DIRECTIONAL QUANTILE CLASSIFIERS

By Alessio Farcomeni[†] and Marco Geraci[‡] and Cinzia Viroli[§]

University of Rome "Tor Vergata"[†], *University of South Carolina*[‡], *University of Bologna*[§]

Abstract   We introduce classifiers based on directional quantiles. We derive theoretical results for selecting optimal quantile levels given a direction, and, conversely, an optimal direction given a quantile level. We also show that the misclassification rate is infinitesimal if population distributions differ by at most a location shift and if the number of directions is allowed to diverge at the same rate of the problem's dimension. We illustrate the satisfactory performance of our proposed classifiers in both small and high dimensional settings via a simulation study and a real data example. The code implementing the proposed methods is publicly available in the `R` package `Qtools`.

1. Introduction.   The idea of using quantiles in classification is relatively recent and largely unexplored. The median classifier for high-dimensional problems proposed by Hall, Titterington and Xue (2009), which calculates the L_1 distance of the coordinates of a multivariate data point from componentwise medians (rather than centroids), is particularly advantageous when data exhibit heavy-tailed or skewed distributions. Building on Hall, Titterington and Xue's (2009) idea, Hennig and Viroli (2016a) proposed quantile classifiers which hinge on the sum of distances from componentwise quantiles at some generic level $\theta \in (0,1)$. The ensemble quantile classifier by Lai and McLeod (2020) assigns weights to the componentwise distances by minimising a regularised loss function, where the regularisation parameter is determined by cross-validation.

In all the studies mentioned above, quantiles are calculated marginally for each input variable (componentwise). This implies that their calculation ignores the possible interdependence among variables. In this study, we consider directional quantiles for multivariate distributions (Kong and Mizera, 2012) to address such a limitation. Our choice is motivated by several reasons. First, as already mentioned, the dependence among variables is

---

taken into account by computing linear combinations of input variables. Second, directional quantiles have a simple interpretation since the projections' weights embody the relative importance of the variables involved in the classification problem. Finally, in the special case of p canonical directions (with p equal to the number of variables), the use of directional quantiles leads to the componentwise quantile classifier (Hennig and Viroli, 2016a), and thus inherits asymptotic optimal properties as shown in Appendix. Directional quantiles have already found application in risk classification problems (Geraci et al., 2020) and proved to be a worthwhile alternative to risk classification based on componentwise quantile thresholds.

In general, the application of our methods does not require any assumption on the shape of the population distributions. We derive asymptotic theoretical properties of the proposed classifier, under the assumption that distributions for alternative populations differ by at most a location-shift. While this assumption may be unrealistic in practice, empirical results support the merit of the proposed classifier also when the distributions differ by shape and not just by location.

The rest of the paper is organised as follows. In the next section, we introduce notation and basic definitions, followed by our proposal of directional quantile classifiers. Theoretical results are stated in Section 3. We report the results of a simulation study in Section 4 and of a real data analysis in Section 5. Concluding remarks are given in Section 6. All proofs of theoretical results are reported in Appendix 6. A software implementation of our approach can be found in the package Qtools (Geraci, 2016), freely available on the Comprehensive R Archive Network (R Core Team, 2020).

2. Methods.

2.1. *Notation and definitions.* Let $\mathbf{X}^{(1)} = \left(X_1^{(1)}, X_2^{(1)}, \ldots, X_p^{(1)} \right)^\top$ and $\mathbf{X}^{(2)} = \left(X_1^{(2)}, X_2^{(2)}, \ldots, X_p^{(2)} \right)^\top$ denote two p-variate random variables with absolutely continuous distributions $F^{(1)}$ and $F^{(2)}$ defined on the same space $\mathcal{X} \subseteq \mathbb{R}^p$ for two populations $\Pi^{(1)}$ and $\Pi^{(2)}$, respectively. The marginal distributions of the components of $\mathbf{X}^{(k)}$ are denoted by $F_j^{(k)}$, for $j = 1, 2, \ldots$ and $k = 1, 2$. Further, $I(\cdot)$ denotes the indicator function which is equal to 1 if its argument is true, and 0 otherwise.

Our goal is to assign a new observation $\mathbf{y} = (y_1, y_2, \ldots, y_p)^\top$ to either $\Pi^{(1)}$ or $\Pi^{(2)}$ according to how *close* the point is to one or the other. In quantile-based classification (Hennig and Viroli, 2016a), the distance is first calculated for each component of $\mathbf{y}$ using the asymmetrically weighted loss

function

$$(1) \qquad \Phi^{(k)}(\theta; y_j) = \{\theta + (1 - 2\theta)I(y_j - Q_{X_j}^{(k)}(\theta) < 0)\}|y_j - Q_{X_j}^{(k)}(\theta)|$$

for $j = 1, 2, \ldots, p$ and $k = 1, 2$, where $Q_{X_j}^{(k)}(\theta)$ is the componentwise quantile at level $\theta \in (0, 1)$ for the kth population, which can be obtained by inversion of $F_j^{(k)}$. Subsequently, $\mathbf{y}$ is assigned to $\Pi^{(1)}$ if the discrepancy

$$(2) \qquad d(\mathbf{y}, \theta) = \sum_{j=1}^{p} \{\Phi^{(2)}(\theta; y_j) - \Phi^{(1)}(\theta; y_j)\}$$

is positive, and to $\Pi^{(2)}$ otherwise. The quantile classifier reduces to the componentwise median classifier of Hall, Titterington and Xue (2009) for $\theta = 0.5$. An extension of (2) to more than two populations is straightforward.

The classification rule based on (2) does not acknowledge the possible interdependence among the variables, since quantiles are obtained marginally for each variable. We address this limitation by using directional quantiles for multivariate data (Kong and Mizera, 2012). We now explain our idea informally and, in the next section, give a rigorous treatment.

Define $\mathbf{u}$ to be a vector with unit norm in $\mathbb{R}^p$. Throughout this paper, our focus will be on the *projected* random variables $\mathbf{u}^\top \mathbf{X}^{(k)} \equiv Z^{(k)}$, $k = 1, 2$, defined on $\mathcal{Z} \subseteq \mathbb{R}$. By assumption, the $Z^{(k)}$'s are continuous. We denote the corresponding distribution and density functions with $G^{(k)}(\cdot; \mathbf{u})$ and $g^{(k)}(\cdot; \mathbf{u})$, respectively.

Our goal is to develop a classifier where the quantities in (1) are opportunely redefined on the corresponding *projections* along $\mathbf{u}$ to capture the multivariate nature of the distributions, namely

$$(3) \quad \Phi^{(k)}(\theta; \mathbf{u}^\top \mathbf{y}) = \{\theta + (1 - 2\theta)I(\mathbf{u}^\top \mathbf{y} - Q_Z^{(k)}(\theta; \mathbf{u}) < 0)\}|\mathbf{u}^\top \mathbf{y} - Q_Z^{(k)}(\theta; \mathbf{u})|$$

for k = 1,2, where $Q_Z^{(k)}(\theta; \mathbf{u}) \equiv Q_{\mathbf{u}^\top \mathbf{X}}^{(k)}(\theta)$ is the θth quantile of $Z^{(k)}$. The latter is obtained by inverting $G^{(k)}$ and it can be recognised as the θth *directional quantile* of $\mathbf{X}^{(k)}$ in the direction $\mathbf{u}$ (Kong and Mizera, 2012).

