



Class Maps for Visualizing Classification Results

Jakob Raymaekers, Peter J. Rousseeuw & Mia Hubert

To cite this article: Jakob Raymaekers, Peter J. Rousseeuw & Mia Hubert (2022) Class Maps for Visualizing Classification Results, *Technometrics*, 64:2, 151-165, DOI: [10.1080/00401706.2021.1927849](https://doi.org/10.1080/00401706.2021.1927849)

To link to this article: <https://doi.org/10.1080/00401706.2021.1927849>



© 2021 The Author(s). Published with license by Taylor & Francis Group, LLC.



[View supplementary material](#)



Published online: 25 Jun 2021.



[Submit your article to this journal](#)



Article views: 2920



[View related articles](#)



[View Crossmark data](#)



Citing articles: 1 [View citing articles](#)

Class Maps for Visualizing Classification Results

Jakob Raymaekers, Peter J. Rousseeuw, and Mia Hubert

Section of Statistics and Data Science, Department of Mathematics, KU Leuven, Belgium

ABSTRACT

Classification is a major tool of statistics and machine learning. A classification method first processes a training set of objects with given classes (labels), with the goal of afterward assigning new objects to one of these classes. When running the resulting prediction method on the training data or on test data, it can happen that an object is predicted to lie in a class that differs from its given label. This is sometimes called label bias, and raises the question whether the object was mislabeled. The proposed class map reflects the probability that an object belongs to an alternative class, how far it is from the other objects in its given class, and whether some objects lie far from all classes. The goal is to visualize aspects of the classification results to obtain insight in the data. The display is constructed for discriminant analysis, the k-nearest neighbor classifier, support vector machines, logistic regression, and coupling pairwise classifications. It is illustrated on several benchmark datasets, including some about images and texts.

ARTICLE HISTORY

Received September 2020
Accepted April 2021

KEYWORDS

Discriminant analysis;
k-Nearest neighbors;
Mislabeling; Pairwise
coupling; Support vector
machines

1. Introduction

Classification is a major tool of statistics and machine learning. For an extensive introduction to classification methods, see Hastie, Tibshirani, and Friedman (2017). A classification method first processes a training set of objects with given classes (labels), with the goal of afterward assigning new objects to one of these classes. When running the resulting prediction method on the training data or on validation data or test data, it can happen that an object is predicted to lie in a class that differs from its given label. This is sometimes called label bias, and raises the question whether the object might have been mislabeled. Our goal is to visualize aspects of the data classification to obtain insight in the data. For consistency, we will depict the predictions in the vertical direction throughout.

We start with a simple illustrative example. The floral bud data originate from Wouters et al. (2015) and were kindly provided to us by Dr. Bart De Ketelaere. They arose from an experiment in a pear orchard with the goal of detecting floral buds in their environment with branches, bud scales, and supports for the trees. The data contain 550 observations and 6 variables. There are four classes: “branch” (49 members), “bud” (363), “scale” (94), and “support” (44). The classification in Section 3 yields the confusion matrix

$$\begin{pmatrix} 46 & 0 & 0 & 3 \\ 0 & 355 & 2 & 6 \\ 2 & 0 & 90 & 2 \\ 6 & 1 & 0 & 37 \end{pmatrix}$$

in which the rows are the given classes and the columns are the predicted classes.

After a classification is carried out, we can display the result in a stacked bar chart or a mosaic plot (Hartigan and Kleiner

1981; Friendly 1994). Figure 1 shows such a stacked mosaic plot, which graphically represents the confusion matrix. The classes are represented by colors reminiscent of their meaning. The given classes are on the horizontal axis, and the predicted labels are on the vertical axis. The area of each rectangle is proportional to the number of objects in it. The display immediately shows that the given classes have different numbers of objects. Several variations of this plot are possible. One could rank the vertical labels in the order of the original classes, but we choose to put the given class at the bottom so that the lower part of each bar reflects the objects that were classified in accordance with their label. For the other labels in each bar we take the order of the remaining original classes. Here, we see that buds are often classified correctly but that there is some confusion between branch and support.

