Real-time discriminant analysis in the presence of label and measurement noise

Iwein Vranckx, Jakob Raymaekers, Bart De Ketelaere,
Peter J. Rousseeuw, Mia Hubert

KU Leuven, BE-3001 Heverlee, Belgium

November 10, 2020

Abstract

Quadratic discriminant analysis (QDA) is a widely used classification technique. Based on a training dataset, each class in the data is characterized by an estimate of its center and shape, which can then be used to assign unseen observations to one of the classes. The traditional QDA rule relies on the empirical mean and covariance matrix. Unfortunately, these estimators are sensitive to label and measurement noise which often impairs the model's predictive ability. Robust estimators of location and scatter are resistant to this type of contamination. However, they have a prohibitive computational cost for large scale industrial experiments. We present a novel QDA method based on a recent real-time robust algorithm. We additionally integrate an anomaly detection step to classify the most atypical observations into a separate class of outliers. Finally, we introduce the label bias plot, a graphical display to identify label and measurement noise in the training data. The performance of the proposed approach is illustrated in a simulation study with huge datasets, and on real datasets about diabetes and fruit.

Keywords: Label bias, Minimum covariance determinant, Mislabeling, Outliers, Robust classification.

1 Introduction

Supervised classification is a very common task in statistics and machine learning. Given a training dataset of labeled instances, the goal is to train a classifier such that new observations can be classified into one of the known

classes (groups). Many classification techniques exist, see e.g. Hastie et al. (2009); Murphy (2012); James et al. (2013) for an overview. We will focus on discriminant analysis (DA), one of the oldest and well-studied techniques. DA is based on the underlying assumption that the data follow a mixture of multivariate normal distributions. In its basic form DA has several attractive properties, most notably its conceptual and computational simplicity.

Traditional DA relies on the empirical mean and covariance matrix of each class. Despite its nice properties, it is highly sensitive to violations of the mixture density assumption. Two sources of these violations which commonly occur in practice are *label noise* and *measurement noise*. Label noise, also called mislabeling, occurs when instances in the training data have been given a wrong label, so their recorded label differs from their actual one. See Frénay and Verleysen (2013) for a comprehensive survey. These instances can heavily affect the classification result, since they essentially encourage the classifier to associate characteristics of one class with the label of another class. As training data is often labeled manually, the labels are prone to human error and some degree of mislabeling is likely to occur in practice.

The second type of noise, measurement noise, occurs when observations in the training data have deviating measurements. Such outliers may affect the mean and covariance matrix of their class which is then characterized poorly, causing classical DA to underperform.

Several proposals have been made to make DA robust against label and measurement noise, see for example Chork and Rousseeuw (1992); Croux and Dehon (2001); Hubert and Van Driessen (2004). These methods all rely on robust estimators of location and scatter, which work well but need substantial computation time for large datasets. The context of this paper is that of real-time classification in an industrial setting, such as food sorting or classification of plastics and glass. Typically, huge amounts of product are scanned in an automated inspection process. Robust discriminant analysis is an absolute must in this setting, since these datasets are typically corrupted by both label and measurement noise. Unfortunately, none of the previously mentioned robust algorithms can handle the sheer volume of data that is generated by these classification tasks.

In this paper we address this issue by incorporating the recently introduced RT-DetMCD method (De Ketelaere et al., 2020), a real-time robust estimator of location and scatter, into the discriminant analysis framework. We further integrate an anomaly detection step: if an observation does not match any of the known classes it will be classified into a separate outlier class, explicitly revealing significant measurement errors in the training set. The resulting approach allows us to combine a high degree of robustness against label and measurement noise with a low computation time. We also introduce a graphical display to identify label and measurement noise in the training data. Using extensive simulations we show that the proposed ap-

proach works well at huge datasets, even with high noise rates. The accuracy of the method is also illustrated on two real datasets, in which we identify and interpret several sets of atypical observations.

The remainder of the paper is organized as follows. In Section 2 we describe our real-time robust classifier, which incorporates the anomaly detection step. Section 3 introduces the new graphical display. The simulation study in Section 4 compares the performance of the proposed method with that of classical discriminant analysis under label and measurement noise. Section 5 illustrates the method on two real datasets. Finally, the main conclusions are summarized in Section 6.

2 Real-time robust QDA

2.1 Discriminant analysis

Suppose we have a p -variate random vector X which describes the data generated by an experiment. Assume X follows a multivariate normal mixture model with G classes (subpopulations), i.e. the density of X can be written as $f(\mathbf{x}) = \sum_{g=1}^G p_g f_g(\mathbf{x})$ where p_g denotes the prior probability of class $g = 1, \dots, G$ and $f_g \sim N(\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g)$ is the p -variate normal density of class g described by the location vector $\boldsymbol{\mu}_g$ and scatter matrix $\boldsymbol{\Sigma}_g$. The aim is to divide the p -dimensional space into G regions which correspond to the classes. Based on these regions, new cases can be classified into one of the classes. To find these regions DA uses the Bayes discriminant rule, which assigns an unknown observation $\mathbf{x}$ to the class g for which $\ln(p_g f_g(\mathbf{x}))$ is highest among all G classes. Using the density of the multivariate normal distribution one obtains the quadratic discriminant analysis (QDA) rule: assign $\mathbf{x}$ to the class g for which the quadratic discriminant score $d(\mathbf{x}, \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, p_g)$ is highest, with

$$d(\mathbf{x}, \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, p_g) = -\frac{1}{2} \ln |\boldsymbol{\Sigma}_g| - \frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_g)^\top \boldsymbol{\Sigma}_g^{-1} (\mathbf{x} - \boldsymbol{\mu}_g) + \ln(p_g). \quad (1)$$

If the covariance matrices $\boldsymbol{\Sigma}_g$ of all classes are equal they can be replaced by a common covariance matrix $\boldsymbol{\Sigma}$, which leads to linear discriminant scores and the corresponding linear discriminant analysis (LDA) method. However, as this homoskedasticity assumption is often not realistic in practical settings, we will concentrate on QDA.

2.2 Classical discriminant analysis

The quadratic discriminant scores (1) are computed based on the prior probabilities p_g , the means $\boldsymbol{\mu}_g$ and the covariance matrices $\boldsymbol{\Sigma}_g$, which all have to be estimated from the data. Suppose we have a multivariate dataset $\mathbf{X}$ of n observations in p dimensions sampled from G different classes, and a class label vector $\mathbf{y}$ of length n with y_i in $\{1, \dots, G\}$ for all $i = 1, \dots, n$. Denote

by n_g the number of observations from class g in the data. The set of observations and labels $(\mathbf{X}, \mathbf{y})$ is called the training set. In order to estimate the unknown parameters $\boldsymbol{\mu}_g$ and $\boldsymbol{\Sigma}_g$, classical QDA (CQDA) uses the empirical mean $\bar{\mathbf{x}}_g$ and empirical covariance matrix $\mathbf{S}_g$ of each class. The unknown membership probabilities are estimated using the relative frequencies of each class in the training data: $\hat{p}_{g,C} = n_g/n$ where C stands for classical. The CQDA rule then assigns $\mathbf{x}$ to the class g for which $d(\mathbf{x}, \bar{\mathbf{x}}_g, \mathbf{S}_g, \hat{p}_{g,C})$ is highest. Note that the data dimension p should be below $\min_g n_g$ since otherwise $\mathbf{S}_g$ becomes singular.