By working with projections, we basically summarise a multivariate problem as a univariate one. Clearly, one difficulty to address is how many and which directions should be considered. To this end, we should note that not all the directions are equally useful for classification. To exemplify, consider Figure 1, which depicts bivariate normal samples from two independent populations centered at (1,1) and (3,3), respectively, and same variance. We want to assign the new observation $\mathbf{y} = (1.3, 3.4)^\top$ to one of the two populations. The log-density at $\mathbf{y}$ of two bivariate normal distributions with sample
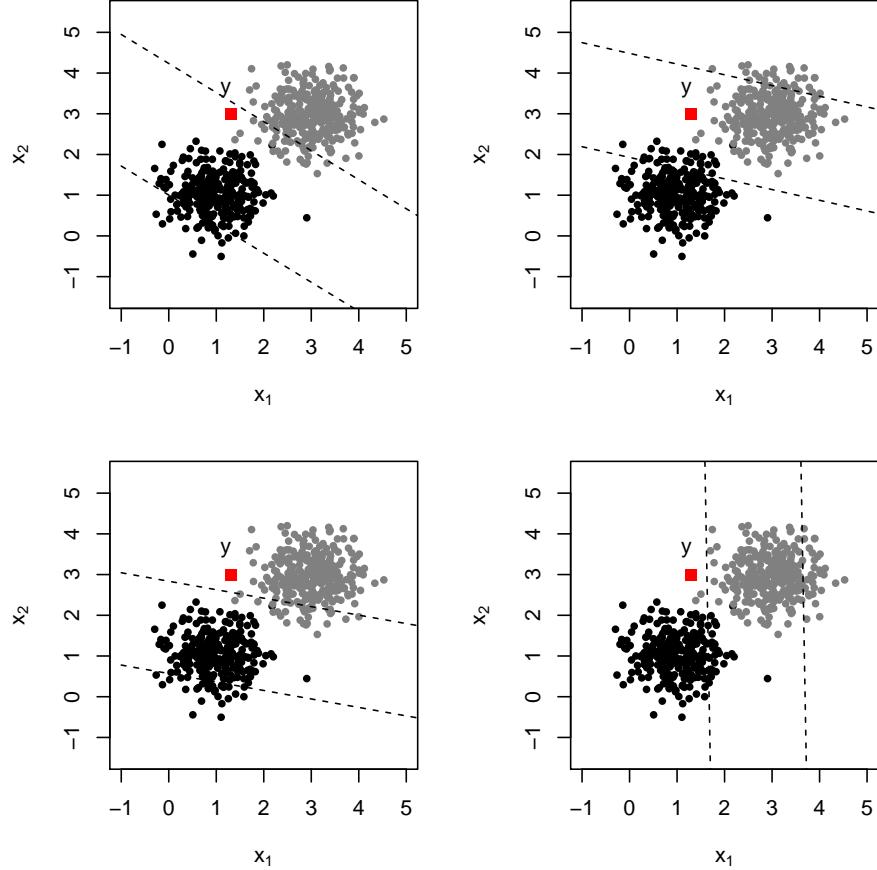
FIGURE 1. *Simulated data depicting bivariate normal samples from two independent distributions (black and grey dots). The red filled squares mark the point with coordinates* $(1.3, 3.4)$*, while dashed lines mark directions.*

means and covariance matrices separately estimated from the two samples, is -8.8 and -5.7, respectively. This suggests that $\mathbf{y}$ has been generated more likely from F_2 than from F_1.

Now compute $\Phi^{(k)}(0.9; \mathbf{u}^\top \mathbf{y})$, $k = 1, 2$, as in (3) for four normalised directions. The results are reported in Table 1. Based on a principle of minimum distance, we assign $\mathbf{y}$ to F_2, thus consistently with a maximum likelihood principle, for three, though not all four, directions.

TABLE 1
Distance $\Phi^{(k)}(0.9; \mathbf{u}^\top \mathbf{y})$, $k = 1, 2$, calculated for simulated data using four different directions $\mathbf{u}$.

$\mathbf{u}^\top$	$\Phi^{(1)}$	$\Phi^{(2)}$
$(-0.58, -0.81)$	0.27	0.01
$(0.25, 0.97)$	1.58	0.07
$(-0.20, -0.98)$	0.30	0.08
$(1.00, 0.02)$	0.03	0.23

2.2. *Directional quantile classifier.* Let $\vartheta = \{\theta_1, \theta_2, \ldots, \theta_R\}$ be a set of R distinct quantile levels on $(0, 1)$. Also, define the set $\upsilon_r = \{\mathbf{u}_{r1}, \mathbf{u}_{r2}, \ldots, \mathbf{u}_{rS_r}\}$ containing S_r normalised directions associated with θ_r, $r = 1, \ldots, R$, and let $\upsilon = \{\upsilon_1, \upsilon_2, \ldots, \upsilon_R\}$. (Note that for convenience one may set $S_r = S$ for $r = 1, \ldots, R$.)

As mentioned in the previous section, we need to be wary of particular directions that may lead us to a classification error. Therefore, we introduce weights ω_{rs} associated with each direction $\mathbf{u}_{rs}$ to decrease (or increase) their relative importance. Let $\boldsymbol{\omega} = (\omega_{11}, \ldots, \omega_{1S_1}, \ldots, \omega_{RS_R})^\top$ denote the vector of all such weights. We propose the discrepancy

$$(4) \qquad d(\mathbf{y}, \vartheta, \upsilon, \boldsymbol{\omega}) = \sum_{r=1}^{R} \sum_{s=1}^{S_r} \omega_{rs} \{\Phi^{(2)}(\theta_r; \mathbf{u}_{rs}^\top \mathbf{y}) - \Phi^{(1)}(\theta_r; \mathbf{u}_{rs}^\top \mathbf{y})\},$$

where $\Phi^{(k)}$ is defined in (3). Then our *directional quantile classifier* (DQC) assigns the observation $\mathbf{y}$ to $\Pi^{(1)}$ if $d(\mathbf{y}, \vartheta, \upsilon, \boldsymbol{\omega}) > 0$, or to $\Pi^{(2)}$ otherwise. Note that if $R = 1$, $S_r = p$, $\omega_{rs} = 1$, and $\upsilon = \{e_1, e_2, \ldots, e_p\}$ the standard basis in $\mathbb{R}^p$, then (4) reduces to (2).

A difficulty associated with the calculation of (4) is the selection of quantile levels, directions, and weights in the training data, say $\mathbf{x}$, that give the best performance on the test data, say $\mathbf{y}$. For some prior probabilities π_1 and π_2, let

$$\psi(\mathbf{x}, \vartheta, \upsilon, \boldsymbol{\omega}) = \pi_1 \int_{\mathcal{X}} I\{d(\mathbf{x}, \vartheta, \upsilon, \boldsymbol{\omega}) > 0\} \, dF^{(1)}(\mathbf{x})$$

$$(5) \qquad + \pi_2 \int_{\mathcal{X}} I\{d(\mathbf{x}, \vartheta, \upsilon, \boldsymbol{\omega}) \leq 0\} \, dF^{(2)}(\mathbf{x})$$

denote the population probability of correct classification by the DQC. Note that maximising (5) is equivalent to minimising the theoretical misclassification rate. For any given level θ and direction $\mathbf{u}$, the optimal misclassification rate is obtained when

$$\pi_1 \int_{\mathcal{X}} \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) \, dF^{(1)}(\mathbf{x}) < \pi_1 \int_{\mathcal{X}} \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) \, dF^{(1)}(\mathbf{x})$$

and
$$\pi_2 \int_{\mathcal{X}} \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) \, dF^{(2)}(\mathbf{x}) < \pi_2 \int_{\mathcal{X}} \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) \, dF^{(2)}(\mathbf{x}),$$

which is equivalent to minimise

$$\pi_1 \int_{\mathcal{X}} \left\{ \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) - \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) \right\} \, dF^{(1)}(\mathbf{x})$$

(6)
$$+ \pi_2 \int_{\mathcal{X}} \left\{ \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) - \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) \right\} \, dF^{(2)}(\mathbf{x}).$$

In the general problem with K populations, the minimum misclassification rate is obtained when

(7)
$$\sum_{k=1}^{K} \pi_k \int_{\mathcal{X}} \Phi^{(k)}(\theta; \mathbf{u}^\top \mathbf{x}) \, dF^{(k)}(\mathbf{x}) < \sum_{k=1}^{K} \pi_k \int_{\mathcal{X}} \min_{k' \neq k} \Phi^{(k')}(\theta; \mathbf{u}^\top \mathbf{x}) \, dF^{(k)}(\mathbf{x}).$$

Let $\Delta^{(k)}(\mathbf{x}, \theta, \mathbf{u}) = \Phi^{(k)}(\theta; \mathbf{u}^\top \mathbf{x}) - \min_{k' \neq k} \Phi^{(k')}(\theta; \mathbf{u}^\top \mathbf{x})$. Given a sample of n observations $\mathbf{x}_i$ and corresponding class labels $\ell_i \in \{1, 2, \ldots, K\}$, we aim to solve

(8)
$$\min_{\vartheta, \upsilon, \boldsymbol{\omega}} \sum_{k=1}^{K} \sum_{i:\ell_i=k} \sum_{r=1}^{R} \sum_{s=1}^{S_r} \omega_{rs} \Delta^{(k)}(\mathbf{x}_i, \theta_r, \mathbf{u}_{rs}).$$

Problem (8) may seem daunting, but luckily we can solve for $\boldsymbol{\omega}$ rather easily. Given ϑ and υ, problem (8) is linear with unit-norm constraints and can be minimised by using the Lagrange multiplier method. This problem has a closed-form solution given by $\hat{\boldsymbol{\omega}} = (\hat{\omega}_{11}, \ldots, \hat{\omega}_{1S_1}, \ldots, \hat{\omega}_{RS_R})^\top$ with generic rsth element

$$\hat{\omega}_{rs} = \frac{\tilde{\Delta}_{rs}}{\sqrt{\sum_{r=1}^{R} \sum_{s=1}^{S_r} \tilde{\Delta}_{rs}^2}},$$

where $\tilde{\Delta}_{rs} = \sum_{k=1}^{K} \sum_{i:\ell_i=k} \Delta^{(k)}(\mathbf{x}_i, \theta_r, \mathbf{u}_{rs})$.