Figure 1 does not yet give us an idea why some object is predicted to belong to a different class. Is it because the object lies in or near a region where classes overlap? Or is it deeply inside its predicted class and far from its given class, arousing suspicion that its original label was wrong? Or is it actually far from both its given and predicted classes? To assist with these questions we will propose a display that incorporates additional information.

Section 2 gives general definitions and illustrates the display on an example. The subsequent sections apply it to discriminant analysis, the k-nearest neighbor method, support vector machines, logistic regression, and pairwise coupling.

2. Basic Principles of a Class Map

Suppose we have objects denoted by their index i , where $i = 1, \dots, n$, and there are classes (labels, groups) g with $g =$

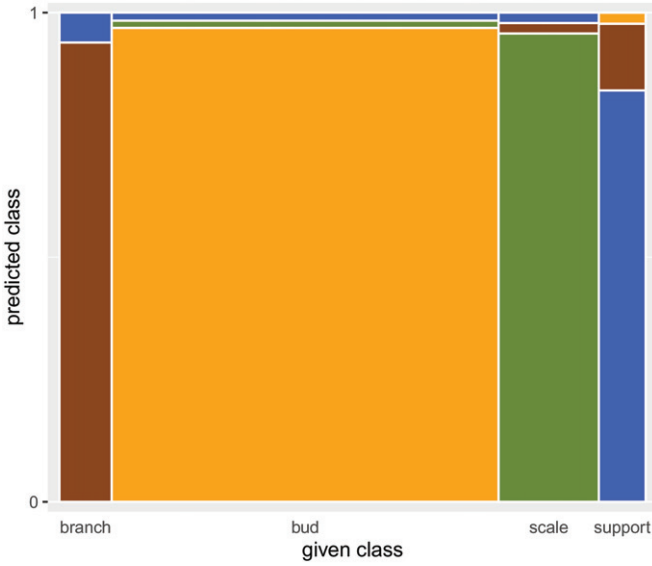


Figure 1. A stacked mosaic plot of a classification of the floral bud data. The given classes (labels) are on the horizontal axis, and the predicted classes are on the vertical axis. The area of each rectangle is proportional to the number of objects in it.

$1, \dots, G$. The target is thus a discrete variable with G levels. Consider a case i in the training set or a test set. Either explicitly or implicitly, most classifiers provide posterior probabilities $\hat{p}(i, g)$ of object i belonging to each of the classes g , with $\sum_{g=1}^G \hat{p}(i, g) = 1$ for each i . For instance, in discriminant analysis the posterior probabilities are based on the estimated densities of the classes, whereas the k -nearest neighbor classifier estimates them by the class frequencies in the k -neighborhood of i .

The object i is then classified according to the rule

$$\text{assign object } i \text{ to class } \operatorname{argmax}_{g=1, \dots, G} \hat{p}(i, g). \quad (1)$$

Now assume that object i has a known given label g_i . We wish to measure to what extent the given label g_i agrees with the classification. For this purpose we define the highest $\hat{p}(i, g)$ attained by a class *different from* g_i as

$$\tilde{p}(i) = \max\{\hat{p}(i, g) ; g \neq g_i\}. \quad (2)$$

The class attaining this maximum can be seen as the best alternative class. If $\hat{p}(i, g_i) > \tilde{p}(i)$ it follows that g_i attains the overall highest value of $\hat{p}(i, g)$ so the classifier agrees with the given class g_i . On the other hand, if $\hat{p}(i, g_i) < \tilde{p}(i)$ the classifier will not assign object i to class g_i .

We now compute the conditional posterior probability of the best alternative class when comparing it with the given class g_i as

$$\text{PAC}(i) = \frac{\tilde{p}(i)}{\hat{p}(i, g_i) + \tilde{p}(i)}. \quad (3)$$

The abbreviation PAC stands for Probability of the Alternative Class. Note that when $\text{PAC}(i) < 0.5$ the classifier does predict the given class g_i , whereas $\text{PAC}(i) > 0.5$ indicates that the best alternative class outperforms g_i in the eyes of the classifier. $\text{PAC}(i) \approx 0$ corresponds to a class that fits very well, and $\text{PAC}(i) \approx 1$ to a class that does not fit well at all. In cluster analysis, the silhouette plot of Rousseeuw (1987) is based on the

same idea of quantifying how well an object is positioned in its class by a comparison with the best alternative class.