Although CQDA remains a popular classification method, it is known that it is very sensitive to mislabeling and outliers, as it is based on classical estimators of location and scatter. To illustrate this, we consider a bivariate toy example with two classes to which we apply the CQDA classification rule. We generated data from two bivariate normal distributions, depicted in the top left panel of Figure 1. The dataset contains 80 observations of class 1 (orange) and 100 observations of class 2 (blue). The tolerance ellipses correspond to the points $\mathbf{x}$ whose Mahalanobis distance

$$\text{MD}(\mathbf{x}, \bar{\mathbf{x}}_g, \mathbf{S}_g) = \sqrt{(\mathbf{x} - \bar{\mathbf{x}}_g)^\top \mathbf{S}_g^{-1} (\mathbf{x} - \bar{\mathbf{x}}_g)} \quad (2)$$

equals $\sqrt{\chi_{2,0.99}^2}$, the square root of the 0.99 quantile of the χ^2 distribution with $p = 2$ degrees of freedom. They visualize the shape of the empirical covariance matrices and fit the data nicely. The grey curve is the quadratic decision boundary obtained by CQDA. We see that it separates the classes quite well, misclassifying only three orange instances into the blue class. This is natural as there is some overlap between the classes.

In the top right panel we have introduced mislabeling by replacing the label of four observations from each class by the label of the other class. The blown up ellipses and the modified decision boundary show how much CQDA is affected. (The dotted curve is the decision boundary based on the uncontaminated data.) Although the mislabeled cases have no outlying measurements, they are outlying with respect to the class they are (incorrectly) assigned to. This changes the empirical mean and covariance matrix and therefore also the CQDA discriminant scores.

In a second experiment we replaced five observations from class 1 and eight observations from class 2 by outlying points, thereby introducing measurement noise. The blue outliers are positioned such that CQDA based on the clean data would assign them to the orange class. The result is shown in the lower left panel of Figure 1. The outliers have perturbed the classification, as both the decision boundary and the tolerance ellipses have changed substantially. Because of this the cluster of blue outliers is now classified into the blue class, instead of the orange class in which the uncontaminated boundary would put them.

Finally, the bottom right panel shows the dataset with both mislabeling and outliers, affecting the decision boundary and the ellipses even more.

Figure 1: The effect of label and measurement noise on classical quadratic discriminant analysis.

One must thus be careful when applying CQDA in practice, particularly in situations where mislabeling and/or contamination occur frequently. Note that when the dimension is above 3, the effects in Figure 1 are no longer visible by eye.

2.3 Real-time robust discriminant analysis

In order to make DA more reliable in the presence of label and/or measurement noise, several robust alternatives have been proposed. The most common strategy is to replace the classical estimators by robust counterparts. For example, Chork and Rousseeuw (1992) apply the Minimum Volume Ellipsoid estimator introduced in Rousseeuw (1984), He and Fung (2000) and

Croux and Dehon (2001) rely on S-estimators for linear discriminant analysis, whereas Hubert and Van Driessen (2004) propose to use the Minimum Covariance Determinant (MCD) estimator of Rousseeuw (1984). While this approach yields a more reliable version of DA, all these robust estimators become increasingly computationally demanding at large datasets. Particularly in industrial settings where large amounts of data have to be processed near-instantly, the existing algorithms become infeasible.

To address this issue we propose to incorporate the recently introduced real-time deterministic MCD (RT-DetMCD) method of (De Ketelaere et al., 2020). It is a parallel algorithm for the MCD estimator which runs very fast. For each class g the MCD estimator aims to find the subset of h_g observations whose sample covariance matrix has the lowest determinant. The raw MCD estimates of location and scatter are then computed as the classical mean and covariance matrix of the h_g observations in this subset. The number h_g should be chosen such that $\lfloor (n_g + p + 1)/2 \rfloor \leq h_g < n_g$ and such that $n_g - h_g$ is above the actual number of cases in class g that are contaminated by label or measurement noise. Since this number is unknown, we take $h_g = \lfloor (n_g + p + 1)/2 \rfloor$ to be able to withstand up to 50% of outliers in each class. Note that $h_g = 0.75 n_g$ is also often recommended, as it yields more efficient estimates and a degree of robustness that is sufficient in applications with less than 25% of contamination in each class.

The MCD estimator can be computed using the FastMCD (Rousseeuw and Van Driessen, 1999) or DetMCD (Hubert et al., 2012) algorithms, but their computational cost is still too high for our purpose. The RT-DetMCD estimator takes a different approach by first splitting up the data into q blocks which are processed in parallel. In each block an improved version of the DetMCD algorithm is run, yielding q estimates of location and scatter. These estimates are then combined through a novel robust pooling strategy which results in a single raw location and scatter estimate. This raw estimate allows a reweighting step in which the classical mean and covariance are computed on all unflagged observations, yielding the final reweighted location and scatter estimate. More details on the algorithm can be found in the Additional Material or (De Ketelaere et al., 2020).

The RT-DetMCD algorithm is incorporated into QDA as follows. For each class $g = 1, \dots, G$ in the training data we compute the RT-DetMCD estimates of location and scatter, denoted by $\hat{\boldsymbol{\mu}}_{g,R}$ and $\hat{\boldsymbol{\Sigma}}_{g,R}$ where the subscript R stands for robust. In order to estimate the class probabilities p_g robustly we use the following procedure (Hubert and Van Driessen, 2004): first we compute the robust distance of every observation $\mathbf{x}_i$ to its own class y_i by

$$\text{RD}_{i,y_i} = \text{RD}(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{y_i,R}, \hat{\boldsymbol{\Sigma}}_{y_i,R}) = \sqrt{(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_{y_i,R})^\top \hat{\boldsymbol{\Sigma}}_{y_i,R}^{-1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_{y_i,R})}. \quad (3)$$

Next we flag outliers within each class: we consider $\mathbf{x}_i$ an outlier when the robust distance with respect to its own class is too large, i.e. when $\text{RD}_{i,y_i} >$

$\sqrt{\chi_{p,0.99}^2}$. Note that these outliers can be the result of label and/or measurement noise. We now drop these outliers from the calculation of the membership probabilities, yielding robust estimates of the membership probabilities $\hat{p}_{g,R} = \tilde{n}_g/\tilde{n}$ where $\tilde{n}_g$ denotes the number of non-outliers in class g and $\tilde{n} = \sum_{g=1}^G \tilde{n}_g$. Including this $\hat{p}_{g,R}$ together with $\hat{\boldsymbol{\mu}}_{g,R}$ and $\hat{\boldsymbol{\Sigma}}_{g,R}$ into (1) yields robust discriminant scores.

On top of this we add another feature. Most classifiers have the disadvantage that a new case $\mathbf{x}$ will always be assigned to one of the known classes. In practice however, it is possible that the new case belongs to a different class which was not present in the training data. We therefore incorporate an anomaly detection step into QDA by assigning a new observation to the ‘overall outlier’ class with label 0 if its robust distance with respect to *all* classes is too large. A similar idea was used in the SIMCA method Wold and Sjöström (1977). Any unknown observation $\mathbf{x}$ is thus assigned to the class g for which the discriminant score is highest, under the condition that it does not deviate too much from all known classes. In the latter case, $\mathbf{x}$ will be given the label 0.

Putting all of this together, we obtain the proposed real-time robust quadratic discriminant analysis (RT-RQDA) classification given by

if $\min_g \text{RD}(\mathbf{x}, \hat{\boldsymbol{\mu}}_{g,R}, \hat{\boldsymbol{\Sigma}}_{g,R}) > \sqrt{\chi_{p,0.99}^2}$ **then** assign $\mathbf{x}$ to class 0;
else assign $\mathbf{x}$ to the class g for which $d(\mathbf{x}, \hat{\boldsymbol{\mu}}_{g,R}, \hat{\boldsymbol{\Sigma}}_{g,R}, \hat{p}_{g,R})$ is highest.