We now turn to how to choose directions and quantile levels. A crude solution would consist in doing a multidimensional grid search on $p + 1$ dimensions. However, such a solution would become computationally prohibitive even at modest values of p. Thankfully, we are able to mitigate the computational cost of a naïve numerical solution with some theoretical results (Section 3); in particular, with Theorem 1, which guarantees that for each projection there exists (at least) a quantile that leads to the optimal Bayes misclassification probability, and Theorem 2, which, conversely,

identifies the best direction for a given quantile level. Unfortunately, a theoretical result for the simultaneous optimisation with respect to θ and $\mathbf{u}$ does not exist. Nevertheless, we show that our DQC is asymptotically optimal (i.e. the misclassification rate goes to zero) when the number of directions increases with p and n (Theorem 3) under certain assumptions.

In summary, there are different possible approaches including randomly selecting one or more directions and using the optimal quantile levels associated with those directions; or spanning a grid of quantile levels and using the optimal directions associated with those quantiles. After some empirical investigation, we found that a strategy that gives satisfactory results in different settings is as follows. First, we define a grid of θ values spanning the unit interval and, for each of these values, randomly draw a set of normalised directions from the hyperplane that is identified as optimal according to Theorem 2. The performance of a DQC based on each single θ value is evaluated using five-fold cross-validation. In the end, we use a single quantile level (optimal according to cross-validation), with the corresponding directions sampled from the optimal hyperplane. In particular, this strategy improves over the use of an asymptotically optimal quantile level when n is small. Moreover, when p is not too large, a similar strategy can be used to select an approximately optimal hyperplane.

3. Theoretical results. In this section, we present theoretical results concerning our DQC. The proofs of lemmas and theorems are reported in the Appendix.

3.1. *Optimal quantile level θ.* We derive the theoretical rate of correct classification as a function of θ, for given $\mathbf{u}$. We assume $K = 2$ populations, although results can be generalised to $K > 2$.

LEMMA 1. For given $\mathbf{u}$, let $Q_\alpha(\theta; \mathbf{u}) = \min\{Q_Z^{(1)}(\theta; \mathbf{u}), Q_Z^{(2)}(\theta; \mathbf{u})\}$ with corresponding inverse $G_\alpha(\cdot; \mathbf{u})$, density $g_\alpha(\cdot; \mathbf{u})$, and prior probability of correct classification π_α, and let $Q_\beta(\theta; \mathbf{u}) = \max\{Q_Z^{(1)}(\theta; \mathbf{u}), Q_Z^{(2)}(\theta; \mathbf{u})\}$ with corresponding inverse $G_\beta(\cdot; \mathbf{u})$, density $g_\beta(\cdot; \mathbf{u})$, and prior probability of correct classification π_β. The probability of correct classification of the directional quantile classifier is

$$(9) \qquad \psi(\theta) = \pi_\alpha G_\alpha(\tilde{Q}(\theta; \mathbf{u}); \mathbf{u}) + \pi_\beta\{1 - G_\beta(\tilde{Q}(\theta; \mathbf{u}); \mathbf{u})\}$$

where $\tilde{Q}(\theta; \mathbf{u}) = \theta Q_\alpha(\theta; \mathbf{u}) + (1 - \theta)Q_\beta(\theta; \mathbf{u})$. Analogously, the theoretical misclassification rate is

$$(10) \qquad 1 - \psi(\theta) = \pi_\alpha\{1 - G_\alpha(\tilde{Q}(\theta; \mathbf{u}); \mathbf{u})\} + \pi_\beta G_\beta(\tilde{Q}(\theta; \mathbf{u}); \mathbf{u}).$$
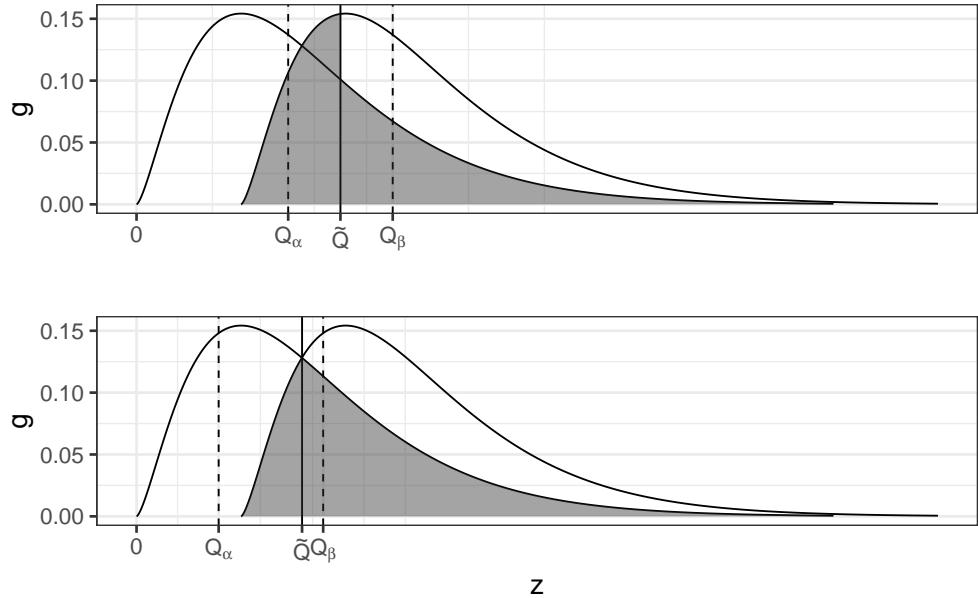
FIGURE 2. *Misclassification probability (shaded grey area) with two location-shifted skewed distributions according to median classifier (upper panel) and the optimal quantile classifier (lower panel).*

THEOREM 1. Assume that the density functions $g_\alpha(z; \mathbf{u})$ and $g_\beta(z; \mathbf{u})$ exist for z and are nonzero on the same compact domain $\mathcal{Z}$. Further assume that there is a point z_0 with $\pi_\alpha g_\alpha(z_0; \mathbf{u}) = \pi_\beta g_\beta(z_0; \mathbf{u})$ so that $\pi_\alpha g_\alpha(z; \mathbf{u}) > \pi_\beta g_\beta(z; \mathbf{u})$ for z on one side of z_0 and $\pi_\alpha g_\alpha(z; \mathbf{u}) < \pi_\beta g_\beta(z; \mathbf{u})$ for z on the other side of z_0. Then the quantile classifier using the quantile $\tilde{Q}(\theta; \mathbf{u})$ that minimises the theoretical misclassification probability achieves the optimal Bayes misclassification probability.

The consistency of the classifier may be illustrated with an example. Consider a two class decision problem where one population is a location-shift version of the other. Figure 2 shows two distributions which have both the same right skewness. The quantiles $Q_\alpha(\theta)$ and $Q_\beta(\theta)$ are marked by dashed lines. The median classifier (Hall, Titterington and Xue, 2009) in the upper panel leads to a non-optimal misclassification probability equal to 0.30. However, the misclassification probability is reduced to 0.28 by setting $\theta = 0.202$.

3.2. *Optimal direction* $\mathbf{u}$. The next lemma and theorem give the optimal direction that minimises the misclassification rate at a given θ.

LEMMA 2. Let z be a realisation of either $Z^{(1)}$ or $Z^{(2)}$, then

$$\Phi^{(2)}(\theta; z) - \Phi^{(1)}(\theta; z) \le Q_Z^{(2)}(\theta) - Q_Z^{(1)}(\theta),$$

where $\Phi^{(k)}(\theta; z) = \theta \max(\eta^{(k)}, 0) + (1 - \theta) \max(-\eta^{(k)}, 0)$ and $\eta^{(k)} = z - Q_Z^{(k)}(\theta)$, $k = 1, 2$.