Note that the denominator of (3) is strictly positive by construction. When there are $G = 2$ classes, the alternative class is the only other class, the denominator of (3) simplifies to 1, and $\text{PAC}(i)$ becomes an unconditional probability. The PAC bears some similarity to the notion of margin given by $m(i) = \hat{p}(i, g_i) - \tilde{p}(i)$, whose distribution is used by Cai, Zhang, and Peng (2009) to measure the classification capability of an ensemble of classifiers.

The second component of the proposed display reflects how far the object i is from its given class g_i , again in the eyes of the classifier. For this we start by computing a distance $D(i, g_i)$ of i to g_i . The choice of D needs to be relevant for the type of classifier, because different classifiers make different assumptions about the data. For instance, discriminant analysis assumes roughly elliptical point clouds, k -nearest neighbors assume no particular shape but focus on local inter-point dissimilarities, and support vector machines allow for disconnected classes. Next we estimate the cumulative distribution function (cdf) of $D(x, g_i)$ where x is a random object generated from class g_i . Then we define the *farness* of the object i to its class g_i as that cdf in $D(i, g_i)$, that is,

$$\text{farness}(i) = P[D(x, g_i) \leq D(i, g_i)]. \quad (4)$$

Therefore, $\text{farness}(i)$ lies in the $[0, 1]$ range, just like $\text{PAC}(i)$. Our implementation estimates the cdf of D by pooling the $D(i, g_i)$ for $i = 1, \dots, n$ and applying the function `transfo` of the R-package `cellwise` (Raymaekers and Rousseeuw 2020). This method is described in Raymaekers and Rousseeuw (2021a) and in Section A.1 of the Supplementary Material.

We propose to draw the so-called *class map* for each class g , which plots $\text{PAC}(i)$ versus $\text{farness}(i)$ for all objects i in g . Figure 2 is such a class map where g is the class “bud” in Figure 1. Both axes are expressed in probabilities. Most points lie in the gray region where $\text{PAC} < 0.5$, meaning that the classifier has assigned them to their given class, “bud,” and they are shown in orange, the color of the class “bud” in Figure 1. The lower $\text{PAC}(i)$, the more the classifier is convinced they belong to their class. For instance, point a has $\text{PAC}(i) \approx 0$. There are also some points above the gray region, where $\text{PAC} > 0.5$ so the classifier has assigned them to their best alternative class instead of “bud,” with much conviction when PAC is high. Point b is like that, and its green color indicates that it was assigned to the class “scale.” Points with high $\text{PAC}(i)$ do not sit well inside their given class, the graphical analogy being that of a fish out of water (the gray zone). There can also be boundary cases, such as point c with $\text{PAC} \approx 0.5$, meaning that the classifier does not have much conviction to assign it to either its given class or its alternative class.

Let us now consider $\text{farness}(i)$ as well. This is a probability, but the tick marks on the horizontal axis are shown on the scale of the quantiles of the standard Gaussian distribution restricted to $[0, 4]$, in order to better distinguish the objects with relatively high farness. The dashed vertical line is at 0.99 here, but the user can choose a different value if so desired. (It is not a canonical cutoff like that of $\text{PAC}(i)$, where 0.5 is the boundary between assigning i to its given class or not.) Points a, b, and c have