In order to illustrate this classifier we reconsider the toy example of Subsection 2.2. Figure 2 shows the result of training RT-RQDA on the clean dataset and on the datasets corrupted with label and/or measurement noise. The top left panel shows that RT-RQDA acts like classical QDA in the absence of noise, as the decision boundary and the tolerance ellipses are almost identical to those in the top left panel of Figure 1. The three other panels contain the same contamination as in Figure 1, but RT-RQDA now yields decision boundaries and tolerance ellipses that are very similar to those from the uncontaminated data. This robustness to label and measurement noise will be studied more extensively in the simulation study in Section 4.

3 The label bias plot

While the robust discriminant analysis method introduced above performs well under label and measurement noise, by itself it gives little insight into the presence or absence of such noise in individual observations. For that purpose we construct a graphical display which visualizes label and/or measurement noise in the data.

Suppose we have an observation $\mathbf{x}_i$ with observed class label y_i which is assigned to the class $\hat{y}_i$, so the discriminant score $d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{\hat{y}_i,R}, \hat{\boldsymbol{\Sigma}}_{\hat{y}_i,R}, \hat{p}_{\hat{y}_i,R})$ is the highest. If $\hat{y}_i = y_i$ the QDA classifier has assigned $\mathbf{x}_i$ to its given class,

Figure 2: The effect of label and measurement noise on robust quadratic discriminant analysis.

hence $d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{\hat{y}_i, R}, \hat{\boldsymbol{\Sigma}}_{\hat{y}_i, R}, \hat{p}_{\hat{y}_i, R}) = d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{y_i, R}, \hat{\boldsymbol{\Sigma}}_{y_i, R}, \hat{p}_{y_i, R})$. If on the other hand $d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{\hat{y}_i, R}, \hat{\boldsymbol{\Sigma}}_{\hat{y}_i, R}, \hat{p}_{\hat{y}_i, R}) > d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{y_i, R}, \hat{\boldsymbol{\Sigma}}_{y_i, R}, \hat{p}_{y_i, R})$ it follows that $\hat{y}_i \neq y_i$, so the observation $\mathbf{x}_i$ was predicted to belong to a class different from its given one. The larger this difference, the more the DA classifier wants to assign $\mathbf{x}_i$ to its predicted class $\hat{y}_i$ instead of its given class. A high difference could be caused by one or more of the following:

- the observation $\mathbf{x}_i$ may have been mislabeled (label noise);
- the observation $\mathbf{x}_i$ may be outlying with respect to its given class y_i (measurement noise);
- there may be overlap between the classes y_i and $\hat{y}_i$ making them difficult to separate.

To quantify label noise we define the *label bias* of the observation $\mathbf{x}_i$ as

$$\text{LB}(\mathbf{x}_i) = \sqrt{d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{\hat{y}_i, R}, \hat{\boldsymbol{\Sigma}}_{\hat{y}_i, R}, \hat{p}_{\hat{y}_i, R}) - d(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{y_i, R}, \hat{\boldsymbol{\Sigma}}_{y_i, R}, \hat{p}_{y_i, R})}. \quad (4)$$

The label bias is exactly zero for well-classified cases, and when it is strictly positive the predicted label differs from the given label. A high label bias raises suspicion that the given label might be wrong (mislabeling).

For quantifying potential measurement noise we use the robust distance (3) of $\mathbf{x}_i$ to its given class y_i . This measures how close $\mathbf{x}_i$ lies to the center of its class, relative to the scatter of its class. As described in Subsection 2.3, high robust distances can be used to flag outliers.

We now introduce the *label bias plot* (LB-plot) of class g as the scatter plot of the points

$$\left(\text{RD}(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{y_i, R}, \hat{\boldsymbol{\Sigma}}_{y_i, R}), \text{LB}(\mathbf{x}_i) \right) \quad (5)$$

for all $\mathbf{x}_i$ that belong to same class g , i.e. with $y_i = g$. Note that we defined LB in (4) as a square root, because the discriminant scores contain the squared robust distances and we want to put the axes of the label bias plot on the same footing.

In the label bias plot the points are colored according to their predicted class label $\hat{y}_i$. As one plot is made for each class, we obtain G different LB-plots. High values of LB suggest mislabeling, whereas points with high RD are outlying with respect to their observed class.

Figure 3 presents the LB-plots of the artificial data with mislabeling and outliers shown in the lower right panel of Figure 2. The top panel shows the LB-plots resulting from CQDA, whereas the bottom panel shows those for RT-RQDA.

Note that each LB-plot has two dashed lines, corresponding to cutoffs on the axes. For the robust distances on the horizontal axis we use the cutoff value $c = \sqrt{\chi_{p, 0.99}^2}$ that we already used in Subsection 2.3 to flag outliers in a class. For the LB on the vertical axis we use the cutoff $\sqrt{\ln(2)} = 0.83$. Points above this cutoff have a mixture model likelihood $p_g f_g(\mathbf{x}_i)$ that is at least twice as high for their predicted class $\hat{y}_i$ than for their given class y_i .

Finally, we depict the overall outliers, i.e. the observations for which $\min_g \text{RD}(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_{g, R}, \hat{\boldsymbol{\Sigma}}_{g, R}) > c$, as empty circles rather than dots. These observations are outlying with respect to every class in the data, suggesting they may be gross errors or members of a different population.

We first discuss the LB-plots of RT-RQDA in the lower half of Figure 3, since they are the most informative. Looking at the LB-plot for class 1 (orange), we immediately note that most points have $\text{LB} = 0$ and lie to the left of c . These are the regular observations from class 1 that are classified to belong to class 1. Next, we see five points with $\text{LB} = 0$ and a robust distance above the cutoff c . They are displayed as empty circles, indicating that they are also outlying with respect to class 2. They correspond to the orange outliers in Figure 2, which were generated to be closer to class 1 than to class 2. Hence they are not misclassified but the LB-plot indicates that they do not sit well in the orange class. The plot

Figure 3: LB-plots of classical and robust QDA applied to artificial data with label and measurement noise.

also contains seven points with a strictly positive label bias, meaning they have been assigned to the blue class although their given label is orange. Three of them have a relatively low LB, and a low robust distance indicating that they are within the robust tolerance ellipse of class 1. The LB-plot thus tells us that while the classifier assigns these points to class 2, they are also quite close to class 1. This suggests overlap, and indeed in Figure 2 these three points lie in the region where the ellipses overlap. There are also three points of class 1 with a high LB and $RD > c$. This suggests label or measurement noise, and these three points were indeed deliberately mislabeled in Figure 2. Note that the LB-plot does not allow us to determine the root cause of their outlyingness: it could be that the observations belong to class 1 but due to measurement noise they ended up closer to class 2, or they could actually belong to class 2 but received label 1 due to label noise. In either case the points would show up in the upper right quadrant of the LB-plot. Since the points are plotted as dots and not as empty circles, they are

not outlying with respect to class 2. This lends more credence to label noise than measurement noise. Finally there is one point with $LB \approx 2$ and a robust distance which is slightly smaller than c . This is a borderline case, which could be due to overlap or mislabeling. On the scatter plot in Figure 2 we see that one of the mislabeled points of class 1 indeed belongs to the overlapping region.