THEOREM 2. Let $\mathbf{W} = (W_1, W_2, \ldots, W_p)^\top$ be a p-variate random variable such that $Q_{W_j}(\theta) = 0$, for $j = 1, \ldots, p$, and let $\boldsymbol{\mu}^{(k)} = (\mu_1^{(k)}, \mu_2^{(k)}, \ldots, \mu_p^{(k)})^\top$ be a vector of constants, $k = 1, 2$. We assume that $\mathbf{X}^{(k)} = \mathbf{W} + \boldsymbol{\mu}^{(k)}$ and its probability distribution function is $F^{(k)}$, for $k = 1, 2$. Moreover, assume that $Q_Z^{(2)}(\theta; \mathbf{u}) > Q_Z^{(1)}(\theta; \mathbf{u})$, where $Q_Z^{(k)}(\theta; \mathbf{u})$ is the θ-quantile of $Z^{(k)} \equiv \mathbf{u}^\top \mathbf{X}^{(k)}$. (Notice that there is no loss of generality with this assumption since the case $Q_Z^{(2)}(\theta; \mathbf{u}) \le Q_Z^{(1)}(\theta; \mathbf{u})$ can be reformulated as $Q_Z^{(2)}(\theta; -\mathbf{u}) > Q_Z^{(1)}(\theta; -\mathbf{u})$.) Under these assumptions, the normalised direction $\mathbf{u}$ that minimises the misclassification error (6) is

$$(11) \qquad \hat{\mathbf{u}} = \frac{\boldsymbol{\mu}^{(2)} - \boldsymbol{\mu}^{(1)}}{\|\boldsymbol{\mu}^{(2)} - \boldsymbol{\mu}^{(1)}\|}$$

The generalization of Theorem 2 to $K > 2$ populations involves $K(K-1)/2$ optimal directions for each of all the possible pairwise comparisons.

3.3. *Asymptotic misclassification rate*. In this section, we show that under certain assumptions, the correct classification probability converges to unity when the number of dimensions grows to infinity along with the sample size and the number of projections. The proof is built following a strategy similar to that used in Hall, Titterington and Xue (2009, Theorem 2), although our premises start from milder assumptions. In particular the projections are not required to obey the "$\psi - mixing$ condition" (Bradley, 2005), which is rather strict in practice. Our theorem is developed for any $\theta_r \in (0, 1)$, unit weights $\omega_{rs} = 1$, and $R = 1$. Thus, the asymptotic result holds for sub-components of the summation in (8), which are then weighted and summed to minimise the misclassification rate. Hence, the overall criterion inherits the optimal properties of its additive components.

As we did with the theorems in the previous sections, we present this theorem for $K = 2$ classes. Its extension to $K > 2$ classes requires contrasting each class against the remaining $K - 1$ classes, consistently with (7).

THEOREM 3.   Consider a set of directions $\upsilon = \{\mathbf{u}_1, \ldots, \mathbf{u}_S\}$ sampled from a unit p-sphere and let $n = \max(n_1, n_2)$, with n_1 and n_2 denoting the sample sizes of the two groups in the training set. Assume

(i) For a constant $A_1 > 0$, $S \geq A_1 n$.
(ii) The p variables $X_1^{(k)}, X_2^{(k)}, \ldots, X_p^{(k)}$ have each the same distribution as $W_1 + \mu_1^{(k)}, W_2 + \mu_2^{(k)}, \ldots, W_p + \mu_p^{(k)}$, respectively. Moreover, $Q_{W_j}(\theta) = 0 \; \forall j$ and $\sup\limits_{j \geq 1} \mathrm{Var}(W_j) = A_2 < +\infty$.
(iii) The first moments of the projections are uniformly bounded in a strong sense. This implies that $\forall c > 0$ and $\forall \mathbf{u}$, $\exists \mathbf{v}$ with $|\mathbf{u}^\top \mathbf{v}| > c$ such that

$$\inf_{s \geq 1} \inf_{|\mathbf{u}_s^\top \mathbf{v}| > c} \theta \, \mathrm{E} \, |\mathbf{u}_s^\top \mathbf{W} + \mathbf{u}_s^\top \mathbf{v}| - (1 - \theta) \, \mathrm{E} \, |\mathbf{u}_s^\top \mathbf{W}| > 0.$$

(iv) For some $\epsilon > 0$, the proportion of values $s \in \{1, 2, \ldots, S\}$ for which

$$|\theta \mathbf{u}_s^\top \boldsymbol{\mu}^{(2)} - (1 - \theta) \mathbf{u}_s^\top \boldsymbol{\mu}^{(1)}| > \epsilon$$

multiplied by $n^{1/2}$, say $n^{1/2} \sharp \mathcal{K}_\epsilon$, is of larger order than S, which means $S \left(n^{1/2} \sharp \mathcal{K}_\epsilon \right)^{-1}$ goes to zero as n and S increase.

Under the previous assumptions, the directional quantile classifier $\mathcal{C}$ based on

$$d(\mathbf{y}, \theta, \upsilon, \boldsymbol{\omega}) = \sum_{s=1}^{S} \{ \Phi^{(2)}(\theta; \mathbf{u}_s^\top \mathbf{y}) - \Phi^{(1)}(\theta; \mathbf{u}_s^\top \mathbf{y}) \},$$

makes the correct choice asymptotically. More specifically, as $p \to \infty$, the classifier $\mathcal{C}$ makes the correct decision with probability

$$P^{(1)}\{\mathcal{C}(\mathbf{Y}) = 1\} + P^{(2)}\{\mathcal{C}(\mathbf{Y}) = 2\}$$

converging to 1 if both n_1 and n_2 diverge with p, where $P^{(k)}$, $k = 1, 2$ denotes the probability computed under the assumption that Y is drawn from population k.

4. Simulation study.   We assessed the performance of the proposed classifier in a simulation study under three scenarios with two populations. In the first scenario, observations were generated independently from a multivariate Student's t distribution with 3 degrees of freedom, with either uncorrelated or correlated variables. In the second scenario, observations were generated as in the first scenario, but each variable was subsequently transformed according to $x \mapsto \log(|x|)$ to induce asymmetry. In both cases, the two populations differed by a location shift equal to 0.4. Finally, in

the third scenario observations were generated as in the first scenario, but each variable was subsequently transformed according to $x \mapsto \log(|x|)$ or to $x \mapsto -\log(|x|)$ depending on whether observations belonged to one or the other population, respectively.

Data were generated for each combination of overall sample size $n \in \{50, 100, 500\}$ (with $n/2$ observations in each class) and dimension $p \in \{10, 50, 100, 500\}$. All in all, this resulted in $3 \times 2 \times 3 \times 4 = 72$ simulation cases. The scale matrix used in the multivariate t distribution with correlated variables was generated randomly for each p using the function rcorrmatrix with default settings as provided in the package clusterGeneration (Qiu and Joe, 2015; Joe, 2006). This resulted in non-constant pairwise correlations on the interval $(-1, 1)$. Observations in the training and test datasets were generated in the same way. Data generation under each setting was replicated 100 times.

We compared the directional quantile classifier (DQC) in terms of misclassification rate on the test data with that of the centroid classifier (Centroid) (Tibshirani et al., 2002), median classifier (Median) (Hall, Titterington and Xue, 2009), componentwise quantile classifier (CQC) (Hennig and Viroli, 2016a), ensemble quantile classifier (EQC) (Lai and McLeod, 2020), Fisher's linear discriminant analysis (LDA), k-nearest neighbour (KNN) (Cover and Hart, 1967), penalised logistic regression (PLR) (Park and Hastie, 2008), support vector machines (SVM) (Cortes and Vapnik, 1995; Wang, Zhu and Zou, 2008), and naïve Bayes classifier (Bayes) (Hand and Yu, 2001). Tuning parameters for PLR, KNN, and SVM where selected using cross-validation. For the CQC, the Galton correction was used to reduce skewness and optimal quantile was selected by minimising the error rate on the training set (Hennig and Viroli, 2016a).

We used the package Qtools (Geraci, 2016, 2020) for the directional quantile classifier; the package quantileDA (Hennig and Viroli, 2016b) for the centroid, median and componentwise quantile classifiers; the package eqc (Lai and McLeod, 2019) for the ensemble quantile classifier; the package MASS (Venables and Ripley, 2002) for linear discriminant analysis; the package class (Venables and Ripley, 2002) for k-nearest neighbour; the package e1071 (Meyer et al., 2019) for support vector machines and Bayes classifiers; and the package stepPlr (Park and Hastie, 2018) for penalised logistic regression. All analyses were carried out in R version 4.0.0 (R Core Team, 2020).

The misclassification rates averaged over 100 replications for all simulation cases are reported in Tables 2-4. The results indicate that the performance of our proposed classifier improves as n and p increase, in agreement with

the theoretical results. In the first two scenarios, our classifier outperforms the competitors in both scenarios when variables are uncorrelated. When variables are correlated, the proposed classifier still performs very well, even if it is not uniformly the best. In the third scenario where class distributions have different shapes, the performance of our classifier is often, but not always, the best.