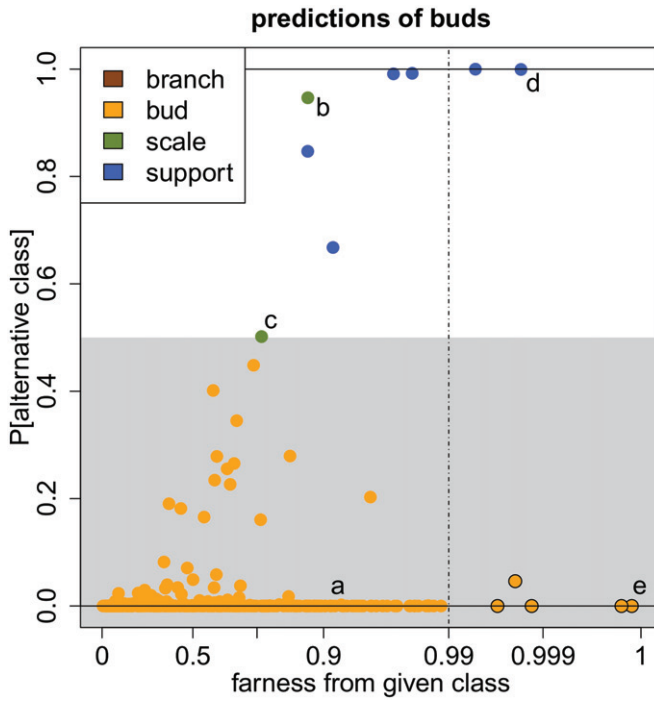


Figure 2. Example of a class map.

typical farness values, meaning they are relatively close to “bud.” We can say that point a is in the best quadrant, since it is assigned to its own class where it is not far from the center. Point b is not far from the center of “bud” but nevertheless assigned to “scale” with much conviction, so there is likely some overlap between these two classes, with b closer to “scale” than to “bud,” and c roughly equally far from both.

What about the remaining two quadrants? Point d is far from “bud,” and at the same time assigned to “support” (blue color) with much conviction. In this quadrant, the situation is quite clear: point d has two arguments in favor of belonging to its alternative class “support.” In the final quadrant, point e lies quite far from “bud” but it is still assigned to “bud” with much conviction. This can happen when a case is far from its class but in the opposite direction of the best alternative class.

A final piece of information incorporated in the class map is whether a point with unusually high $\text{farness}(i) = \text{farness}(i, g_i)$ to its own class g_i is also far from the other classes. To find out we compute the “overall farness” of each object i as

$$O(i) = \min_{g=1, \dots, G} \text{farness}(i, g). \quad (5)$$

When $O(i)$ exceeds the cutoff (here 0.99) it can be considered an outlier, and we put more emphasis on that point by using a different plotting symbol which has a black border around it. For such a point no class seems really suitable. This happens for point e, but not for point d.

When we are dealing with training data, the given class labels g_i are known, which allows us to compute (2)–(4). For cases without given label, these are not defined. But we can also have test data or validation data with given labels. Then we apply the trained model to the new data, yielding new posterior probabilities, the prediction based on (1), as well as $\text{PAC}(i)$ from (2) to (3). For $\text{farness}(i)$ one needs to compute the $D(i, g_i)$ with

the same parameters as in the training data, and ensure that (4) uses the distribution fitted on the training data. So for each object i in the test set, $\text{PAC}(i)$, $\text{farness}(i)$, and $O(i)$ only depend on the training set and that case i , and not on the remainder of the test set. The test set may even consist of a single case.

The main purpose of class maps is to learn more about the objects in a classification. For instance, class maps can give an indication of potential mislabeling. When cases have a high PAC, and especially when their farness is large too, one may want to check whether their original label was correct. When mislabeling is detected in the training data, it may be beneficial to retrain the classifier without the mislabeled cases.

Class maps may also reveal subgroups in the data, which can indicate that a class is not really homogeneous, and should perhaps be split up. (An example of a subgroup will be shown in Figure 7.) Class maps may also alert us to the possibility that the classifier is less suitable for the data at hand. In Section 8, we will illustrate this in an example where several classifiers are applied to the same dataset.

3. Discriminant Analysis

One of the oldest and the best understood classification techniques is discriminant analysis (DA), intended for objects that can be represented as points \mathbf{x}_i with d numerical coordinates. The underlying model is the normal mixture model as described in for example, McLachlan (2004, chap. 3). It assumes the points in class g follow a normal distribution $N(\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g)$ with unknown class mean $\boldsymbol{\mu}_g$ and covariance matrix $\boldsymbol{\Sigma}_g$, and unknown class probability π_g . To train the classifier we compute estimates $\hat{\boldsymbol{\mu}}_g$ and $\hat{\boldsymbol{\Sigma}}_g$ which are typically the empirical mean and covariance of the points \mathbf{x}_i in class g . The class probability π_g is estimated as $\hat{\pi}_g := n_g/n$, where n_g is the number of objects in class g . The estimated density of the mixture distribution is then $\sum_g \hat{\pi}_g \hat{f}_g(\mathbf{x})$ with \hat{f}_g the multivariate normal density