The LB-plot of class 2 indicates that most points are well classified, but there are 12 points with large LB and large RD. Four of them are plotted as dots, indicating that they do not have a large robust distance with respect to class 1. These correspond to the mislabeled points in Figure 2. The other eight points in the top right of the LB-plot are displayed as empty circles and are therefore outlying with respect to all classes. They correspond to the small blue cluster of outliers in Figure 2, which indeed lies closer to the orange class yet it is also outlying to that class.

The LB-plots based on RT-RQDA thus gave us insight into how well the cases were classified, whether there is overlap between the classes, and whether there are any suspicious labels or outliers. This information becomes even more useful in higher dimensions, when there is no version of Figure 2 to look at.

The LB-plots based on CQDA are shown in the top panels of Figure 3. The horizontal axis now shows the classical Mahalanobis distances (2) based on the sample mean and covariance matrix. The label bias is defined as in (4) but now based on the classical discriminant scores. Although we see some similarities with the LB-plots from RT-RGDA, there are several important differences. For class 1, only four out of the five overall outliers are flagged with empty circles, and only one of the mislabeled points has a distance above c . The other mislabeled points look less suspicious, instead the plot rather suggests they are in an overlapping region. In the LB-plot of class 2 only one point has a distance above c . Indeed, in the bottom right panel of Figure 1 the CQDA tolerance ellipses were so inflated that they engulf all blue points but one. In general, the label and measurement noise is less visible in the LB-plots based on CQDA than in those based on RT-RQDA. The LB-plots of CQDA would lead us to think that the misclassifications are mainly due to a large overlap between the classes, which we know is not the ground truth in these generated data. In practice we therefore recommend to use the LB-plots based on the robust estimates whenever available.

4 Simulation study

In this section we compare the performance of the proposed RT-RQDA classifier to that of classical QDA through a simulation study. We will denote the true classes as π_g for $g = 1, \dots, G$.

4.1 Simulation setup

We first generate uncontaminated (clean) data $\mathbf{X}$ with label vector $\mathbf{y}$. We create three classes π_1 , π_2 and π_3 of $p = 5$ dimensional observations. The observations of class π_g are sampled from a normal distribution $N(\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g)$, where the class centers are given by $\boldsymbol{\mu}_1 = (6, 0, 0, 0, 0)^\top$, $\boldsymbol{\mu}_2 = (0, 0, 6, 0, 0)^\top$ and $\boldsymbol{\mu}_3 = (0, 0, 0, 0, 6)^\top$.

The covariance matrices of the classes are the diagonal matrices $\Sigma_1 = I_5$, $\Sigma_2 = \text{diag}(1, 2, 3, 4, 5)$ and $\Sigma_3 = \text{diag}(1, 1, 1, 5, 10)$. The number of observations in each class is set to $n_1 = 250,000$, $n_2 = 350,000$ and $n_3 = 400,000$ with $n = 1,000,000$ to emulate industrial data sizes. The true prior probabilities of the classes π_1 , π_2 and π_3 are thus 0.25, 0.35 and 0.4.

In order to evaluate the performance of the method in more realistic scenarios, we also consider three simulation setups with contamination. In the first we only include label noise. In this setting the data is corrupted according to the *noisy completely at random* (NCAR) process Frénay and Verleysen (2013). For each class π_g we randomly relabel $\varepsilon_\ell n_g$ observations: $\varepsilon_\ell n_g / 2$ instances receive one of the other labels $k \neq g$, while the other half acquire the remaining label. We choose the label noise fraction $\varepsilon_\ell = 20\%$.

The second setup contains measurement noise but no label noise. The measurement noise is generated as follows. For class π_1 the outliers are sampled from cluster contamination at $N(\mu_1^*, \Sigma_1^*)$ with $\mu_1^* = (-6, 0, 0, 0, 0)^\top$ and $\Sigma_1^* = \frac{1}{10}\Sigma_1$. For class π_2 we consider point contamination by concentrating all outliers into a single point $\mu_2^* = (0, 0, -15, 0, 20)^\top$. For the third class π_3 we generate shift contamination according to $N(\mu_3^*, \Sigma_3)$ with $\mu_3^* = (14, 0, 0, 0, -6)^\top$. We replace $\varepsilon_m n_g$ random observations of each class g by outliers of the corresponding types, where ε_m denotes the fraction of measurement noise, which we set to 20% in this simulation. Note that the contamination has been generated in such a way that the contaminating points are outlying with respect to all classes. Hence, we expect a good classifier to predict them as ‘overall outliers’, that is, assign them to the additional class π_0 .

In the final setup we combine both types of noise by introducing $\varepsilon_\ell / 2$ label noise and $\varepsilon_m / 2$ measurement noise, both as described above.

4.2 Evaluation of results

Each simulation setup is replicated 50 times and the performance of RT-RQDA is compared to CQDA in several ways. First of all, we compare the misclassification errors of both methods by computing the average confusion matrix over all replications. The standard confusion matrix will not allow us to properly analyze the results under label and measurement noise. Therefore, we will present extended confusion matrices, for which we need some notation. Consider class π_g . The observations in this class can be split up in up to 4 subclasses, depending on the contamination scheme. The largest subclass is that of the clean observations, which were generated from $N(\mu_g, \Sigma_g)$ and received the correct label. This subclass is denoted $\pi_{g,g}$. Secondly, there are the observations which were generated from the clean distribution of class π_g but which received a wrong label. We denote these as $\pi_{g,k}$ and $\pi_{g,\ell}$ where k and ℓ are the labels of the other two classes. Finally, there are the observations with measurement error denoted as $\pi_{g,0}$.

Instead of creating confusion matrices with three rows, one for each class, we split up each class into its nonempty subclasses. For the experiment with clean data, we still have only three rows. The experiment with only label noise has three subclasses in each class, and hence yields an extended confusion matrix with nine rows. When only measurement noise occurs each class has two subclasses, yielding

an extended confusion matrix with six rows. The experiment with both label and measurement noise has all twelve rows. By extending the confusion matrices in this way, we obtain very precise insight into what happens with every type of data point, corrupted or otherwise. Finally, we also add an additional column to the confusion matrix, denoted π_0 , which shows the percentage of points classified as an overall outlier using the outlier detection rule of RT-RQDA.

In addition to the extended confusion matrices we also report the average Kullback–Leibler (KL) divergence. The KL divergence of the estimated scatter matrix $\hat{\Sigma}_g$ from the true covariance matrix Σ_g of class π_g is defined as

$$\text{KL}_g = \text{trace}(\hat{\Sigma}_g \Sigma_g^{-1}) - p - \ln |\hat{\Sigma}_g \Sigma_g^{-1}|. \quad (6)$$

The determinants of the true covariance matrices are $|\Sigma_1| = 1$, $|\Sigma_2| = 120$ and $|\Sigma_3| = 50$. The estimated determinants $|\hat{\Sigma}_g|$ are also reported. Ideally, the estimated determinants should be close to the theoretical ones as this suggests that they accurately describe the volume of the clean data cloud. Finally, for each class we also report how many observations are flagged as outlying with respect to their own class. These are the observations $\mathbf{x}_i$ for which $\text{RD}_{i,y_i} > c$ where RD_{i,y_i} is the robust distance (3). We normalize this number by the total number of noisy observations in class π_g given by $(\varepsilon_\ell + \varepsilon_m)n_g$, and denote the resulting quantity by α_g . Ideally, α_g should be close to one, as this indicates the label and measurement noise are well detected. Our simulations were done in MATLAB on an Intel Core i7-8700K processor based computer with 16 GB of 3.70GHz RAM. The same hardware was used in all experiments.