5. Clinical trial on Crohn's disease. We analyse data from a matched case-control study in first-degree relatives (FDRs) of Crohn's disease (CD) patients originally published by Sorrentino et al. (2014). The goal of the study was to identify asymptomatic FDRs with early CD signs using several intestinal inflammatory markers. The latter included hemoglobin, erythrocyte sedimentation rate, C-reactive protein, fecal calprotectin, and average mature ileum score. In our analysis, we grouped subjects into 2 classes, one with signs of inflammation ($n_1 = 9$ subjects with early or frank CD) and one with normal values of markers ($n_2 = 26$ subjects with no signs of inflammation, including healthy controls). In a separate analysis, we augment the dataset with 45 artificial markers generated from independent standard normal distributions to investigate the impact of uninformative noise on the performance of the DQC. We approach data analysis with leave-one-out validation and evaluate the misclassification rate as the proportion of subjects that are misclassified when each is left out of analysis.

We estimated the classification error for all the classifiers as included in our simulation study (Section 4). The results are reported in Table 5. The proposed DQC outperforms its competitors in both the original ($p = 5$) and noisy ($p = 50$) versions of the dataset.

6. Conclusions. We proposed directional quantile classifiers whose predictive ability is consistently good in both simulation and real data studies, on small and large dimensional classification problems. In particular, the empirical results show that our approach either outperforms its competitors or, when this is not the case, its performance is still in the ballpark of that of the best classifiers. Such a reliable behaviour across different scenarios is not shared by the other selected classifiers. Moreover, the directional quantile classifiers enjoy optimal theoretical properties under certain assumptions.

A limitation of the approach is that the number of directions needed to span a p-sphere with a regular grid becomes prohibitive already at modest values of p. On the other hand, our theoretical results indicate that one can sample directions from an optimal hyperplane, thus reducing the computational burden, but not at the expense of the classifier's performance. Our strategy allows us to balance the importance of quantile levels and directions

TABLE 2

Misclassification rates averaged over 100 replications for ten classifiers (DQC, directional quantile classifier; Centroid, centroid classifier; Median, median classifier; CQC, componentwise quantile classifier; EQC, ensemble quantile classifier; LDA, linear discriminant analysis; KNN, k-nearest neighbour; PLR, penalised logistic regression; SVM, support vector machines; Bayes, naïve Bayes) in the first scenario where populations have symmetric distributions.

	Uncorrelated				Correlated			
Dimension p	10	50	100	500	10	50	100	500
Sample size $n = 50$								
DQC	0.334	0.187	0.120	0.020	0.315	0.202	0.128	0.028
Centroid	0.355	0.232	0.168	0.049	0.349	0.277	0.189	0.059
Median	0.372	0.230	0.153	0.043	0.357	0.252	0.170	0.047
CQC	0.362	0.273	0.220	0.180	0.367	0.284	0.222	0.177
EQC	0.373	0.240	0.172	0.044	0.339	0.253	0.168	0.055
LDA	0.365	0.382	0.295	0.313	0.245	0.252	0.308	0.339
KNN	0.362	0.287	0.263	0.212	0.360	0.300	0.271	0.211
PLR	0.348	0.199	0.134	0.023	0.275	0.154	0.103	0.025
SVM	0.413	0.252	0.140	0.046	0.401	0.263	0.140	0.049
Bayes	0.390	0.333	0.302	0.225	0.395	0.327	0.287	0.237
Sample size $n = 100$								
DQC	0.306	0.145	0.089	0.015	0.283	0.146	0.076	0.017
Centroid	0.325	0.181	0.114	0.025	0.331	0.210	0.132	0.036
Median	0.334	0.194	0.129	0.032	0.339	0.211	0.136	0.033
CQC	0.343	0.213	0.151	0.076	0.341	0.223	0.157	0.092
EQC	0.337	0.213	0.135	0.039	0.329	0.193	0.138	0.040
LDA	0.338	0.236	0.393	0.182	0.226	0.055	0.105	0.240
KNN	0.346	0.214	0.184	0.102	0.325	0.223	0.191	0.108
PLR	0.329	0.182	0.113	0.019	0.240	0.092	0.056	0.020
SVM	0.370	0.176	0.106	0.032	0.382	0.153	0.069	0.034
Bayes	0.367	0.284	0.227	0.179	0.371	0.291	0.242	0.179
Sample size $n = 500$								
DQC	0.286	0.128	0.069	0.010	0.263	0.126	0.058	0.010
Centroid	0.300	0.145	0.080	0.014	0.288	0.154	0.077	0.016
Median	0.320	0.173	0.101	0.020	0.315	0.178	0.099	0.019
CQC	0.327	0.176	0.108	0.023	0.320	0.183	0.103	0.025
EQC	0.324	0.177	0.106	0.021	0.296	0.149	0.085	0.021
LDA	0.302	0.160	0.106	0.367	0.196	0.027	0.000	0.036
KNN	0.326	0.163	0.098	0.018	0.283	0.166	0.097	0.022
PLR	0.301	0.161	0.104	0.013	0.194	0.044	0.018	0.008
SVM	0.300	0.142	0.084	0.018	0.238	0.066	0.020	0.014
Bayes	0.329	0.198	0.145	0.077	0.326	0.199	0.140	0.076

TABLE 3
Misclassification rates averaged over 100 replications for ten classifiers (DQC, directional quantile classifier; Centroid, centroid classifier; Median, median classifier; CQC, componentwise quantile classifier; EQC, ensemble quantile classifier; LDA, linear discriminant analysis; KNN, k-nearest neighbour; PLR, penalised logistic regression; SVM, support vector machines; Bayes, naïve Bayes) in the second scenario where populations have distributions with same skewness.

	Uncorrelated				Correlated			
Dimension p	10	50	100	500	10	50	100	500
Sample size $n = 50$								
DQC	0.313	0.170	0.096	0.052	0.306	0.169	0.095	0.059
Centroid	0.323	0.212	0.145	0.097	0.330	0.220	0.140	0.113
Median	0.334	0.206	0.147	0.105	0.334	0.215	0.140	0.106
CQC	0.350	0.235	0.187	0.234	0.360	0.245	0.193	0.248
EQC	0.340	0.180	0.089	0.015	0.333	0.178	0.102	0.019
LDA	0.317	0.383	0.238	0.228	0.315	0.397	0.233	0.237
KNN	0.382	0.275	0.210	0.064	0.364	0.282	0.213	0.078
PLR	0.313	0.183	0.095	0.002	0.322	0.186	0.098	0.004
SVM	0.330	0.224	0.150	0.021	0.332	0.240	0.150	0.021
Bayes	0.378	0.281	0.220	0.153	0.377	0.272	0.223	0.161
Sample size $n = 100$								
DQC	0.293	0.129	0.060	0.012	0.280	0.118	0.058	0.010
Centroid	0.310	0.168	0.106	0.057	0.307	0.161	0.104	0.065
Median	0.328	0.177	0.110	0.067	0.316	0.171	0.106	0.071
CQC	0.317	0.173	0.110	0.135	0.314	0.160	0.113	0.137
EQC	0.310	0.135	0.071	0.006	0.296	0.126	0.062	0.006
LDA	0.301	0.218	0.374	0.084	0.281	0.203	0.395	0.089
KNN	0.358	0.242	0.188	0.047	0.353	0.256	0.186	0.045
PLR	0.300	0.163	0.079	0.001	0.284	0.153	0.078	0.001
SVM	0.318	0.177	0.089	0.005	0.333	0.168	0.079	0.005
Bayes	0.330	0.229	0.167	0.095	0.334	0.225	0.166	0.098
Sample size $n = 500$								
DQC	0.273	0.097	0.038	0.000	0.265	0.097	0.035	0.000
Centroid	0.282	0.119	0.059	0.007	0.275	0.116	0.056	0.008
Median	0.295	0.128	0.074	0.017	0.286	0.124	0.069	0.019
CQC	0.272	0.099	0.053	0.019	0.261	0.092	0.050	0.018
EQC	0.267	0.088	0.035	0.001	0.244	0.079	0.032	0.001
LDA	0.279	0.116	0.060	0.374	0.266	0.114	0.057	0.372
KNN	0.323	0.206	0.140	0.016	0.310	0.207	0.140	0.015
PLR	0.279	0.121	0.060	0.000	0.266	0.119	0.056	0.000
SVM	0.283	0.109	0.046	0.000	0.274	0.107	0.044	0.000
Bayes	0.273	0.129	0.080	0.020	0.266	0.125	0.077	0.021

TABLE 4

Misclassification rates averaged over 100 replications for ten classifiers (DQC, directional quantile classifier; Centroid, centroid classifier; Median, median classifier; CQC, componentwise quantile classifier; EQC, ensemble quantile classifier; LDA, linear discriminant analysis; KNN, k-nearest neighbour; PLR, penalised logistic regression; SVM, support vector machines; Bayes, naïve Bayes) in the third scenario where populations have distributions with opposite skewness.