$$\hat{f}_g(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^d \det(\hat{\boldsymbol{\Sigma}}_g)}} e^{-\text{MD}^2(\mathbf{x}, \hat{\boldsymbol{\mu}}_g, \hat{\boldsymbol{\Sigma}}_g)/2}, \quad (6)$$

where $\text{MD}^2(\mathbf{x}, \hat{\boldsymbol{\mu}}_g, \hat{\boldsymbol{\Sigma}}_g) := (\mathbf{x} - \hat{\boldsymbol{\mu}}_g)' \hat{\boldsymbol{\Sigma}}_g^{-1} (\mathbf{x} - \hat{\boldsymbol{\mu}}_g)$ is the squared Mahalanobis distance.

The quadratic discriminant analysis (QDA) model allows the $\boldsymbol{\Sigma}_g$ to be different. On the other hand, the linear discriminant analysis (LDA) model assumes that all $\boldsymbol{\Sigma}_g$ are equal, so one needs a single covariance estimate. We can for instance compute $\hat{\boldsymbol{\Sigma}}$ from the pooled data $\tilde{\mathbf{x}}_i := \mathbf{x}_i - \hat{\boldsymbol{\mu}}_{g_i}$ for $i = 1, \dots, n$ where g_i is the given label of object i . We then set all $\hat{\boldsymbol{\Sigma}}_g := \hat{\boldsymbol{\Sigma}}$ in (6).

In either case, for a given object \mathbf{x}_i the posterior probability of any class g is thus

$$\hat{p}(i, g) = \frac{\hat{\pi}_g \hat{f}_g(i)}{\sum_{j=1}^G \hat{\pi}_j \hat{f}_j(i)}$$

so $\sum_{g=1}^G \hat{p}(i, g) = 1$, and we assign i by the maximum a posteriori classification rule (1). Applying Equations (2) and (3)

yields the conditional probability of the best alternative class \tilde{g}_i :

$$\text{PAC}(i) = \frac{\hat{\pi}_{\tilde{g}_i} \hat{f}_{\tilde{g}_i}(i)}{\hat{\pi}_{\tilde{g}_i} \hat{f}_{\tilde{g}_i}(i) + \hat{\pi}_{\hat{g}_i} \hat{f}_{\hat{g}_i}(i)} \quad (7)$$

in which the denominators of the posterior probabilities have canceled.

The second ingredient of the class map is the farness. In the setting of discriminant analysis, it is natural to set the distance $D(i, g)$ of object i to class g equal to

$$D(i, g) := \text{MD}(\mathbf{x}_i, \hat{\boldsymbol{\mu}}_g, \hat{\boldsymbol{\Sigma}}_g) \quad (8)$$

which is the well-known unsquared Mahalanobis distance in Equation (6) used by the classifier. We then compute the farness of object i to any class g by formula (4).

We classified the floral bud data described in the introduction by QDA, yielding the stacked plot in Figure 1 and the class map of “bud” in Figure 2. Figure 3 shows the class map of each given class separately. We already interpreted that of buds in Section 2. In the class map of scales, most points have a very low $\text{PAC} \approx 0$, meaning that they were assigned to their given class with much conviction. In the class map of branches, we see less conviction, and the three points with $\text{PAC}(i) > 0.5$ are in blue so they were assigned to class “support.” Since their farness is unexceptional, they can be assumed to lie in a region where “branch” and “support” overlap. Conversely, in the class map of supports we see six brown points, so we may conclude that “branch” and ‘support’ are not that well-separated.

When we have a labeled test set (validation set) the same formulas can be applied. When case i is in the test set its prediction \hat{g}_i is still given by (1), and we use the formulas (7) and (8) where the $\hat{\pi}_g$, $\hat{\boldsymbol{\mu}}_g$, $\hat{\boldsymbol{\Sigma}}_g$ and \hat{f}_g are those of the training set.