4.3 Results on clean data

We first consider the setting of clean data, for which the results are presented in Table 1. The confusion matrix has only three rows, one for each class, since no label noise or measurement noise was introduced. Since we have set the outlier cutoff to $\sqrt{\chi_{p,0.99}^2}$, perfect classification corresponds with diagonal values equal to 0.99 and a π_0 column equal to 1%. We see that both classification methods perform very well on clean data. Note that CQDA is somewhat more efficient in this clean setting as it has lower KL values than RT-RQDA, and the determinants of the estimated covariance matrices are closer to the true values. The tolerance ellipsoids of RT-RQDA have somewhat lower volumes, resulting in slightly higher values in the column of π_0 .

4.4 Results under label noise

The results of the simulation setup with 20% label noise are summarized in Table 2. For CQDA the performance metrics show large increases in the KL divergences and in the determinants of the estimated covariance matrices. Although CQDA fails to accurately estimate the model parameters, its classification error remains rather low as can be inferred from its extended confusion matrix. This is mainly due to the fact that the true classes are well separated. Unlike CQDA, RT-RQDA shows a very stable behavior, with virtually unchanged estimates of the model parameters and a confusion matrix that is similar to the clean setting. Finally,

Table 1: Classical and robust discriminant analysis results for uncontaminated data.

	CQDA				RT-RQDA			
	π_1	π_2	π_3	π_0	π_1	π_2	π_3	π_0
Extended confusion matrices								
$\pi_{1,1}$	0.990	0.000	0.000	0.010	0.986	0.000	0.000	0.013
$\pi_{2,2}$	0.000	0.981	0.010	0.009	0.000	0.978	0.009	0.013
$\pi_{3,3}$	0.000	0.004	0.986	0.009	0.000	0.005	0.983	0.013
Performance metrics								
KL	<0.001	<0.001	<0.001		0.007	0.007	0.007	
$ \hat{\Sigma} $	1.001	120.100	49.994		0.764	91.750	38.195	

note that the α_g values suggest that CQDA identifies only a small fraction of the mislabeled observations as outliers with respect to the class of their given label, whereas RT-RQDA flags nearly all of them.

Table 2: Classical and robust discriminant analysis results with 20% label noise.

	CQDA				RT-RQDA			
	π_1	π_2	π_3	π_0	π_1	π_2	π_3	π_0
Extended confusion matrices								
$\pi_{1,1}$	0.999	0.000	0.000	0.000	0.986	0.000	0.000	0.013
$\pi_{1,2}$	0.999	0.000	0.000	0.000	0.986	0.000	0.000	0.013
$\pi_{1,3}$	0.999	0.000	0.000	0.000	0.986	0.000	0.000	0.013
$\pi_{2,1}$	0.001	0.975	0.021	0.004	0.000	0.979	0.008	0.012
$\pi_{2,2}$	0.001	0.975	0.020	0.004	0.000	0.979	0.008	0.012
$\pi_{2,3}$	0.001	0.975	0.021	0.004	0.000	0.979	0.008	0.012
$\pi_{3,1}$	0.003	0.002	0.992	0.003	0.000	0.005	0.982	0.012
$\pi_{3,2}$	0.003	0.002	0.992	0.003	0.000	0.005	0.982	0.013
$\pi_{3,3}$	0.003	0.002	0.992	0.003	0.000	0.005	0.982	0.013
Performance metrics								
KL	11.123	2.394	2.315		0.007	0.007	0.007	
$ \hat{\Sigma} $	143.230	1047.600	508.410		0.766	97.165	38.609	
α	0.329	0.133	0.383		0.949	0.926	1.082	

4.5 Results under measurement noise

Table 3 reports the results when the data contains 20% of measurement noise. Since by construction all contaminated points are outlying to all classes, they should be classified into π_0 . The extended confusion matrix now has six rows because each class has a subclass generated as outliers ($\pi_{g,0}$) and a subclass generated as clean data ($\pi_{g,g}$). We see major differences between CQDA and RT-RQDA.

First note that CQDA barely detects outliers as indicated by the low values in the column of π_0 . Instead, it classifies the outliers as regular observations. The performance metrics in the table reveal the problem with CQDA: the estimates of the parameters are affected so strongly by the outliers that the resulting model does not detect them anymore. This illustrates the so-called masking effect (Maronna et al., 2019). In contrast, RT-RQDA correctly detects the introduced noise as outlying observations. A second important effect is that with CQDA the classification of the clean data suffers. In particular, the regular observations of class π_1 and class π_3 have a misclassification rate of roughly 20% and 40%. The proposed RT-RQDA method instead has a stable performance which is comparable to that in the clean scenario of Table 1.

Table 3: Classical and robust discriminant analysis results with 20% measurement noise.

	CQDA				RT-RQDA			
	π_1	π_2	π_3	π_0	π_1	π_2	π_3	π_0
Extended confusion matrices								
$\pi_{1,0}$	1.000	0.000	0.000	0.000	0.000	0.000	0.000	1.000
$\pi_{1,1}$	0.794	0.000	0.205	0.000	0.989	0.000	0.000	0.011
$\pi_{2,0}$	0.000	1.000	0.000	0.000	0.000	0.000	0.000	1.000
$\pi_{2,2}$	0.000	0.977	0.005	0.018	0.000	0.980	0.009	0.010
$\pi_{3,0}$	0.193	0.000	0.782	0.026	0.000	0.000	0.000	1.000
$\pi_{3,3}$	0.193	0.191	0.614	0.003	0.000	0.004	0.985	0.010
Performance metrics								
KL	42.500	47.844	37.411		0.001	0.001	0.001	
$ \hat{\Sigma} $	18.029	2180.300	1732.700		0.905	108.720	45.223	
α	0.065	0.074	0.046		1.046	1.045	1.045	

4.6 Results under label and measurement noise

Finally, Table 4 presents the simulation results with label and measurement noise occurring simultaneously. Also here RT-RQDA exhibits a stable behavior. All of the far outliers are detected and allocated to the outlier class π_0 . The misclassification errors are similar to those for clean data and the model parameters are estimated accurately, as indicated by the low values of the KL divergence and the determinants of the estimated covariance matrices.

In contrast, CQDA is heavily affected by the noise. It fails to detect the outliers as can be seen from the low percentages in the column of class π_0 and the low values of α_g . The classification of the clean data is also affected, in particular for subclass $\pi_{3,3}$ which has a misclassification error of 18%. The estimated parameters deviate substantially from their true values as evidenced by the elevated KL divergences and high determinants.

From the simulation study we conclude that the proposed RT-RQDA method is more reliable than CQDA in the presence of label and/or measurement noise while maintaining a competitive performance on clean data.

Table 4: Classical and robust discriminant analysis results with 10% label and 10% measurement noise.