	Uncorrelated				Correlated			
Dimension p	10	50	100	500	10	50	100	500
				Sample size $n = 50$				
DQC	0.199	0.171	0.166	0.159	0.237	0.172	0.110	0.023
Centroid	0.228	0.176	0.169	0.160	0.362	0.265	0.190	0.066
Median	0.321	0.283	0.273	0.264	0.359	0.240	0.166	0.045
Quantile	0.236	0.112	0.087	0.073	0.371	0.279	0.215	0.181
EQC	0.315	0.279	0.256	0.234	0.349	0.239	0.162	0.051
LDA	0.277	0.450	0.253	0.161	0.248	0.270	0.298	0.349
KNN	0.277	0.213	0.192	0.173	0.365	0.284	0.214	0.074
PLR	0.259	0.252	0.213	0.173	0.317	0.189	0.100	0.003
SVM	0.231	0.175	0.170	0.159	0.338	0.240	0.157	0.018
Bayes	0.229	0.132	0.123	0.106	0.373	0.288	0.227	0.145
				Sample size $n = 100$				
DQC	0.188	0.162	0.165	0.166	0.195	0.133	0.071	0.016
Centroid	0.214	0.167	0.167	0.166	0.336	0.212	0.128	0.033
Median	0.314	0.287	0.285	0.275	0.341	0.215	0.132	0.032
Quantile	0.214	0.086	0.071	0.058	0.346	0.226	0.159	0.091
EQC	0.296	0.254	0.256	0.241	0.334	0.198	0.132	0.039
LDA	0.237	0.300	0.456	0.174	0.222	0.056	0.110	0.238
KNN	0.246	0.204	0.191	0.184	0.351	0.251	0.189	0.044
PLR	0.234	0.252	0.235	0.187	0.284	0.156	0.079	0.001
SVM	0.221	0.169	0.170	0.166	0.323	0.174	0.083	0.004
Bayes	0.187	0.112	0.105	0.095	0.343	0.232	0.169	0.092
				Sample size $n = 500$				
DQC	0.182	0.166	0.162	0.159	0.177	0.111	0.053	0.010
Centroid	0.209	0.170	0.165	0.160	0.283	0.153	0.076	0.015
Median	0.312	0.288	0.283	0.279	0.312	0.178	0.099	0.020
Quantile	0.203	0.069	0.052	0.041	0.316	0.182	0.104	0.024
EQC	0.282	0.249	0.241	0.236	0.293	0.148	0.085	0.021
LDA	0.212	0.194	0.212	0.474	0.194	0.027	0.000	0.032
KNN	0.193	0.173	0.176	0.178	0.311	0.206	0.138	0.015
PLR	0.213	0.201	0.226	0.237	0.266	0.118	0.057	0.001
SVM	0.209	0.172	0.167	0.163	0.274	0.107	0.043	0.001
Bayes	0.164	0.102	0.096	0.086	0.269	0.126	0.080	0.019

Table 5

Cross-validated estimates of the misclassification rates for the Crohn's disease dataset ($p = 5$) and its noisy version ($p = 50$) using ten classifiers (DQC, directional quantile classifier; Centroid, centroid classifier; Median, median classifier; CQC, componentwise quantile classifier; EQC, ensemble quantile classifier; LDA, linear discriminant analysis; KNN, k-nearest neighbour; PLR, penalised logistic regression; SVM, support vector machines; Bayes, naïve Bayes).

	$p = 5$	$p = 50$
DQC	0.229	0.229
Centroid	0.286	0.286
Median	0.400	0.400
CQC	0.314	0.343
EQC	0.314	0.314
LDA	0.257	0.543
KNN	0.371	0.343
PLR	0.286	0.343
SVM	0.257	0.257
Bayes	0.286	0.257

used for classification by means of weights, which can be optimised using a convenient closed-form expression.

APPENDIX A - PROOFS OF THEOREMS

A.1. Proofs of Lemma 1 and Theorem 1.

PROOF. The proofs of Lemma 1 and Theorem 1 follow the arguments given in Hennig and Viroli (2016a, Supplementary Material). Here, we briefly sketch the main idea. The optimal value θ that minimises the theoretical misclassification probability can be obtained by setting the first derivative of (10) to zero, from which

$$\pi_\alpha g_\alpha\{\tilde{Q}(\theta; \mathbf{u})\} = \pi_\beta g_\beta\{\tilde{Q}(\theta; \mathbf{u})\}.$$

By assumption, there exists $\theta \in (0,1)$ such that $\tilde{Q}(\theta; \mathbf{u}) = z_0$. Hence, the identity above is satisfied because $Q_\alpha(\theta; \mathbf{u})$ and $Q_\beta(\theta; \mathbf{u})$ are continuous functions of θ that converge to the lower and upper bound of $\mathcal{Z}$ for θ approaching either 0 or 1, respectively. Furthermore, under the assumptions of Theorem 1, the optimal Bayesian classifier has a single decision boundary at $\tilde{Q}(\theta; \mathbf{u})$. $\square$

A.2. Proof of Lemma 2.

PROOF. Without loss of generality, assume $Q_Z^{(1)}(\theta) \leq Q_Z^{(2)}(\theta)$. Let $\Delta(\theta; z) = \Phi^{(2)}(\theta; z) - \Phi^{(1)}(\theta; z)$ and consider three possible, distinct cases: $z \leq Q_Z^{(1)}(\theta)$, $Q_Z^{(1)}(\theta) < z < Q_Z^{(2)}(\theta)$, and $Q_Z^{(2)}(\theta) \leq z$.

If $z \leq Q_Z^{(1)}(\theta)$, then

$$\Delta(\theta; z) = (1-\theta)\{Q_Z^{(2)}(\theta) - z\} - (1-\theta)\{Q_Z^{(1)}(\theta) - z\}$$
$$= (1-\theta)\{Q_Z^{(2)}(\theta) - Q_Z^{(1)}(\theta)\} \leq Q_Z^{(2)}(\theta) - Q_Z^{(1)}(\theta)$$

by definition. If $Q_Z^{(1)}(\theta) < z < Q_Z^{(2)}(\theta)$, then

$$\Delta(\theta; z) = (1-\theta)\{Q_Z^{(2)}(\theta) - z\} - \theta\{z - Q_Z^{(1)}(\theta)\}$$
$$= \theta\{Q_Z^{(1)}(\theta) - Q_Z^{(2)}(\theta)\} + Q_Z^{(2)}(\theta) - z$$
$$\leq \theta Q_Z^{(1)}(\theta) - Q_Z^{(2)}(\theta) \leq Q_Z^{(2)}(\theta) - Q_Z^{(1)}(\theta).$$

Finally, if $Q_Z^{(2)}(\theta) \leq z$, then

$$\Delta(\theta; z) = \theta\{z - Q_Z^{(2)}(\theta)\} - \theta\{z - Q_Z^{(1)}(\theta)\}$$
$$\leq Q_Z^{(2)}(\theta) - Q_Z^{(1)}(\theta).$$

$\square$

A.3. Proof of Theorem 2.

PROOF. By Lemma 2, the differences $\Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) - \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x})$ and $\Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) - \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x})$ are upper bounded by $Q_Z^{(2)}(\theta; \mathbf{u}) - Q_Z^{(1)}(\theta; \mathbf{u})$ since $Q_Z^{(2)}(\theta; \mathbf{u}) > Q_Z^{(1)}(\theta; \mathbf{u})$. Therefore the quantity in (6), which is to be minimised with respect to $\mathbf{u}$ subject to $\|\mathbf{u}\| = 1$, is uniformly bounded above by

$$
\begin{aligned}
&\pi_1 \int_{\mathcal{X}} \left\{ \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) - \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) \right\} \, \mathrm{d}F^{(1)}(\mathbf{x}) \\
&+ \pi_2 \int_{\mathcal{X}} \left\{ \Phi^{(2)}(\theta; \mathbf{u}^\top \mathbf{x}) - \Phi^{(1)}(\theta; \mathbf{u}^\top \mathbf{x}) \right\} \, \mathrm{d}F^{(2)}(\mathbf{x}) + \lambda(\mathbf{u}^\top \mathbf{u} - 1) \\
&\leq Q_Z^{(2)}(\theta; \mathbf{u}) - Q_Z^{(1)}(\theta; \mathbf{u}) + \lambda(\mathbf{u}^\top \mathbf{u} - 1) \\
&= (Q_W(\theta; \mathbf{u}) + \mathbf{u}^\top \boldsymbol{\mu}^{(2)}) - (Q_W(\theta; \mathbf{u}) + \mathbf{u}^\top \boldsymbol{\mu}^{(1)}) + \lambda(\mathbf{u}^\top \mathbf{u} - 1) \\
&= \mathbf{u}^\top(\boldsymbol{\mu}^{(2)} - \boldsymbol{\mu}^{(1)}) + \lambda(\mathbf{u}^\top \mathbf{u} - 1).
\end{aligned}
$$

To find $\mathbf{u}$, we minimise the Lagrangian function $\mathbf{u}^\top(\boldsymbol{\mu}^{(2)} - \boldsymbol{\mu}^{(1)}) + \lambda(\mathbf{u}^\top \mathbf{u} - 1)$ which has solution given in (11). □

A.4. Proof of Theorem 3.

