2^2 distribution with two degrees of freedom. Note that outlyingness in this type of plot differs from cellwise outlyingness, since the former refers to two variables only, whereas the latter uses all 11 variables. So it is not unusual to see some red points inside the ellipse, and some black points outside it.

In the plot of **miles/gallon** versus **torque** we see two cases that stand out a lot, the BMW i3 and the Vauxhall Ampera. These are electric cars with a small gasoline engine to extend their range, explaining their huge **miles/gallon** values. For each of these cars, the red vertical line connects the observed point (x_{ij}, x_{ik}) to its imputed point $(\hat{x}_{ij}, \hat{x}_{ik})$ plotted in blue. That the line is vertical means that their **miles/gallon** cell was flagged whereas their **torque** cell was not, that is, $\hat{x}_{ik} \neq x_{ik}$ and $\hat{x}_{ij} = x_{ij}$. At the bottom right we see the Bugatti Veyron which has the highest **torque**. Only its **torque** cell was flagged, so the red line to the imputed point is horizontal here.

The right panel of Figure 4 shows the variables **width** and **acceleration**. The width of the Land Rover is flagged as this is a wide all terrain vehicle, and the width of the

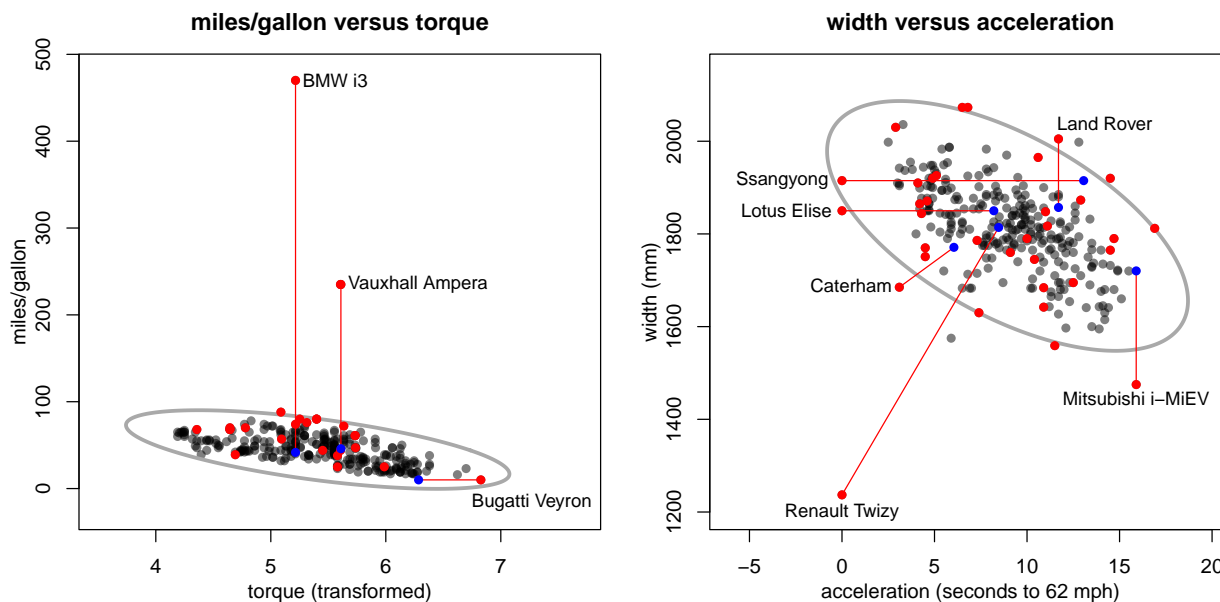


Figure 4: Top Gear data: bivariate plots of (left) miles/gallon versus torque, and (right) width versus acceleration. The 99% tolerance ellipse is given by the cellMCD estimates $\hat{\mu}$ and $\hat{\Sigma}$ restricted to the variables in the bivariate plot, and the red lines go to the imputed data shown in blue.

electrical Mitsubishi i-MiEV stands out in the opposite direction. The acceleration time of the Ssangyong Rodius and Lotus Elise is outlying on the left. In fact, Top Gear lists their acceleration time as 0 which is physically impossible: presumably the true value was missing and encoded as 0 instead of NA. The same is true for the Renault Twizy. Note that also the `width` cell of the Twizy is flagged, so the red line to its imputed point is slanted instead of horizontal. The Caterham also has both cells flagged, as seen from its slanted line.

We will now compare the results of the cellwise robust cellMCD method with those of the existing casewise MCD method on some benchmark datasets. For this we used all single class datasets with numerical variables in the `robustbase` package (Maechler et al., 2022) that have enough data points per dimension according to the rule of thumb $n/d > 5$. Many of these datasets came from Rousseeuw and Leroy (1987). To make their variables roughly gaussian in the center we first ran the function `transfo()` from the R package `cellwise` on them, which already checks whether no variables are discrete (as was the case for dataset `lactic`) and whether no variables have over 25% of marginal outliers (which eliminated datasets `pension` and `telef`). On the resulting transformed datasets we ran the

Table 1: Correlation between robust distances based on cellMCD and casewise MCD estimates

dataset	n	d	corr
aircraft (regressors only)	23	4	0.961
alcohol	44	7	0.861
animals2	65	2	0.992
bushfire	38	5	0.977
cloud	18	2	0.970
delivery (regressors only)	25	2	0.979
delivery (also response)	25	3	0.851
exAM	12	2	0.983
hbk (regressors only)	75	3	1.000
hbk (also response)	75	4	1.000
kootenay	13	2	1.000
milk	86	8	0.997
phosphor	18	3	0.969
pilot	20	2	0.893
radarImage	1573	5	1.000
salinity (regressors only)	28	3	0.849
salinity (also response)	28	4	0.887
starsCYG	47	2	0.998

casewise MCD by the function `covMcd()` in its default version, and we applied the default `cellMCD()` as well.

In order to compare the results it is convenient to compute the robust distances of the data points given by

$$RD_i = \sqrt{(\mathbf{x}_i - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}})} \quad (23)$$

where $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})$ are the casewise MCD estimates or the cellwise MCD estimates. We then computed the Pearson correlation between both sets of distances, which emphasizes the largest distances. Table 1 lists the datasets, their sample size n and dimension d , and the correlation. Many of these correlations turned out to be quite high. We conclude that, at least for these relatively small datasets with casewise outliers, cellMCD managed to detect

the main structure. Moreover, it provided information on which cells were responsible.