	CQDA				RT-RQDA			
	π_1	π_2	π_3	π_0	π_1	π_2	π_3	π_0
Extended confusion matrices								
$\pi_{1,0}$	0.000	0.000	1.000	0.000	0.000	0.000	0.000	1.000
$\pi_{1,1}$	0.940	0.000	0.060	0.000	0.988	0.000	0.000	0.012
$\pi_{1,2}$	0.940	0.000	0.060	0.000	0.988	0.000	0.000	0.012
$\pi_{1,3}$	0.940	0.000	0.060	0.000	0.987	0.000	0.000	0.012
$\pi_{2,0}$	0.000	1.000	0.000	0.000	0.000	0.000	0.000	1.000
$\pi_{2,1}$	0.004	0.980	0.011	0.005	0.000	0.981	0.009	0.011
$\pi_{2,2}$	0.004	0.980	0.011	0.005	0.000	0.980	0.009	0.011
$\pi_{2,3}$	0.004	0.980	0.011	0.005	0.000	0.980	0.009	0.011
$\pi_{3,0}$	0.491	0.000	0.373	0.136	0.000	0.000	0.000	1.000
$\pi_{3,1}$	0.110	0.065	0.823	0.002	0.000	0.005	0.984	0.011
$\pi_{3,2}$	0.110	0.065	0.822	0.002	0.000	0.005	0.983	0.011
$\pi_{3,3}$	0.110	0.065	0.822	0.002	0.000	0.005	0.983	0.011
Performance metrics								
KL	24.547	21.399	17.698		0.003	0.004	0.003	
$ \hat{\Sigma} $	141.900	4596.400	2259.200		0.837	106.090	42.160	
α	0.251	0.159	0.246		1.014	0.987	1.058	

5 Applications to real data

5.1 Diabetes data

We start by analyzing a benchmark data set which arose out of a study about the onset of diabetes mellitus Smith et al. (1988). The data were downloaded from Leisch and Dimitriadou (2015). The subjects (cases) are females from a community of Pima Native Americans, with classes ‘positive’ (diabetic) and ‘negative’. For each subject eight variables were recorded: number of pregnancies, plasma glucose concentration, diastolic blood pressure, triceps skin fold thickness, serum insulin, a body mass index (BMI), diabetes pedigree function, and age. Following Ripley (1996) we removed the cases with physically impossible measurements (such as zero blood pressure), leaving us with 392 subjects. Since several variables were skewed, we robustly transformed all variables toward central normality by the technique of Raymaekers and Rousseeuw (2020) with its default options. We also verified that each of the eight variables has a positive relation to diabetes, in the sense that higher values increase the odds of the disease.

Applying RT-RQDA yielded the LB-plots in Figure 4. The left panel shows the negative class. The correctly classified cases are shown as blue points with label bias $LB = 0$. The red points were misclassified into the positive class. Some of them have a high label bias above the dashed horizontal line, meaning that RT-RQDA considers it more than twice as likely for them to be diabetic rather than negative. The most extreme ones are marked (a), (b) and (c). Case (a) has high glucose, insulin and age, while (b) has high blood pressure and pedigree. Case

(c) has high glucose, pedigree, and the highest insulin of the entire dataset. All of these characteristics point to diabetes, arousing suspicion that these subjects were misclassified or have a high probability of being diagnosed in the future. We

Figure 4: LB-plots of the classes of negative (left) and positive (right) subjects in the diabetes data, using RT-RQDA.

also see two points with high robust distance but lower LB. Subject (d) has high blood pressure and very high BMI. Subject (e) is correctly assigned to the negative class since $LB = 0$, but has a high distance partly due to an extremely low blood pressure of 30. Note that (a), (d) and (e) have hollow plot symbols, indicating that their distance to the other class is above the distance cutoff as well, so they are not close to either class. On the other hand, (c) has a high distance from the negative class but its solid plot symbol indicates that it is not so far from the other class.

The LB-plot of the subjects labeled positive for diabetes is in the right panel of Figure 4. The subjects with highest label bias are marked (f) and (g). Subject (f) has very low triceps, insulin and BMI, which is in line with the negative class. Subject (g) has extremely low blood pressure and rather low glucose and insulin. These would point to the negative class, hence the high LB. The high BMI of (g) correctly suggests the positive class but is unusual, which increases the robust distance on the horizontal axis. Point (h) corresponds to the highest blood pressure in the entire dataset as well as the highest BMI. Subject (i) has the highest triceps in the dataset as well as the highest pedigree. From the plot symbols of (g), (h) and (i) we conclude that they lie far away from the centers of both classes.

5.2 Fruit data

We analyze the fruit dataset originally collected by Colin Greensill (Central Queensland University, Rockhampton, Australia), parts of which were analyzed in Hubert and Van Driessen (2004); Vanden Branden and Hubert (2005). The original dataset contains the result of a spectroscopy experiment conducted on $n = 2818$

cantaloupe melons of six different cultivars. Each of the spectra was measured on 256 wavelengths. We selected four of the six cultivars (labeled “D”, “H”, “Ha” and “E”) for illustrating our methods. The motivation for this selection is twofold. First, the largest cultivar (named “ES”) was not included because it did not have a roughly elliptical distribution, which makes its classification by QDA less appropriate. Second, the smallest cultivar (named “M”) was not included because we do not have additional information on it, making the interpretation of the results more difficult. The selected four cultivars yield a dataset of size 1774×256 which consists of four classes of sizes 490, 180, 500 and 988. It is known that the data from cultivar Ha was collected using different illumination setups and so one might expect subgroups within this class. The previous analyses of Hubert and Van Driessen (2004); Vanden Branden and Hubert (2005) did in fact identify one subgroup within this cultivar corresponding with a different illumination setup.

We aim to classify the spectra into the four classes corresponding to the cultivars. It is worth noting that the spectra were all recorded on the same wavelengths and scale, and they do not contain missing values. Therefore we did not apply any preprocessing method. Before applying RT-RQDA we reduced the dimensionality by means of robust principal component analysis (Hubert et al., 2005), and selected $p = 3$ components as they explain over 97.5% of the variability. In the context of classification, there may be more appropriate techniques for dimensionality reduction such as approaches based on (local) linear discriminant analysis (Martinez and Kak, 2001; Sugiyama, 2007). However, it is recommended to use a robust method in order to avoid the loss of information in the reduced data due to the presence of outliers. Because of its simplicity and robustness and for the purpose of illustration, we decided on robust PCA. Figure 5 shows the scores resulting from this dimensionality reduction step. We clearly see that class Ha consists of two subgroups, possibly due to the different illumination setups of the spectroscopy experiment. In addition we see that classes E and H may also contain subgroups.

We apply RT-RQDA to the data and construct the LB-plots of all classes. Figure 6 shows the LB-plot of cultivar D as well as a clarifying pairs plot containing only the points of cultivar D. All points in this LB-plot belong to class D, but some have $LB > 0$ meaning they are assigned to other classes indicated by their color. The red points in the LB-plot suggest that there is some overlap with class H. This overlap was to be expected from the original plot of the scores in Figure 5. In addition to the overlap with cultivar H, class D has quite a few outliers as can be seen from the green and red empty circles in the LB-plot. This is confirmed by the position of these points in the right panel of Figure 6, where the same colors and plotting symbols were used as in the left panel.

The LB-plot of cultivar H is shown in Figure 7. We see right away that this cultivar consists of two subgroups. One subgroup is found in the bottom left corner of the LB-plot. This subgroup contains the bulk of the data, and has a little bit of overlap with cultivar D as indicated by the handful of green points with $LB > 0$. The other subgroup is found in the top right portion of the LB-plot. These green empty circles have a high label bias and a high robust distance. The latter indicates that they lie quite far from the bulk of the data, which is confirmed in the plot of (PC2, PC3) in the right panel of Figure 7. The fact that these objects

Figure 5: Pairwise scatter plot of the scores of the fruit data after robust principal component analysis.

Figure 6: LB-plot of cultivar D (left) with a pairs plot of the PC scores of this class using the same colors and plotting symbols (right).