PROOF. Let $Q_Z^{(k)}(\theta; \mathbf{u}_s)$ be the empirical quantile computed on the projected training data $\mathbf{u}_s^\top \mathbf{X}^{(k)}$. We write

$$
\Phi^{(k)}(\theta; \mathbf{u}_s^\top \mathbf{Y}) = \gamma_s^{(k)}(\theta) |\mathbf{u}_s^\top \mathbf{Y} - Q_Z^{(k)}(\theta; \mathbf{u}_s)|,
$$

where $\gamma_s^{(k)}(\theta) = \theta + (1 - 2\theta) I\{\mathbf{u}_s^\top \mathbf{Y} < Q_Z^{(k)}(\theta; \mathbf{u}_s)\}$. Let $\boldsymbol{\mu}_y$ denote the vector of quantiles of $\mathbf{Y}$, and put $\boldsymbol{\mu}_y^{(k)} = \boldsymbol{\mu}^{(k)} - \boldsymbol{\mu}_y$ for $k = 1, 2$ and write $\mathbf{V} = \mathbf{Y} - \boldsymbol{\mu_y}$. By the triangular inequality

$$
\begin{aligned}
&\gamma_s^{(2)}(\theta) |\mathbf{u}_s^\top \mathbf{Y} - Q_Z^{(2)}(\theta; \mathbf{u}_s)| - \gamma_s^{(1)}(\theta) |\mathbf{u}_s^\top \mathbf{Y} - Q_Z^{(1)}(\theta; \mathbf{u}_s)| \\
&= \gamma_s^{(2)}(\theta) |\mathbf{u}_s^\top \mathbf{V} - \mathbf{u}_s^\top \boldsymbol{\mu}_y^{(2)}| - \gamma_s^{(1)}(\theta) |\mathbf{u}_s^\top \mathbf{V} - \mathbf{u}_s^\top \boldsymbol{\mu}_y^{(1)}| \\
&+ \tau_2 |Q_Z^{(2)}(\theta; \mathbf{u}_s) - \mathbf{u}_s^\top \boldsymbol{\mu}^{(2)}| + \tau_1 |Q_Z^{(1)}(\theta; \mathbf{u}_s) - \mathbf{u}_s^\top \boldsymbol{\mu}^{(1)}|,
\end{aligned}
$$

where τ_1 and τ_2 satisfy $|\tau_k| \leq 1$, $k = 1, 2$. Hence

$$
\begin{aligned}
T_1 &\equiv \sum_{s=1}^{S} \gamma_s^{(2)}(\theta) |\mathbf{u}_s^\top \mathbf{Y} - Q_Z^{(2)}(\theta; \mathbf{u}_s)| - \gamma_s^{(1)}(\theta) |\mathbf{u}_s^\top \mathbf{Y} - Q_Z^{(1)}(\theta; \mathbf{u}_s)| \\
&= T_2 + \tau_1 R_1 + \tau_2 R_2,
\end{aligned}
$$

where $T_2 = \sum_{s=1}^{S} \gamma_s^{(2)}(\theta)|\mathbf{u}_s^{\top}\mathbf{V} - \mathbf{u}_s^{\top}\boldsymbol{\mu}_y^{(2)}| - \gamma_s^{(1)}(\theta)|\mathbf{u}_s^{\top}\mathbf{V} - \mathbf{u}_s^{\top}\boldsymbol{\mu}_y^{(1)}|$, $R_1 = \sum_{s=1}^{S} |Q_Z^{(1)}(\theta; \mathbf{u}_s) - \mathbf{u}_s^{\top}\boldsymbol{\mu}^{(1)}|$ and $R_2 = \sum_{s=1}^{S} |Q_Z^{(2)}(\theta; \mathbf{u}_s) - \mathbf{u}_s^{\top}\boldsymbol{\mu}^{(2)}|$.

Given the convergence of the empirical quantiles to the respective population quantiles, it follows that

$$P^{(1)}(T_1 > c_1 - 2c_2 Sn^{-1/2}) \geq P^{(1)}(T_2 > c_1) - P(R_1 > c_2 Sn^{-1/2}) - P(R_2 > c_2 Sn^{-1/2})$$

$$\geq P^{(1)}(T_2 > c_1) - 2\sum_{s=1}^{S} e^{-2n_1\delta_s^{(1)}} - 2\sum_{s=1}^{S} e^{-2n_2\delta_s^{(2)}}$$

for any $c_1, c_2 > 0$, where

$$\delta_s^{(k)} = \left[\min\left\{F^{(k)}\left(\mathbf{u}_s^{\top}\boldsymbol{\mu}^{(k)} + \frac{c_2 S}{n^{1/2}}\right) - \theta, \theta - F^{(k)}\left(\mathbf{u}_s^{\top}\boldsymbol{\mu}^{(k)} - \frac{c_2 S}{n^{1/2}}\right)\right\}\right]^2.$$

Now define

$$d_s = \mathrm{E}\left\{\gamma_s^{(2)}(\theta)|\mathbf{u}_s^{\top}(\mathbf{V} - \boldsymbol{\mu}_y^{(2)})| - \gamma_s^{(1)}(\theta)|\mathbf{u}_s^{\top}(\mathbf{V} - \boldsymbol{\mu}_y^{(1)})|\right\}.$$

Given $\epsilon > 0$, let $\mathcal{K}_\epsilon$ denote the set of indices $s \in \{1, 2, \ldots, S\}$ such that

$$|\gamma_s^{(2)}(\theta)\mathbf{u}_s^{\top}\boldsymbol{\mu}_2 - \gamma_s^{(1)}(\theta)\mathbf{u}_s^{\top}\boldsymbol{\mu}_1| > \epsilon$$

$\forall \theta \in (0, 1)$. Under the assumption that $\mathbf{Y}$ has distribution $F^{(1)}$, we have

$$d_s = \mathrm{E}\left\{\gamma_s^{(2)}(\theta)|\mathbf{u}_s^{\top}(\mathbf{Y} - \boldsymbol{\mu}_2)| - \gamma_s^{(1)}(\theta)|\mathbf{u}_s^{\top}(\mathbf{Y} - \boldsymbol{\mu}_1)|\right\}$$

$$= \gamma_s^{(2)}(\theta)\,\mathrm{E}_1\,|\mathbf{u}_s^{\top}(\mathbf{Z} + \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)| - \gamma_s^{(1)}(\theta)\,\mathrm{E}_1\,|\mathbf{u}_s^{\top}\mathbf{Z}|,$$

where E_1 is the expectation under $P^{(1)}$.

Therefore, by assumption (iii) and provided $c \geq \epsilon$, we have

$$\sum_{s \in \mathcal{K}_\epsilon} d_s \geq a(c)(\sharp\mathcal{K}_c)$$

where $a(c) > 0$, with $a(c) = \gamma_s^{(2)}(\theta)\,\mathrm{E}_1\,|\mathbf{u}_s^{\top}(\mathbf{Z} + \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)| - \gamma_s^{(1)}(\theta)\,\mathrm{E}_1\,|\mathbf{u}_s^{\top}\mathbf{Z}|$ in view of (iii). As a consequence, for $\mathrm{E}_1(T_2) = \sum_{s=1}^{S} d_s$ and $\epsilon \to 0$, and $\forall c$, we have

(A.1) $$\mathrm{E}_1(T_2) \geq a(c)(\sharp\mathcal{K}_c),$$

where $\sharp A$ denotes the cardinality of the set A.