6 Simulation results

In this section we evaluate the performance of cellMCD by a simulation study. The clean data is generated as n points from a d -variate Gaussian distribution with mean $\boldsymbol{\mu} = \mathbf{0}$. Since there is no affine equivariance, letting $\boldsymbol{\Sigma}$ be the identity matrix is not sufficient. Instead we use the types ‘‘A09’’ and ‘‘ALYZ’’. The entries of the A09 correlation matrix are given by $\Sigma_{ij} = 0.9^{|i-j|}$, yielding both small and large correlations. The ALYZ type are randomly generated correlation matrices following the procedure of Agostinelli et al. (2015) and typically have mostly small absolute correlations. We consider three combinations of sample size and dimension (n, d) : (100, 10), (400, 20), and (800, 40).

In this clean data, we then replace a fraction ε in $\{0.1, 0.2\}$ of cells by contaminated cells. These are generated as follows. First, for each column in the data matrix we randomly sample $n\varepsilon$ indices of cells to be contaminated. Then we proceed in a rowwise fashion. For each row, say (z_1, \dots, z_d) , we collect the indices of the cells to be contaminated. Denote this set of size k by $K = \{j_1, \dots, j_k\}$. We next replace the cells $(z_{j_1}, \dots, z_{j_k})$ by the k -dimensional vector $\gamma\sqrt{k}\mathbf{u}/\text{MD}(\mathbf{u}, \boldsymbol{\mu}_K, \boldsymbol{\Sigma}_K)$ where $\boldsymbol{\mu}_K$ and $\boldsymbol{\Sigma}_K$ are $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ restricted to the indices in K . The scalar $\gamma > 0$ quantifies the distance of the outlying cells to the center of the distribution, and we vary γ over $1, \dots, 10$. The vector \mathbf{u} is the normed eigenvector of $\boldsymbol{\Sigma}_K$ with the smallest eigenvalue. In each row, the outlying cells are thus structurally outlying in the subspace generated by the variables in K . Therefore, these cells will often not be marginally outlying, especially when $|K|$ is large and γ is relatively small, which makes them hard to detect. The R-package `cellwise` (Raymaekers and Rousseeuw, 2022) contains the function `generateData` which generates the contaminated data according to this procedure.

We compare the proposed method cellMCD to the following alternative estimators:

- **Grank, Spearman**: the Gaussian and Spearman rank-based estimators used in Öllerer and Croux (2015) and Croux and Öllerer (2016);
- **GKnpd**: the Gnanadesikan-Kettenring estimator used in Tarr et al. (2016);
- **2SGS**: the two-step generalized S-estimator of Agostinelli et al. (2015);

- **DI**: the detection-imputation algorithm of Raymaekers and Rousseeuw (2021b).

In order to evaluate the performance of the different estimators, we compute the Kullback-Leibler discrepancy between the estimated $\hat{\Sigma}$ and the true Σ given by

$$\text{KL}(\hat{\Sigma}, \Sigma) = \text{tr}(\hat{\Sigma}\Sigma^{-1}) - d - \log(\det(\hat{\Sigma}\Sigma^{-1})) .$$

For each setting of the simulation parameters we generate 100 random datasets, and average the Kullback-Leibler discrepancy over these 100 replications.

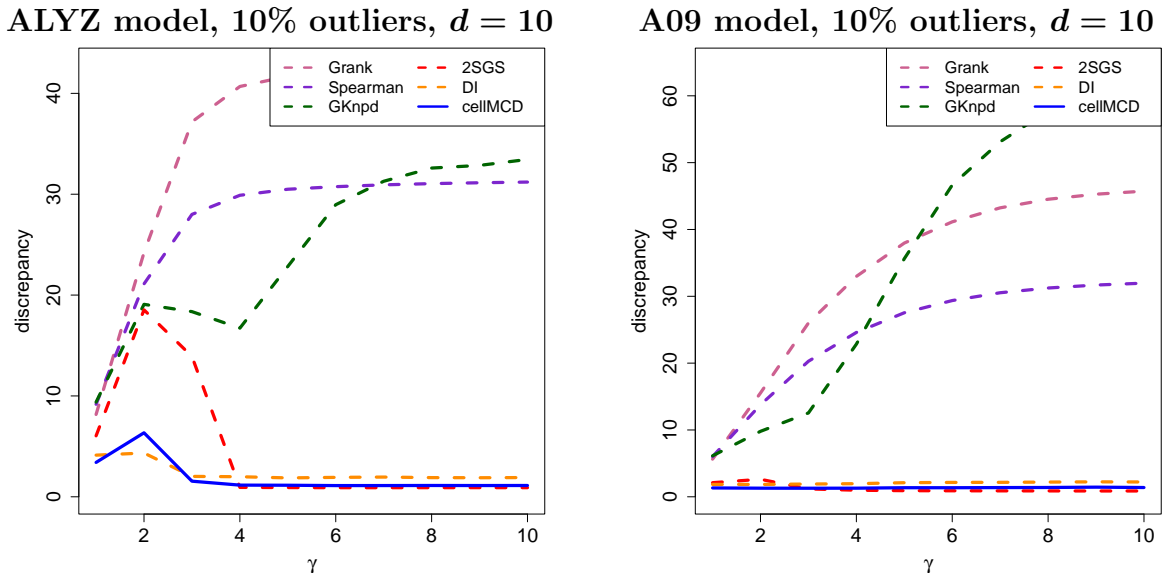



Figure 5: Discrepancy of estimated covariance matrices for $d = 10$ and $n = 100$.


We discuss the results for $\varepsilon = 0.1$. (The results for $\varepsilon = 0.2$ are similar and can be found in Section A.3 of the Supplementary Material.) Figure 5 presents the results for $n = 100$ and $d = 10$. Both cellMCD and DI perform well, as does 2SGS provided $\gamma \geq 4$. Note that the performances of Grank, Spearman and GKnpd do not improve as γ increases. While these estimators bound the influence that a single cell can have on the estimation, the effect remains substantial as the cell becomes more outlying. This is in contrast to 2SGS, DI and cellMCD in which far outliers get a zero weight.

The top panels of Figure 6 show the results for $n = 400$ and $d = 20$. The curves of the GKnpd, Spearman and Grank estimators came out much higher, and would squeeze the differences of the remaining methods when plotted, so we do not show them. The relative performances of the others are similar to Figure 5. The 2SGS method still does well when $\gamma > 4$, but now suffers more for low γ . The performances of DI and cellMCD are again very close, with cellMCD often doing slightly better.