Figure 7: LB-plot of cultivar H (left) with a pairs plot of the PC scores of this class using the same colors and plotting symbols (right).

are plotted as empty circles in the LB-plot indicates that they are also far from the other classes. It turns out that this subgroup corresponds exactly to the first 60 spectra measured for this cultivar. This strongly suggests that the experimental setup was changed after 60 measurements.

Figure 8: LB-plot of cultivar Ha (left) with a pairs plot of the PC scores of this class using the same colors and plotting symbols (right).

Figure 8 shows the LB-plot of cultivar Ha. It is known that during the spectroscopy experiments on this cultivar the illumination setup was changed twice. The initial illumination setup produced the first 180 points, which can be seen in the top right of the LB-plot. It corresponds to all points with a robust distance

over 20. Therefore the initial illumination has produced a different data distribution from the others. This is confirmed by the clear separation in the pairwise plots in the right panel of Figure 8. The remaining 320 spectra of cultivar Ha are found in the bottom left of the LB-plot. The blue points are correctly assigned to class Ha, and the orange ones to class E. It is not immediately clear whether these orange points form a subgroup which should be considered separately, but in the pairs plots we see more of a gradual transition than a clean separation.

The LB-plot of cultivar E is shown in Figure 9. Also here we see two distinct classes. The majority of the points has $LB = 0$ and a relatively small robust distance, suggesting a homogeneous subgroup. This is confirmed by the orange points in the right hand panel of Figure 9. The other subgroup is found in the top right portion of the LB-plot, and clearly separated from the bulk of the data in the pairwise plot of (PC2,PC3). It turns out that the spectra of this cultivar were collected on two different dates. The first 200 were measured roughly one month before the remaining 788. Furthermore, the outlying subgroup visible on the plot corresponds almost exactly with the first 200, suggesting that the subgroups cannot be considered as coming from the same population. This may be the result of different experimental setups, which is not unlikely given that the data of the two groups were collected one month apart.

Figure 9: LB-plot of cultivar E (left) with a pairs plot of the PC scores of this class using the same colors and plotting symbols (right).

In order to evaluate the classification performance of RT-RQDA on the fruit data, we split the data at random into a training set and a validation set containing 60% and 40% of the data (i.e. 1294 and 864 observations). Note that the validation set also contains outliers, which we do not want to take into account when evaluating classifiers. The points in the validation set that are outlying with respect to their class are therefore discarded in the calculation of misclassification rates. This procedure yields a validation set of 632 observations, with classes of sizes 179, 44, 106 and 303 for cultivars D, H, Ha and E.

The left part of Table 5 shows the resulting confusion matrix of CQDA. Its misclassification rate is about 22%. In the right part of Table 5 we see that RT-RQDA performs better, with a misclassification rate of 2%. The largest difference between CQDA and RT-RQDA occurs in the classification of cultivar Ha. Here CQDA assigns most non-outlying spectra to cultivar E, indicating that it has failed to characterize cultivar Ha due to its large number of outliers. The RT-RQDA method does classify cultivar Ha accurately, and only shows a slight confusion between cultivars D and H which is not surprising as these classes overlap.

Table 5: Confusion matrix of the validation set based on classical (left) and robust (right) discriminant analysis.

	CQDA				RT-RQDA			
	D	H	Ha	E	D	H	Ha	E
D	0.933	0.034	0.011	0.022	0.950	0.050	0.000	0.000
H	0.500	0.477	0.023	0.000	0.136	0.864	0.000	0.000
Ha	0.000	0.057	0.000	0.943	0.000	0.000	1.000	0.000
E	0.000	0.000	0.000	1.000	0.000	0.000	0.000	1.000

6 Conclusion

Classical quadratic discriminant analysis (CQDA) is known to be sensitive to label noise and measurement noise (outliers). To address this issue we proposed a procedure for real-time robust quadratic discriminant analysis, called RT-RQDA. It incorporates the recent RT-DetMCD algorithm as well as an anomaly detection step, which flags observations that stand out relative to all classes. Robust discriminant scores are obtained from the RT-DetMCD estimates of location and scatter as well as robust membership probabilities. An extensive simulation study showed that the speed of RT-RQDA allows it to handle huge data sets and that it remains effective even if the data are contaminated by label and measurement noise simultaneously.

We also introduced a graphical display, the LB-plot, which gives insight into the presence of label and measurement noise in each class of the training data. The LB-plot makes atypical observations stand out, thereby identifying outliers, potential mislabeling, overlap, and observations which may be hard to classify.

Finally, the proposed method was illustrated on two real datasets. In both applications, the RT-RQDA procedure with the LB-plot correctly identified several atypical observations.

Acknowledgements

We thank Johan Speybroeck and Tim Wynants for their support throughout the project. We also acknowledge the financial support of VLAIO, grant HBC.2016.0208

(Development and implementation of real-time, robust statistical methods with novel applications in food sorting) for making this industrial research possible. The work of Mia Hubert, Jakob Raymaekers and Peter Rousseeuw was supported by the grant C16/15/068 (Statistical methodology for immense lively experiments) of the Research Council of KU Leuven.

References

- Chork, C. and P. J. Rousseeuw (1992). Integrating a high-breakdown option into discriminant analysis in exploration geochemistry. *Journal of Geochemical Exploration* 43(3), 191–203.
- Croux, C. and C. Dehon (2001). Robust linear discriminant analysis using S-estimators. *Canadian Journal of Statistics* 29(3), 473–493.
- 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.
- Frénay, B. and M. Verleysen (2013). Classification in the presence of label noise: a survey. *IEEE Transactions on Neural Networks and Learning Systems* 25(5), 845–869.
- Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning*. New York: Springer.
- He, X. and W. Fung (2000). High breakdown estimation for multiple populations with applications to discriminant analysis. *Journal of Multivariate Analysis* 72, 151–162.
- Hubert, M., P. J. Rousseeuw, and K. Vanden Branden (2005). ROBPCA: a new approach to robust principal component analysis. *Technometrics* 47(1), 64–79.
- Hubert, M., P. J. Rousseeuw, and T. Verdonck (2012). A deterministic algorithm for robust location and scatter. *Journal of Computational and Graphical Statistics* 21(3), 618–637.
- Hubert, M. and K. Van Driessen (2004). Fast and robust discriminant analysis. *Computational Statistics & Data Analysis* 45(2), 301–320.
- James, G., D. Witten, T. Hastie, and R. Tibshirani (2013). *An Introduction to Statistical Learning*. New York: Springer.
- Leisch, F. and E. Dimitriadou (2015). *mlbench*: Machine learning benchmark problems. R package version 2.1-1, CRAN.
- Maronna, R. A., R. D. Martin, V. J. Yohai, et al. (2019). *Robust Statistics: Theory and Methods (with R)*. John Wiley & Sons.

- Martinez, A. M. and A. C. Kak (2001). PCA versus LDA. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(2), 228–233.
- Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*. MIT Press.
- Raymaekers, J. and P. J. Rousseeuw (2020). Transforming variables to central normality, arXiv: 2005.07946.
- Ripley, B. D. (1996). *Pattern Recognition and Neural Networks*. Cambridge University Press.
- Rousseeuw, P. J. (1984). Least median of squares regression. *Journal of the American Statistical Association* 79(388), 871–880.
- Rousseeuw, P. J. and K. Van Driessen (1999). A fast algorithm for the minimum covariance determinant estimator. *Technometrics* 41(3), 212–223.
- Smith, J., J. Everhart, W. Dickson, W. Knowler, and R. Johannes (1988). Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. In *Proceedings of the Symposium on Computer Applications and Medical Care*, pp. 261–265. IEEE.
- Sugiyama, M. (2007). Dimensionality reduction of multimodal labeled data by local Fisher discriminant analysis. *Journal of Machine Learning Research* 8, 1027–1061.
- Vanden Branden, K. and M. Hubert (2005). Robust classification in high dimensions based on the SIMCA method. *Chemometrics and Intelligent Laboratory Systems* 79(1-2), 10–21.
- Wold, S. and M. Sjöström (1977). SIMCA: A method for analyzing chemical data in terms of similarity and analogy. In B. Kowalski (Ed.), *Chemometrics: Theory and Application*, Chapter 12, pp. 243–282. American Chemical Society Publications.