By the Chebychev inequality and provided that $c_1 < \frac{1}{2} \mathrm{E}_1(T_2)$, we have

$$P^{(1)}(T_2 > c_1) \geq 1 - P^{(1)}(|T_2 - \mathrm{E}_1(T_2)| > c_1) \geq 1 - c_1^{-2} \mathrm{E}_1\{T_2 - \mathrm{E}_1(T_2)\}^2$$
$$(A.2) \qquad \geq 1 - c_1^{-2} \mathrm{var}_1(T2) \geq 1 - A_2 c_1^{-2} S,$$

where var_1 denotes the variance under $P^{(1)}$ and the second inequality follows from assumption (ii); more specifically

$$\mathrm{var}_1(T_2) = \mathrm{var}_1 \left\{ \sum_{s=1}^{S} \left(\gamma_s^{(2)}(\theta)|\mathbf{u}_s^\top(\mathbf{V} - \boldsymbol{\mu}_y^{(2)})| - \gamma_s^{(1)}(\theta)|\mathbf{u}_s^\top(\mathbf{V} - \boldsymbol{\mu}_y^{(1)})| \right) \right\}$$
$$\leq \mathrm{var}_1 \left\{ \sum_{s=1}^{S} \left(\gamma_s^{(2)}(\theta)\mathbf{u}_s^\top(\mathbf{V} - \boldsymbol{\mu}_y^{(2)}) - \gamma_s^{(1)}(\theta)\mathbf{u}_s^\top(\mathbf{V} - \boldsymbol{\mu}_y^{(1)}) \right) \right\}$$
$$= \mathrm{var}_1 \left\{ \sum_{s=1}^{S} \left(\gamma_s^{(2)}(\theta)\mathbf{u}_s^\top(\mathbf{W} + \boldsymbol{\mu}^{(1)} - \boldsymbol{\mu}^{(2)}) - \gamma_s^{(1)}(\theta)\mathbf{u}_s^\top\mathbf{W} \right) \right\}$$
$$\leq \sum_{s=1}^{S} A_2 \mathbf{u}_s^\top\mathbf{u}_s + 2 \sum_{s=1}^{S-1} \sum_{s'=s+1}^{S} A_2 \mathbf{u}_s^\top\mathbf{u}_{s'}.$$

Stam (1982) proved that a uniform random variable on the sphere, $\mathbf{U} \in R^p$, converges to a standard Gaussian as $p \to \infty$. Therefore, for $S \to \infty$, by the strong law of large numbers we have

$$\frac{2\sum_{s=1}^{S-1}\sum_{s'=s+1}^{S} A_2 \mathbf{U}_s^\top\mathbf{U}_{s'}}{S(S-1)} \xrightarrow{a.s.} A_2 \mathrm{E}(\boldsymbol{\Xi}_1^\top\boldsymbol{\Xi}_2) = 0,$$

where $\boldsymbol{\Xi}_1$ and $\boldsymbol{\Xi}_2$ are two independent standard Gaussians. This explains why the covariances become negligible in the last part of (A.2) as p increases.

It remains to prove that $c_1 < \frac{1}{2} \mathrm{E}_1(T_2)$. Consider $c_1 = \frac{c_3 S}{n^{1/2}}$, where c_3 is a positive constant. By (A.1), the latter holds if $c_3 S n^{-1/2} < \frac{1}{2}a(c)\mathcal{K}_c$. But this is true because it implies that

$$S\left(n^{1/2}\sharp\mathcal{K}_c\right)^{-1} < \frac{1}{2}a(c)c_3^{-1},$$

where the term on the left goes to zero according to assumption (iv) while $a(c) > 0$, thus $c_3^{-1} > 0$.

For $c_1 = \frac{c_3 S}{n^{1/2}}$, we have

$$P^{(1)}(T_1 > c_3 S n^{-1/2} - 2c_2 S n^{-1/2}) \geq 1 - A_2 \frac{n}{c_3^2 S} - 2 \sum_{s=1}^{S} e^{-2n_1\delta_s^{(1)}} - 2 \sum_{s=1}^{S} e^{-2n_2\delta_s^{(2)}}.$$

We wish to choose c_3 and c_2 such as

$$P^{(1)}(T_1 > 0) \geq 1 - \epsilon.$$

Therefore, we fix ϵ and choose c_3 such that $\frac{A_2}{c_3^2 A_1} \leq \epsilon$, where A_1 is defined in assumption (i). It follows that

$$\frac{A_2 S}{c_1^2} = A_2 \frac{n}{c_3^2 S} \leq \frac{A_2}{c_3^2 A_1} \leq \epsilon.$$

Then we choose c_2 such that $c_3 > 2c_2$ and observe that $2 \sum_{s=1}^{S} e^{-2n_1 \delta_s^{(1)}} + 2 \sum_{s=1}^{S} e^{-2n_2 \delta_s^{(2)}} \to 0$ for $n, S \to \infty$. Since this is true for each $\epsilon > 0$, then $P^{(1)}(T_1 > 0) \to 1$, and similarly $P^{(2)}(T_1 < 0) \to 1$. $\qquad \square$

REFERENCES

BRADLEY, R. C. (2005). Basic properties of strong mixing conditions: A survey and some open questions. *Probability Surveys* **2** 107-144.

CORTES, C. and VAPNIK, V. (1995). Support-vector networks. *Machine Learning* **20** 273-297.

COVER, T. M. and HART, P. E. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory* **13** 21-27.

GERACI, M. (2016). Qtools: A collection of models and other tools for quantile inference. *R Journal* **8** 117-138.

GERACI, M. (2020). Qtools: Utilities for quantiles R package version 1.5.2.

GERACI, M., BOGHOSSIAN, N. S., FARCOMENI, A. and HORBAR, J. D. (2020). Quantile contours and allometric modelling for risk classification of abnormal ratios with an application to asymmetric growth-restriction in preterm infants. *Statistical Methods in Medical Research* **29** 1769-1786.

HALL, P., TITTERINGTON, D. M. and XUE, J. H. (2009). Median-based classifiers for high-dimensional data. *Journal of the American Statistical Society* **104** 1597-1608.

HAND, D. J. and YU, K. (2001). Idiot's Bayes - Not so stupid after all? *International Statistical Review* **69** 385-398.

HENNIG, C. and VIROLI, C. (2016a). Quantile-based classifiers. *Biometrika* **103** 435-446.

HENNIG, C. and VIROLI, C. (2016b). quantileDA: Quantile classifier R package version 1.1.

JOE, H. (2006). Generating random correlation matrices based on partial correlations. *Journal of Multivariate Analysis* **97** 2177-2189.

KONG, L. and MIZERA, I. (2012). Quantile Tomography: Using Quantiles with Multivariate Data. *Statistica Sinica* **22** 1589-1610.

LAI, Y. and MCLEOD, I. (2019). eqc: Ensemble quantile classification R package version 1.2-2.

LAI, Y. and MCLEOD, I. (2020). Ensemble quantile classifier. *Computational Statistics & Data Analysis* **144** 106849.

MEYER, D., DIMITRIADOU, E., HORNIK, K., WEINGESSEL, A. and LEISCH, F. (2019). e1071: Misc functions of the department of statistics, probability theory group (formerly: E1071), TU Wien R package version 1.7-3.

PARK, M. Y. and HASTIE, T. (2008). Penalized logistic regression for detecting gene interactions. *Biostatistics* **9** 30-50.

PARK, M. Y. and HASTIE, T. (2018). stepPlr: L2 penalized logistic regression with stepwise variable selection R package version 0.93.

QIU, W. and JOE, H. (2015). clusterGeneration: Random cluster generation (with specified degree of separation) R package version 1.3.4.

SORRENTINO, D., AVELLINI, C., GERACI, M., DASSOPOULOS, T., ZARIFI, D., VADALÁ DI PRAMPERO, S. F. and BENEVENTO, G. (2014). Tissue studies in screened first-degree relatives reveal a distinct Crohn's disease phenotype. *Inflammatory Bowel Diseases* **20** 1049-1056.

STAM, A. J. (1982). Limit theorems for uniform distributions on spheres in high-dimensional Euclidean spaces. *Journal of Applied Probability* **19** 221-228.

R CORE TEAM (2020). R: A language and environment for statistical computing R Foundation for Statistical Computing, Vienna, Austria.

TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B. and CHU, G. (2002). Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences of the United States of America* **99** 6567-6572.

VENABLES, W. N. and RIPLEY, B. D. (2002). *Modern applied statistics with S*, Fourth ed. Springer, New York, NY.

WANG, L., ZHU, J. and ZOU, H. (2008). Hybrid Huberized support vector machines for microarray classification and gene selection. *Bioinformatics* **24** 412-419.