ALYZ model, 10% outliers, $d = 20$



A09 model, 10% outliers, $d = 20$



ALYZ model, 10% outliers, $d = 40$



A09 model, 10% outliers, $d = 40$




Figure 6: Discrepancy of estimated covariance matrices for $d = 20$ and $n = 400$ (top panels) and for $d = 40$ and $n = 800$ (bottom panels).

The lower panels with $n = 800$ and $d = 40$ are similar, with cellMCD performing best for all values of γ while DI is quite close, and 2SGS only doing well for higher γ .

We are also interested in the performance of these methods on data without outliers. For this we repeated the simulation with $\varepsilon = 0$, again with 100 replications. The variability of each entry of the covariance matrix was measured taking the Fisher information of that entry into account. These results were then averaged over the d^2 matrix entries. Next we divided the MSE of the classical MLE estimator by that of each robust method, yielding the finite-sample efficiencies in Table 2.

Table 2: Finite-sample efficiencies of cellwise robust estimators

method	ALYZ configuration			A09 configuration		
	$d = 10$	$d = 20$	$d=40$	$d = 10$	$d = 20$	$d=40$
cellMCD	0.92	0.92	0.91	0.90	0.94	0.96
2SGS	0.88	0.96	0.99	0.83	0.92	0.95
DI	0.73	0.69	0.60	0.88	0.91	0.92
GKnpd	0.74	0.80	0.81	0.78	0.77	0.79
Grank	0.93	0.97	0.99	0.89	0.90	0.94
Spearman	0.85	0.89	0.90	0.83	0.82	0.85

We see that the efficiency of cellMCD averages over 90%, which is excellent for a highly robust covariance estimator. This is similar to 2SGS, and outperforms DI. As expected Grank has a high efficiency, but we just saw that the bottom three methods in the table perform poorly under contamination. Note that the finite-sample efficiency of cellMCD is much higher than that of the casewise MCD with the same coverage parameter $h = 0.75n$, which is under 0.70 for this range of dimensions d . This is due to the sparsity penalty in (9), which made the number of actually flagged cells much smaller than $0.25n$.

We conclude that cellMCD is about equally robust as DI but with better efficiency, and is about as efficient as 2SGS but with better robustness at contaminated data. Moreover, it is substantially better at contaminated data than the remaining three methods.

7 Discussion

The cellMCD method proposed here has an elegant formulation based on a single objective function, making it better understood than the earlier 2SGS and DI methods. We proved its good breakdown properties, and like the casewise MCD it can be computed by an algorithm based on C-steps that always lower the objective function, which is therefore guaranteed to converge. We have illustrated cellMCD on a real data set where the accompanying graphical displays revealed interesting aspects of the data that aided interpretation. Simulations indicate that cellMCD outperforms earlier cellwise methods, while being conceptually simple and rather fast to compute.

CellMCD is cellwise robust and incorporates a sparsity penalty. This naturally brings to mind the work of Candès et al. (2011). The goals are clearly related, but there are also some differences. The first is that their work assumes that the cellwise outlier pattern \mathbf{W} is drawn uniformly at random, whereas we adopt the robustness paradigm that the outliers may be placed adversarially. Secondly, the method of Candès et al. (2011) is equivariant for transposing the data matrix, so it treats cases and variables in the same way, whereas in our setting they have to be treated differently. We do allow for some rows being flagged entirely, whereas we cannot allow flagging an entire column as this would make $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ not identifiable, which motivates our constraint $\|\mathbf{W}_{.j}\|_0 \geq h$ for $j = 1, \dots, d$. This explains why we penalize and constrain the number of flagged cells by column, rather than penalizing the total number of flagged cells in the data matrix.

The fact that implosion breakdown can happen easily in the cellwise setting, see (13), was not mentioned in the literature before. We feel that, apart from cellMCD, also other cellwise robust covariance estimators could benefit from a constraint such as $\xi_d(\widehat{\boldsymbol{\Sigma}}) \geq a$, or equivalently from a formulation in which $\widehat{\boldsymbol{\Sigma}}$ is a convex combination of two matrices, one of which is the identity matrix with a small coefficient a .

The breakdown results in section 3 are for the exact cellMCD estimator, whereas the algorithm in section 4 yields an approximate solution. But the proof of the breakdown properties remains valid for the outcome of the algorithm, since the monotone decreasing property of the objective prevents it from increasing without bound. Other questions remain for future work, like consistency and asymptotic normality of cellMCD, which is hard since this is not yet fully settled for one of its components, the EM algorithm.

The casewise MCD is typically followed by a reweighting step. This works as follows. First, the estimated covariance matrix $\widehat{\boldsymbol{\Sigma}}$ is multiplied by a correction factor $c_{n,d,h}$ such that $c_{n,d,h}\widehat{\boldsymbol{\Sigma}}$ is roughly unbiased when the original data are generated from a Gaussian distribution. Next, the robust distances RD_i of the data points given by (23) are computed relative to $\widehat{\boldsymbol{\mu}}$ and $c_{n,d,h}\widehat{\boldsymbol{\Sigma}}$. Each case \mathbf{x}_i then gets a weight w_i depending on its RD_i . Typically, the weight is set to 1 when RD_i^2 is below some quantile of the χ_d^2 distribution with d degrees of freedom, and to 0 otherwise. The final estimates are then the weighted mean and the weighted covariance matrix (4). This reweighting step increases the finite-sample efficiency of the estimator.

For cellMCD, the analogous reweighting step would compute the standardized resid-

ual (21) of every cell x_{ij} and compare its square to a quantile of the χ_1^2 distribution with 1 degree of freedom, yielding zero-one weights w_{ij} . With these w_{ij} one would then run the EM algorithm on the original data. But in fact, the result is not very different from the cellMCD result. This is because all the ingredients are already used in cellMCD, which contains the squared standardized residual in (18), the χ_1^2 quantile in (20), and the partial likelihood on which EM is based in (9). So in some sense the components of a reweighting step are already built into cellMCD itself. This explains its rather high finite-sample efficiency in Table 2.

Software availability: The cellMCD method is implemented as the function `cellMCD()`, and the plots in section 5 were drawn by the function `plot_cellMCD()`. Both functions will be incorporated in the R package *cellWise* on CRAN. Its vignette `cellMCD_examples` reproduces all results and figures in section 5.