Additional Material: Description of the RT-DetMCD algorithm

We give a detailed description of the RT-DetMCD algorithm of De Ketelaere et al. (2020) which computes the MCD estimator of Rousseeuw (1984). Given n observations in p dimensions, the objective is to find the subset H of h observations whose sample covariance matrix has the lowest possible determinant.

The raw MCD-estimate of location $\hat{\boldsymbol{\mu}}_H$ is the average of these h points:

$$\hat{\boldsymbol{\mu}}_H = \frac{1}{h} \sum_{i \in H} \mathbf{x}_i \quad (7)$$

whereas the raw scatter estimate is a multiple of their covariance matrix:

$$\hat{\boldsymbol{\Sigma}}_H = \frac{c_\alpha}{h-1} \sum_{i \in H} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_H)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_H)^\top, \quad (8)$$

where c_α is a consistency factor that depends on h and n .

The best known algorithms for computing the MCD estimator are FastMCD (Rousseeuw and Van Driessen, 1999) and DetMCD (Hubert et al., 2012), both of which make use of so-called concentration steps (C-steps) to improve the subset H . Given a candidate subset of h observations H_0 , the C-step computes a more concentrated approximation by calculating the Mahalanobis distance of every observation $\mathbf{x}_i$ based on the location and scatter of the current subset H_0 as

$$\text{RD}_i = \sqrt{(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_{H_0})^\top \hat{\boldsymbol{\Sigma}}_{H_0}^{-1} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_{H_0})}. \quad (9)$$

These distances are then ordered, and the h observations with the lowest distances form the new h -subset H_1 . The new subset was proved to have an equal or lower determinant (Rousseeuw and Van Driessen, 1999), and this C-step procedure can thus be iterated until convergence. While guaranteed to converge, the procedure may only arrive in a local minimum. Therefore, the C-step procedure is run from different starting values, and after convergence the solution with the lowest determinant is retained. FastMCD starts from randomly sampled starting values, whereas DetMCD uses six deterministic starting values using several fast robust estimators in the first step. While FastMCD and DetMCD are popular algorithms that work well on medium sized datasets, their computational cost makes them impractical for real-time applications or very large industrial datasets.

Recently, a parallel algorithm called RT-DetMCD was introduced (De Ketelaere et al., 2020), which is suitable for very large datasets. In addition to the parallelization strategy, RT-DetMCD incorporates three key improvements: new initial estimators, fast C-step updates, and one-pass aggregation. The resulting algorithm is specifically designed for large datasets. It preserves the robustness and conceptual simplicity of the original MCD. As the algorithm is quite lengthy, we focus on the main differences compared with DetMCD. Further details are found in (De Ketelaere et al., 2020). A research-level MATLAB implementation of RT-DetMCD is publicly available from the web page <http://wis.kuleuven.be/statdatascience/robust/software>.

The RT-DetMCD algorithm starts by splitting up the standardized observations $\mathbf{z}_i$ into q blocks $\mathbf{Z}^{(\ell)}$ with $\ell = 1, \dots, q$. For each block the algorithm computes the estimates $\hat{\boldsymbol{\mu}}_{raw}^{(\ell)}$ and $\hat{\boldsymbol{\Sigma}}_{raw}^{(\ell)}$ by applying C-steps starting from two new initial estimators instead of DetMCD's six initial estimators. These q solutions then need to be aggregated in a robust way. They have many dimensions since the symmetric matrices $\hat{\boldsymbol{\Sigma}}_{raw}^{(\ell)}$ contain $p(p-1)/2$ distinct entries, and the $\hat{\boldsymbol{\mu}}_{raw}^{(\ell)}$ have p additional entries. Due to this high dimension, computing a typical robust estimate of the q fits is problematic. Instead the entry-wise median of the q fits is computed, yielding the entry-wise median of the $\hat{\boldsymbol{\mu}}_{raw}^{(\ell)}$ denoted as

$$\hat{\boldsymbol{\mu}}_{med} = (\text{median}_{\ell}((\hat{\boldsymbol{\mu}}_{raw}^{(\ell)})_1), \dots, \text{median}_{\ell}((\hat{\boldsymbol{\mu}}_{raw}^{(\ell)})_p)) \quad (10)$$

and the entry-wise median of all scatter matrices, given by

$$(\hat{\boldsymbol{\Sigma}}_{med})_{jk} = \text{median}_{\ell}((\hat{\boldsymbol{\Sigma}}_{raw}^{(\ell)})_{jk}) \quad (11)$$

for $j, k = 1, \dots, p$. However, as $\hat{\boldsymbol{\Sigma}}_{med}$ is not necessarily positive definite one cannot use it as a final aggregated outcome. Instead, each thread computes the Kullback-Leibler deviation $\text{KL}[(\hat{\boldsymbol{\Sigma}}_{med}, \hat{\boldsymbol{\mu}}_{med}), (\hat{\boldsymbol{\Sigma}}_{raw}^{(l)}, \hat{\boldsymbol{\mu}}_{raw}^{(l)})]$ given by

$$\begin{aligned} \text{KL}[(\mathbf{A}, \mathbf{a}), (\mathbf{B}, \mathbf{b})] &:= \text{trace}(\mathbf{A}\mathbf{B}^{-1} - \mathbf{I}) - \ln |\mathbf{A}\mathbf{B}^{-1}| \\ &\quad + (\mathbf{a} - \mathbf{b})^{\top} \mathbf{B}^{-1} (\mathbf{a} - \mathbf{b}) . \end{aligned} \quad (12)$$

Next, the h -subsets of the blocks with the $[q/2]$ lowest KL deviations are aggregated by single-pass pooling. This yields the raw RT-DetMCD solution $(\hat{\boldsymbol{\Sigma}}_{raw}, \hat{\boldsymbol{\mu}}_{raw})$. Each thread then computes squared distances $\text{RD}^2(\mathbf{z}_i, \hat{\boldsymbol{\mu}}_{raw}, \hat{\boldsymbol{\Sigma}}_{raw})$ and flags the $\mathbf{z}_i$ for which they exceed a quantile of the chi-squared distribution with p degrees of freedom. This yields a binary vector $\boldsymbol{\delta}$ of length n given by

$$\delta_i = \begin{cases} 1 & \text{if } \text{RD}^2(\mathbf{z}_i, \hat{\boldsymbol{\mu}}_{raw}, \hat{\boldsymbol{\Sigma}}_{raw}) \leq \chi_{p,0.975}^2 \\ 0 & \text{otherwise.} \end{cases}$$

Then reweighted estimates are computed as the classical mean and covariance of the observations with $\delta_i = 0$, yielding the final RT-DetMCD estimates $(\hat{\boldsymbol{\Sigma}}_{rew}, \hat{\boldsymbol{\mu}}_{rew})$. These estimates are then used one last time to flag the outliers in the dataset using the same rule as above. By undoing the initial standardization, all location and scale estimates are transformed back to the original coordinate system.