References

- Aerts, S. and I. Wilms (2017). Cellwise robust regularized discriminant analysis. *Statistical Analysis and Data Mining: The ASA Data Science Journal* 10(6), 436–447.
- Agostinelli, C., A. Leung, V. J. Yohai, and R. H. Zamar (2015). Robust estimation of multivariate location and scatter in the presence of cellwise and casewise contamination. *Test* 24, 441–461.
- Alfons, A. (2016). *robustHD: Robust methods for high-dimensional data*. R package version 0.5.1, CRAN.
- Alqallaf, F., S. Van Aelst, V. J. Yohai, and R. H. Zamar (2009). Propagation of outliers in multivariate data. *The Annals of Statistics* 37, 311–331.
- Boudt, K., P. J. Rousseeuw, S. Vanduffel, and T. Verdonck (2020). The minimum regularized covariance determinant estimator. *Statistics and Computing* 30, 113–128.
- Candès, E. J., X. Li, Y. Ma, and J. Wright (2011). Robust principal component analysis? *Journal of the ACM* 58(3), 1–37.

- Copt, S. and M.-P. Victoria-Feser (2004). Fast algorithms for computing high breakdown covariance matrices with missing data. In M. Hubert, G. Pison, A. Struyf, and S. Van Aelst (Eds.), *Theory and Applications of Recent Robust Methods*, Basel, pp. 71–82. Birkhäuser.
- Croux, C. and V. Öllerer (2016). Robust and sparse estimation of the inverse covariance matrix using rank correlation measures. In *Recent Advances in Robust Statistics: Theory and Applications*, pp. 35–55. Springer.
- Danilov, M., V. J. Yohai, and R. H. Zamar (2012). Robust estimation of multivariate location and scatter in the presence of missing data. *Journal of the American Statistical Association* 107, 1178–1186.
- De Ketelaere, B., M. Hubert, J. Raymaekers, P. J. Rousseeuw, and I. Vranckx (2020). Real-time outlier detection for large datasets by RT-DetMCD. *Chemometrics and Intelligent Laboratory Systems* 199, 103957.
- Debruyne, M., S. Höppner, S. Serneels, and T. Verdonck (2019). Outlyingness: Which variables contribute most? *Statistics and Computing* 29, 707–723.
- Dempster, A., N. Laird, and D. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B* 39(1), 1–22.
- Donoho, D. and P. Huber (1983). The notion of breakdown point. In P. Bickel, K. Doksum, and J. Hodges (Eds.), *A Festschrift for Erich Lehmann*, Belmont, pp. 157–184. Wadsworth.
- Filzmoser, P., S. Höppner, I. Ortner, S. Serneels, and T. Verdonck (2020). Cellwise robust M regression. *Computational Statistics & Data Analysis* 147, 106944.
- García-Escudero, L.-A., D. Rivera-García, A. Mayo-Isacar, and J. Ortega (2021). Cluster analysis with cellwise trimming and applications for the robust clustering of curves. *Information Sciences* 573, 100–124.
- Gnanadesikan, R. and J. Kettenring (1972). Robust estimates, residuals, and outlier detection with multiresponse data. *Biometrics* 28, 81–124.
- Higham, N. J. (2002). Computing the nearest correlation matrix – a problem from finance. *IMA Journal of Numerical Analysis* 22, 329–343.

- Hubert, M., M. Debruyne, and P. J. Rousseeuw (2018). Minimum Covariance Determinant and extensions. *Wiley Interdisciplinary Reviews: Computational Statistics* 10(3), e1421.
- Hubert, M., P. J. Rousseeuw, and P. Segaert (2015). Multivariate functional outlier detection. *Statistical Methods & Applications* 24, 177–202.
- Hubert, M., P. J. Rousseeuw, and W. Van den Bossche (2019). MacroPCA: An all-in-one PCA method allowing for missing values as well as cellwise and rowwise outliers. *Technometrics* 61(4), 459–473.
- Hubert, M., P. J. Rousseeuw, and T. Verdonck (2012). A deterministic algorithm for robust location and scatter. *Journal of Computational and Graphical Statistics* 21, 618–637.
- Katayama, S., H. Fujisawa, and M. Drton (2018). Robust and sparse gaussian graphical modelling under cell-wise contamination. *Stat* 7(1), e181.
- Leung, A., V. Yohai, and R. Zamar (2017). Multivariate location and scatter matrix estimation under cellwise and casewise contamination. *Computational Statistics & Data Analysis* 111, 59–76.
- Little, R. and D. Rubin (2020). *Statistical analysis with missing data (third edition)*. John Wiley and Sons, New York.
- Lopuhaä, H. P. and P. J. Rousseeuw (1991). Breakdown points of affine equivariant estimators of multivariate location and covariance matrices. *The Annals of Statistics* 19, 229–248.
- Maechler, M., P. Rousseeuw, C. Croux, V. Todorov, A. Rückstuhl, M. Salibian-Barrera, T. Verbeke, M. Koller, E. Conceicao, and M. di Palma (2022). *robustbase: Basic robust statistics*. R package, CRAN.
- Maronna, R. A. and V. J. Yohai (2008). Robust low-rank approximation of data matrices with elementwise contamination. *Technometrics* 50(3), 295–304.
- Öllerer, V., A. Alfons, and C. Croux (2016). The shooting S-estimator for robust regression. *Computational Statistics* 31(3), 829–844.

- Öllerer, V. and C. Croux (2015). Robust high-dimensional precision matrix estimation. In *Modern Nonparametric, Robust and Multivariate Methods*, eds. K. Nordhausen and S. Taskinen, pp. 325–350. Springer.
- Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. (2011). Scikit-learn: Machine Learning in Python. *The Journal of Machine Learning Research* 12, 2825–2830.
- Raymaekers, J. and P. J. Rousseeuw (2021a). Fast robust correlation for high-dimensional data. *Technometrics* 63(2), 184–198.
- Raymaekers, J. and P. J. Rousseeuw (2021b). Handling cellwise outliers by sparse regression and robust covariance. *Journal of Data Science, Statistics, and Visualization*, Volume 1, Number 3.
- Raymaekers, J. and P. J. Rousseeuw (2022). *cellWise: Analyzing Data with Cellwise Outliers*. R package, CRAN.
- Rousseeuw, P. J. (1984). Least median of squares regression. *Journal of the American Statistical Association* 79, 871–880.
- Rousseeuw, P. J. (1985). Multivariate estimation with high breakdown point. In W. Grossmann, G. Pflug, I. Vincze, and W. Wertz (Eds.), *Mathematical Statistics and Applications*, pp. 283–297. Reidel.
- Rousseeuw, P. J. and C. Croux (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association* 88, 1273–1283.
- Rousseeuw, P. J. and A. Leroy (1987). *Robust Regression and Outlier Detection*. New York: John Wiley.
- Rousseeuw, P. J. and W. Van den Bossche (2018). Detecting Deviating Data Cells. *Technometrics* 60, 135–145.
- Rousseeuw, P. J. and K. Van Driessen (1999). A Fast Algorithm for the Minimum Covariance Determinant Estimator. *Technometrics* 41, 212–223.

- Rousseeuw, P. J. and B. C. van Zomeren (1990). Unmasking multivariate outliers and leverage points. *Journal of the American Statistical Association* 85, 633–651.
- Schreurs, J., I. Vranckx, M. Hubert, J. Suykens, and P. J. Rousseeuw (2021). Outlier detection in non-elliptical data by kernel MRCD. *Statistics and Computing* 31:66, 1–18.
- She, Y., Z. Wang, and J. Shen (2022). Gaining outlier resistance with progressive quantiles: Fast algorithms and theoretical studies. *Journal of the American Statistical Association*, in press.
- Štefelová, N., A. Alfons, J. Palarea-Albaladejo, P. Filzmoser, and K. Hron (2021). Robust regression with compositional covariates including cellwise outliers. *Advances in Data Analysis and Classification* 15, 1–41.
- Su, P., G. Tarr, and S. Muller (2021). Robust variable selection under cellwise contamination, *arXiv preprint 2110.12406*.
- Tarr, G., S. Muller, and N. Weber (2016). Robust estimation of precision matrices under cellwise contamination. *Computational Statistics & Data Analysis* 93, 404–420.
- Todorov, V. (2012). *rrcov: Scalable Robust Estimators with High Breakdown Point*. R package, CRAN.
- Unwin, A. (2019). Multivariate outliers and the O3 plot. *Journal of Computational and Graphical Statistics* 28, 635–643.
- Van Aelst, S., E. Vandervieren, and G. Willems (2011). Stahel-Donoho estimators with cellwise weights. *Journal of Statistical Computation and Simulation* 81(1), 1–27.

Supplementary Material to:

The Cellwise Minimum Covariance Determinant Estimator

Jakob Raymaekers and Peter J. Rousseeuw

A.1 Proofs of the propositions

Proof of Proposition 1. The proof consists of four parts.

Part (a): this follows immediately from the constraint $\xi_d(\widehat{\Sigma}) \geq a$ for $a > 0$.

Part (b): Explosion breakdown of $\widehat{\Sigma}$.

Denote by \mathcal{X}_m the set of all corrupted samples \mathbf{X}^m obtained by replacing at most m cells in each column of \mathbf{X} by arbitrary values, for $m = n - h$. Also denote

$$\mathcal{W}_h = \{\mathbf{W} \in \{0, 1\}^{n \times d} \mid \|\mathbf{W}_{.j}\|_0 \geq h \text{ for all } j = 1, \dots, d\}.$$

Then we can write

$$\mathcal{X}_m = \bigcup_{\mathbf{W}^* \in \mathcal{W}_h} \{\mathbf{X}^m \in \mathcal{X}_m \mid \mathbf{W}_{ij}^* = 1 \Rightarrow \mathbf{X}_{ij}^m = \mathbf{X}_{ij}\}.$$

In other words, we can write the set of all corrupted samples \mathcal{X}_m as a finite union over subsets of corrupted samples with the same contaminating configuration \mathbf{W}^* .

We start by showing the existence of a solution with finite objective function. Consider any such contaminating configuration $\mathbf{W}^* \in \mathcal{W}_h$. Then take the solution $(\widehat{\boldsymbol{\mu}}_{\text{EM}}, \widehat{\boldsymbol{\Sigma}}_{\text{EM}}, \mathbf{W}^*)$ where the location and scatter are the result of the EM-algorithm with fixed missingness pattern given by \mathbf{W}^* . Then

$$\begin{aligned} & \forall \mathbf{X}^m \in \{\mathbf{X}^m \in \mathcal{X}_m \mid \mathbf{W}_{ij}^* = 1 \Rightarrow \mathbf{X}_{ij}^m = \mathbf{X}_{ij}\}: \\ & \text{Obj}(\widehat{\boldsymbol{\mu}}_{\text{EM}}(\mathbf{X}^m), \widehat{\boldsymbol{\Sigma}}_{\text{EM}}(\mathbf{X}^m), \mathbf{W}^*) = \text{Obj}(\widehat{\boldsymbol{\mu}}_{\text{EM}}(\mathbf{X}), \widehat{\boldsymbol{\Sigma}}_{\text{EM}}(\mathbf{X}), \mathbf{W}^*) = M_{\mathbf{W}^*} < \infty \end{aligned}$$

in which $\text{Obj}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{W})$ denotes the objective function (9) of cellMCD. In other words, for all \mathbf{X}^m with the same contaminating configuration, we have a candidate solution with a finite objective function. Since there are finitely many such contaminating distributions, we can always find a candidate solution with a value of the objective function smaller than $M = \max\{M_{\mathbf{W}^*} : \mathbf{W}^* \in \mathcal{W}_h\} < \infty$.

We now show that $\widehat{\Sigma}$ does not explode. By construction, $\xi_d(\widehat{\Sigma}) \geq a$ for some constant $a > 0$. Then we have that

$$\begin{aligned}
\ln |\widehat{\Sigma}^{(\mathbf{w}_i)}| &= \sum_{j=1}^{d(\mathbf{w}_i)} \ln \xi_j(|\widehat{\Sigma}^{(\mathbf{w}_i)}|) \\
&= \ln \xi_1(|\widehat{\Sigma}^{(\mathbf{w}_i)}|) + \sum_{j=2}^{d(\mathbf{w}_i)} \ln \xi_j(|\widehat{\Sigma}^{(\mathbf{w}_i)}|) \\
&\geq \ln \xi_1(|\widehat{\Sigma}^{(\mathbf{w}_i)}|) + \sum_{j=2}^{d(\mathbf{w}_i)} \ln \xi_{d(\mathbf{w}_i)}(|\widehat{\Sigma}^{(\mathbf{w}_i)}|) \\
&\geq \ln \max_j \widehat{\Sigma}_{jj}^{(\mathbf{w}_i)} + \sum_{j=2}^{d(\mathbf{w}_i)} \ln(\xi_d(\widehat{\Sigma})) \\
&\geq \ln \max_j \widehat{\Sigma}_{jj}^{(\mathbf{w}_i)} + (d-1) \ln a
\end{aligned}$$

where we have used that $\xi_1(\widehat{\Sigma}^{(\mathbf{w})}) \geq \max_j \widehat{\Sigma}_{jj}^{(\mathbf{w})}$ for any \mathbf{w} . That is, the largest eigenvalue of any positive semi-definite (sub)matrix is at least as large as its largest diagonal element.

Now we can bound the first term of the objective from below by an increasing function of the largest eigenvalue. First note that we have at least one row i^* for which the j^* -th element of \mathbf{w}_{i^*} is 1, where $j^* = \operatorname{argmax}_j \widehat{\Sigma}_{jj}$. Therefore

$$\begin{aligned}
\sum_{i=1}^n \ln |\widehat{\Sigma}^{(\mathbf{w}_i)}| &= \ln |\widehat{\Sigma}^{(\mathbf{w}_{i^*})}| + \sum_{i \neq i^*} \ln |\widehat{\Sigma}^{(\mathbf{w}_i)}| \\
&\geq \ln \max_j \widehat{\Sigma}_{jj}^{\mathbf{w}_{i^*}} + (d-1) \ln a + (n-1)d \ln a \\
&= \ln \max_j \widehat{\Sigma}_{jj} + (nd-1) \ln a \\
&= \ln \max_{jk} |\widehat{\Sigma}_{jk}| + (nd-1) \ln a \\
&\geq \ln \frac{\xi_1(\widehat{\Sigma})}{d} + (nd-1) \ln a
\end{aligned}$$

where we have used that $\xi_1(\widehat{\Sigma}) \leq d \max_{jk} |\widehat{\Sigma}_{jk}|$, i.e. the largest eigenvalue of a $d \times d$ positive definite matrix is at most d times its largest absolute entry. Also, we have used that $\max_{jk} |\widehat{\Sigma}_{jk}| = \max_j |\widehat{\Sigma}_{jj}|$ since $\widehat{\Sigma}$ is a covariance matrix, so its maximum occurs on the diagonal.

As all other terms of the objective function are bounded from below by zero, we obtain:

$$\begin{aligned} \text{Obj}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \mathbf{W}) &= \sum_{i=1}^n (\ln |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| + d(\mathbf{w}_i) \ln(2\pi) + \text{MD}^2(\mathbf{x}_i^m, \mathbf{w}_i, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})) + \sum_{j=1}^d \xi_j \|\mathbf{1}_d - \mathbf{W}_{.j}\|_0 \\ &\geq \sum_{i=1}^n \ln |\hat{\boldsymbol{\Sigma}}^{(\mathbf{w}_i)}| \geq \ln \frac{\xi_1(\hat{\boldsymbol{\Sigma}})}{d} + (nd - 1) \ln a. \end{aligned}$$

We thus find that the objective function explodes when $\xi_1(\hat{\boldsymbol{\Sigma}}) \rightarrow \infty$. Given that for any possible contaminated dataset there is a candidate solution with objective function less than or equal to $M < \infty$, we conclude that the solution cannot have an exploding eigenvalue.

Part (c): Breakdown of $\hat{\boldsymbol{\mu}}$.

Note that for all $\mathbf{W} \in \mathcal{W}_h$ and each variable j there is at least one $\mathbf{W}_{ij} = 1$, so a cell \mathbf{X}_{ij} that was not replaced. Denote $M_2 := \max_{ij} |\mathbf{X}_{ij}| < \infty$. Then we have

$$\begin{aligned} \text{Obj}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \mathbf{W}) &= \sum_{i=1}^n (\ln |\boldsymbol{\Sigma}^{(\mathbf{w}_i)}| + d(\mathbf{w}_i) \ln(2\pi) + \text{MD}^2(\mathbf{x}_i^m, \mathbf{w}_i, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})) + \sum_{j=1}^d \xi_j \|\mathbf{1}_d - \mathbf{W}_{.j}\|_0 \\ &\geq nd \ln a + \sum_{i=1}^n \text{MD}^2(\mathbf{x}_i^m, \mathbf{w}_i, \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}) \\ &= nd \ln a + \sum_{i=1}^n \left\| \left(\hat{\boldsymbol{\Sigma}}^{(\mathbf{w}_i)} \right)^{-1/2} (\mathbf{x}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}) \right\|_2^2 \\ &\geq nd \ln a + \sum_{i=1}^n \xi_{\min}^2 \left(\left(\hat{\boldsymbol{\Sigma}}^{(\mathbf{w}_i)} \right)^{-1/2} \right) \|\mathbf{x}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}\|_2^2 \\ &= nd \ln a + \sum_{i=1}^n \frac{1}{\xi_{\max}(\hat{\boldsymbol{\Sigma}}^{(\mathbf{w}_i)})} \|\mathbf{x}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}\|_2^2 \\ &\geq nd \ln a + \frac{1}{\xi_{\max}(\hat{\boldsymbol{\Sigma}})} \sum_{i=1}^n \|\mathbf{x}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}\|_2^2 \\ &\geq nd \ln a + \frac{1}{\xi_{\max}(\hat{\boldsymbol{\Sigma}})} (\|\hat{\boldsymbol{\mu}}\|_2^2 - dM_2^2). \end{aligned}$$

In the last line we have used that there is at least one uncontaminated cell in each variable for which $\mathbf{W}_{ij} = 1$, together with the fact that this cell is bounded in absolute value by M_2 . From part (b) we don't have explosion of the covariance matrix, so $\xi_{\max}(\hat{\boldsymbol{\Sigma}}) < \infty$. Should $\|\hat{\boldsymbol{\mu}}\|_2 \rightarrow \infty$ our objective function would explode, but we know it does not.

Part (d): The bound $(n - h + 1)/n$ is sharp. So far we know that $\varepsilon_n^*(\hat{\boldsymbol{\mu}}, \mathbf{X}) \geq (n - h + 1)/n$ and $\varepsilon_n^+(\hat{\boldsymbol{\Sigma}}, \mathbf{X}) \geq (n - h + 1)/n$. We now show that this common lower bound cannot be improved, by constructing an example which causes breakdown. For this we take a contaminating configuration obtained by replacing $n - h + 1$ cells in the first column of the data \mathbf{X} by some value c and leaving all other columns untouched. Unlike before, there is no way to cover all these cells with any $W \in \mathcal{W}_h$. Put $M_2 = \max_{ij} |\mathbf{X}_{ij}|$ as before.

Consider any solution $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \mathbf{W})$ with $\mathbf{W} \in \mathcal{W}_h$. Denote by \mathcal{I} the set of indices of the rows which have a contaminated cell equal to c in their first variable. Denote by subscript o_i the set of variables j for which $\mathbf{w}_{ij} = 1$. By the first order conditions of the EM algorithm, upon convergence of the algorithm we must have $\hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i$ where \mathbf{y}_i are the imputed observations. For the first entry of $\hat{\boldsymbol{\mu}}$ we have:

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_1 &= \frac{1}{n} \sum_{i=1}^n \mathbf{y}_{i1} \\
&= \frac{1}{n} \sum_{i|\mathbf{w}_{i1}=1} \mathbf{X}_{i1}^m + \frac{1}{n} \sum_{i|\mathbf{w}_{i1}=0} E[\mathbf{X}_{i1} | \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \mathbf{W}] \\
&= \frac{1}{n} \sum_{i|\mathbf{w}_{i1}=1} \mathbf{X}_{i1}^m + \frac{1}{n} \sum_{i|\mathbf{w}_{i1}=0} \left(\hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\Sigma}}_{1,o_i} \hat{\boldsymbol{\Sigma}}_{o_i,o_i}^{-1} (\mathbf{X}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}) \right) \\
&= \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}} \mathbf{X}_{i1}^m + \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}^c} \mathbf{X}_{i1}^m + \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=0\}} \left(\hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\Sigma}}_{1,o_i} \hat{\boldsymbol{\Sigma}}_{o_i,o_i}^{-1} (\mathbf{X}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}) \right) \\
&= \frac{c}{n} \#(\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}) + \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}^c} \mathbf{X}_{i1}^m + \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=0\}} \left(\hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\Sigma}}_{1,o_i} \hat{\boldsymbol{\Sigma}}_{o_i,o_i}^{-1} (\mathbf{X}_{i,o_i}^m - \hat{\boldsymbol{\mu}}_{o_i}) \right) \\
&= \frac{c}{n} \#(\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}) + \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}^c} \mathbf{X}_{i1} + \frac{1}{n} \sum_{\{i|\mathbf{w}_{i1}=0\}} \left(\hat{\boldsymbol{\mu}}_1 + \hat{\boldsymbol{\Sigma}}_{1,o_i} \hat{\boldsymbol{\Sigma}}_{o_i,o_i}^{-1} (\mathbf{X}_{i,o_i} - \hat{\boldsymbol{\mu}}_{o_i}) \right).
\end{aligned}$$

Note that we have replaced \mathbf{X}^m by \mathbf{X} in the last line, since all those cells are uncontaminated. By construction of our contaminated data, we have $\#(\{i|\mathbf{w}_{i1}=1\} \cap \mathcal{I}) \geq 1$. Now take a sequence c_k which diverges, i.e. $c_k \rightarrow \infty$ as $k \rightarrow \infty$. Suppose that our estimates $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$ would not break down as $k \rightarrow \infty$. Then the $\hat{\boldsymbol{\mu}}_1$ on the left hand side of the above equality would be bounded. The second term on the right hand side is just an average of uncontaminated data so it is bounded too. The last term on the right hand side would be bounded as well, since it consists of the estimated $\hat{\boldsymbol{\mu}}$, $\hat{\boldsymbol{\Sigma}}$, and the uncontaminated data. (Note that $\hat{\boldsymbol{\Sigma}}_{o_i,o_i}^{-1}$ is bounded since $\xi_1(\hat{\boldsymbol{\Sigma}}_{o_i,o_i}^{-1}) \leq 1/a < \infty$.) However, the first term on the right hand side would diverge. This is a contradiction. We conclude that either the location or the covariance matrix (or both) must diverge as $k \rightarrow \infty$. \square

Proof of Proposition 2. Put $\boldsymbol{\mu} = \mathbf{0}$ without loss of generality. Following Petersen and Pedersen (2012), p. 47, we can write

$$\boldsymbol{\Sigma}^{-1} = \mathbf{A}\mathbf{B}\mathbf{A}^\top$$

with

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\boldsymbol{\Sigma}_{22}^{-1}\boldsymbol{\Sigma}_{21} & \mathbf{I} \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} \mathbf{C}_1^{-1} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_{22}^{-1} \end{bmatrix}.$$

Note that

$$\mathbf{x}^\top \mathbf{A} = \begin{bmatrix} \mathbf{x}_1^\top - \mathbf{x}_2^\top \boldsymbol{\Sigma}_{22}^{-1} \boldsymbol{\Sigma}_{21} & \mathbf{x}_2^\top \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1^\top - \widehat{\mathbf{x}}_1^\top & \mathbf{x}_2^\top \end{bmatrix}$$

and so

$$\begin{aligned} \text{MD}^2(\mathbf{x}, \mathbf{0}, \boldsymbol{\Sigma}) &= \mathbf{x}^\top \boldsymbol{\Sigma}^{-1} \mathbf{x} = (\mathbf{x}^\top \mathbf{A}) \mathbf{B} (\mathbf{x}^\top \mathbf{A})^\top \\ &= (\mathbf{x}_1 - \widehat{\mathbf{x}}_1)^\top \mathbf{C}_1^{-1} (\mathbf{x}_1 - \widehat{\mathbf{x}}_1) + \mathbf{x}_2^\top \boldsymbol{\Sigma}_{22}^{-1} \mathbf{x}_2 \\ &= \text{MD}^2(\mathbf{x}_1, \widehat{\mathbf{x}}_1, \mathbf{C}_1) + \text{MD}^2(\mathbf{x}_2, \mathbf{0}, \boldsymbol{\Sigma}_{22}). \end{aligned}$$

For (15), we verify that

$$|\boldsymbol{\Sigma}^{-1}| = |\mathbf{A}| |\mathbf{B}| |\mathbf{A}| = 1 |\mathbf{C}_1^{-1}| |\boldsymbol{\Sigma}_{22}^{-1}| 1$$

so

$$|\boldsymbol{\Sigma}| = |\mathbf{C}_1| |\boldsymbol{\Sigma}_{22}|.$$

Finally,

$$\begin{aligned} &L(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) - L(\mathbf{x}_2, \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{22}) \\ &= \text{MD}^2(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) - \text{MD}^2(\mathbf{x}_2, \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_{22}) + d \ln(2\pi) - d(\mathbf{x}_2) \ln(2\pi) + \ln |\boldsymbol{\Sigma}| - \ln |\boldsymbol{\Sigma}_{22}| \\ &= \text{MD}^2(\mathbf{x}_1, \widehat{\mathbf{x}}_1, \mathbf{C}_1) + d(\mathbf{x}_1) \ln(2\pi) + \ln |\mathbf{C}_1| = L(\mathbf{x}_1, \widehat{\mathbf{x}}_1, \mathbf{C}_1). \end{aligned}$$

□

Proof of Proposition 3. Part (a) of the C-step repeatedly updates one column of $\widetilde{\mathbf{W}}$, say column j . It sets $\widetilde{\mathbf{w}}_{ij} = 1$ for all i with negative Δ_{ij} . If that number exceeds h the constraint is satisfied, and otherwise it takes the h smallest values of Δ_{ij} . In either case we obtain the lowest sum of the terms of the objective (9) in column j that satisfies the constraint, so that sum has to be less than or equal to before. This remains true after repeating the procedure on other columns.

Part (b) starts by performing the standard E-step. Next, the M-step is carried out and the constraint $\xi_d(\widehat{\Sigma}) \geq a$ is applied by truncating all eigenvalues of $\widehat{\Sigma}$ at a from below. This combination nevertheless reduces the objective (7) or keeps it the same, following section 11.3 of Little and Rubin (2020) on the Gaussian model with a restricted covariance matrix. This is because the E-step is unchanged, whereas the constraint acts on the M-step which is the same as if the result of the E-step came from complete data. For our specific constraint this was also shown in Proposition 1 of Aubry et al. (2021), see in particular their formulas (33) and (34). Since the objective (7) is reduced or stays the same, this also follows for the total objective (9). \square

A.2 The initial estimator DDCW

The C-step iterations of section 4 need initial cellwise robust estimates $\widehat{\mu}^0$ and $\widehat{\Sigma}^0$ of location and covariance. For this purpose we developed an initial estimator called DDCW, described here. Its steps are:

1. Drop variables with too many missing values or zero median absolute deviation, and continue with the remaining columns.
2. Run the DetectDeviatingCells (DDC) method (Rousseeuw and Van den Bossche, 2018) with the constraint that no more than $n - h$ cells are flagged in any variable. DDC also rescales the variables, and may delete some cases. Continue with the remaining imputed and rescaled cases denoted as \mathbf{z}_i .
3. Project the \mathbf{z}_i on the axes of their principal components, yielding the transformed data points $\tilde{\mathbf{z}}_i$.
4. Compute the wrapped location $\widehat{\mu}_w$ and covariance matrix $\widehat{\Sigma}_w$ (Raymaekers and Rousseeuw, 2021a) of these $\tilde{\mathbf{z}}_i$. Next, compute the temporary points $\mathbf{u}_i = (u_{i1}, \dots, u_{id})$ given by $u_{ij} = \max(\min(\tilde{z}_{ij} - (\widehat{\mu}_w)_j, 2), -2)$. Then remove all cases for which the squared robust distance $\text{RD}^2(i) = \mathbf{u}_i' \widehat{\Sigma}_w^{-1} \mathbf{u}_i$ exceeds $\chi_{d,0.99}^2 \text{median}_h(\text{RD}^2(h)) / \chi_{d,0.5}^2$.
5. Project the remaining $\tilde{\mathbf{z}}_i$ on the eigenvectors of $\widehat{\Sigma}_w$ and again compute a wrapped location and covariance matrix.

6. Transform these estimates back to the original coordinate system of the imputed data, and undo the scaling. This yields the estimates $\hat{\boldsymbol{\mu}}^0$ and $\hat{\boldsymbol{\Sigma}}^0$.

Note that DDCW can handle missing values since the DDC method in Step 2 imputes them. The reason for the truncation in the rejection rule in Step 4 is that otherwise the robust distance RD could be inflated by a single outlying cell. Step 4 tends to remove rows which deviate strongly from the covariance structure. These are typically rows which cannot be shifted towards the majority of the data without changing a large number of cells.

Also note that instead of starting from a single initial estimate $(\hat{\boldsymbol{\mu}}^0, \hat{\boldsymbol{\Sigma}}^0)$, one could start from several initial estimates. Iterating C-steps from each (with the same λ_j and $a > 0$) until convergence, one can then keep the solution with the lowest objective (9).


A.3 More simulation results

Figure 7 shows the simulation results for $\varepsilon = 0.2$ in the same settings as in section 6. The results are qualitatively similar.

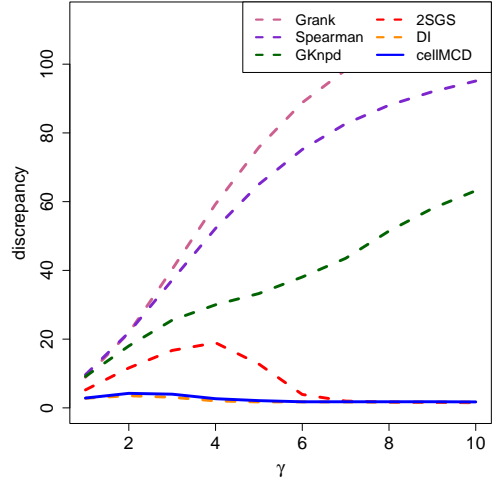
References

- Aubry, A., A. De Maio, S. Marano, and M. Rosamilia (2021). Structured covariance matrix estimation with missing-(complex) data for radar applications via expectation-maximization. *IEEE Transactions on Signal Processing* 69, 5920–5934.
- Petersen, K. B. and M. S. Pedersen (2012). *The Matrix Cookbook*. Technical University of Denmark.


ALYZ model, 20% outliers, $d = 10$




A09 model, 20% outliers, $d = 10$




ALYZ model, 20% outliers, $d = 20$



A09 model, 20% outliers, $d = 20$



ALYZ model, 20% outliers, $d = 40$



A09 model, 20% outliers, $d = 40$




Figure 7: Discrepancy of covariance matrices for $d = 10$ and $n = 100$ (top panels), $d = 20$ and $n = 400$ (middle panels) and for $d = 40$ and $n = 800$ (bottom panels).