We now illustrate class maps on the MNIST benchmark dataset due to LeCun et al. (1998). It contains 70,000 images of handwritten digits, ranging from 0 to 9, of which 60,000 serve as a training set and 10,000 as a test set. Each image has been size-normalized and centered in a fixed-size image of 28×28 pixels on a grayscale.

The top row of Figure 4 shows some randomly sampled images from the dataset, one of each digit, and the bottom row shows the averaged image per class.

To predict digits from images, we first reduced the dimension of the data from 784 to 50 by PCA. This avoids numerical instability due to the inversion of covariance matrices in QDA. The misclassification rate of QDA is quite low here, around 4%. Figure 5 is the stacked mosaic plot of the QDA classification. Unlike Figure 1, here we opted to see the outliers [given by $O(i) > 0.99$ in (5)] as an extra predicted class in dark gray at the top. Fortunately there are not many outliers here.

The top panel of Figure 6 is the class map of the digit 0. The large majority of points have $\text{PAC} \approx 0$, meaning they lie well within the class. Some more exceptional points have been marked, and the corresponding images are shown under the class map. Let us start with points above the gray zone, that is, whose predictions differ from 0. Points a and b have a high PAC but a relatively low farness. We see that QDA predicts them as a 6, which is not surprising when looking at their images. We also see why the farness is not very high, since the images look a lot like a zero, but without fully closing the circle. Point c is a

very small circle, positioned somewhat higher than the average zero. This makes it look as though it could be the top part of the digit 9. QDA indeed considers it a 9, but its low farness indicates that it is not far from the class of zeroes. Point d has a $\text{PAC} > 0.5$, but not quite as high as the other marked points. QDA considers this a 5, and from the image we understand why. It is a 0 with an extra horizontal stroke on top, which gives it several characteristics of a 5. Note that its farness is not extreme, so this image is in between both classes. Points e to g are all classified as 2, and their images are zeroes with an extra line. In f and g the extra line becomes more slanted and pronounced, increasing the resemblance with the digit 2 which also has such a slanted stroke, and increasing their PAC. Note that these two points also have a higher farness and that g has a black border so it is an outlier, indicating that it is rather distant from all classes. Finally, points h, i and j are outliers with a PAC below 0.5, so they are assigned to class 0, but their images are flat zeroes written with thick pen strokes, unlike most zeroes in the data.

The class map of digit 1 is in Figure 7. Unlike the other class maps, all the digits with $\text{PAC} > 0.5$ have a high farness. This may be due to the fact that most images of 1 are just vertical lines, so they look very similar to each other. As a result, deviations from that shape get a high Mahalanobis distance. Out of the 6742 images, QDA predicts 306 in a different class, so their $\text{PAC} > 0.5$. Of these, 104 are predicted in class 2 (yellow points). The fact that these points are relatively concentrated in the class map suggests that they may have some properties in common. Figure 8 shows all the images of 1’s predicted as 2. Most of these digits are indeed 1’s but written with a horizontal line at the bottom (which often occurs in a 2). This is unlike the vast majority of 1’s in the data, as can be inferred from the averaged image of 1 in Figure 4. We have thus identified a subgroup of class 1, which corresponds with a particular writing style that is not recognized by the classifier.

The class maps of the remaining digits, as well as the relevant images, can be found in Section A.2 of the supplementary material.

When classifying the MNIST test data using the trained QDA model, the misclassification rate is about the same as on the training data. Therefore the classification is stable, without indication that the training data was overfitted. The class maps of the test data, shown in Section A.3, also look similar to those of the training data.

4. Classification by k-Nearest Neighbors

Another popular classifier is the *k-nearest neighbor method* (kNN) of Fix and Hodges (1951), which has several appealing properties. It is not restricted to data points with coordinates, as it can take data in the form of dissimilarities $d(i, j)$ between objects. Such a dissimilarity matrix may for instance originate from subjective judgments, in which case there were no coordinates to begin with, and the axioms of a metric need not be satisfied. Of course, if there are coordinates one can always compute dissimilarities from them. This even works when the variables are of mixed types. Kaufman and Rousseeuw (1990, chap. 1) describes how one can compute a dissimilarity matrix from mixed variables of continuous, symmetric binary,

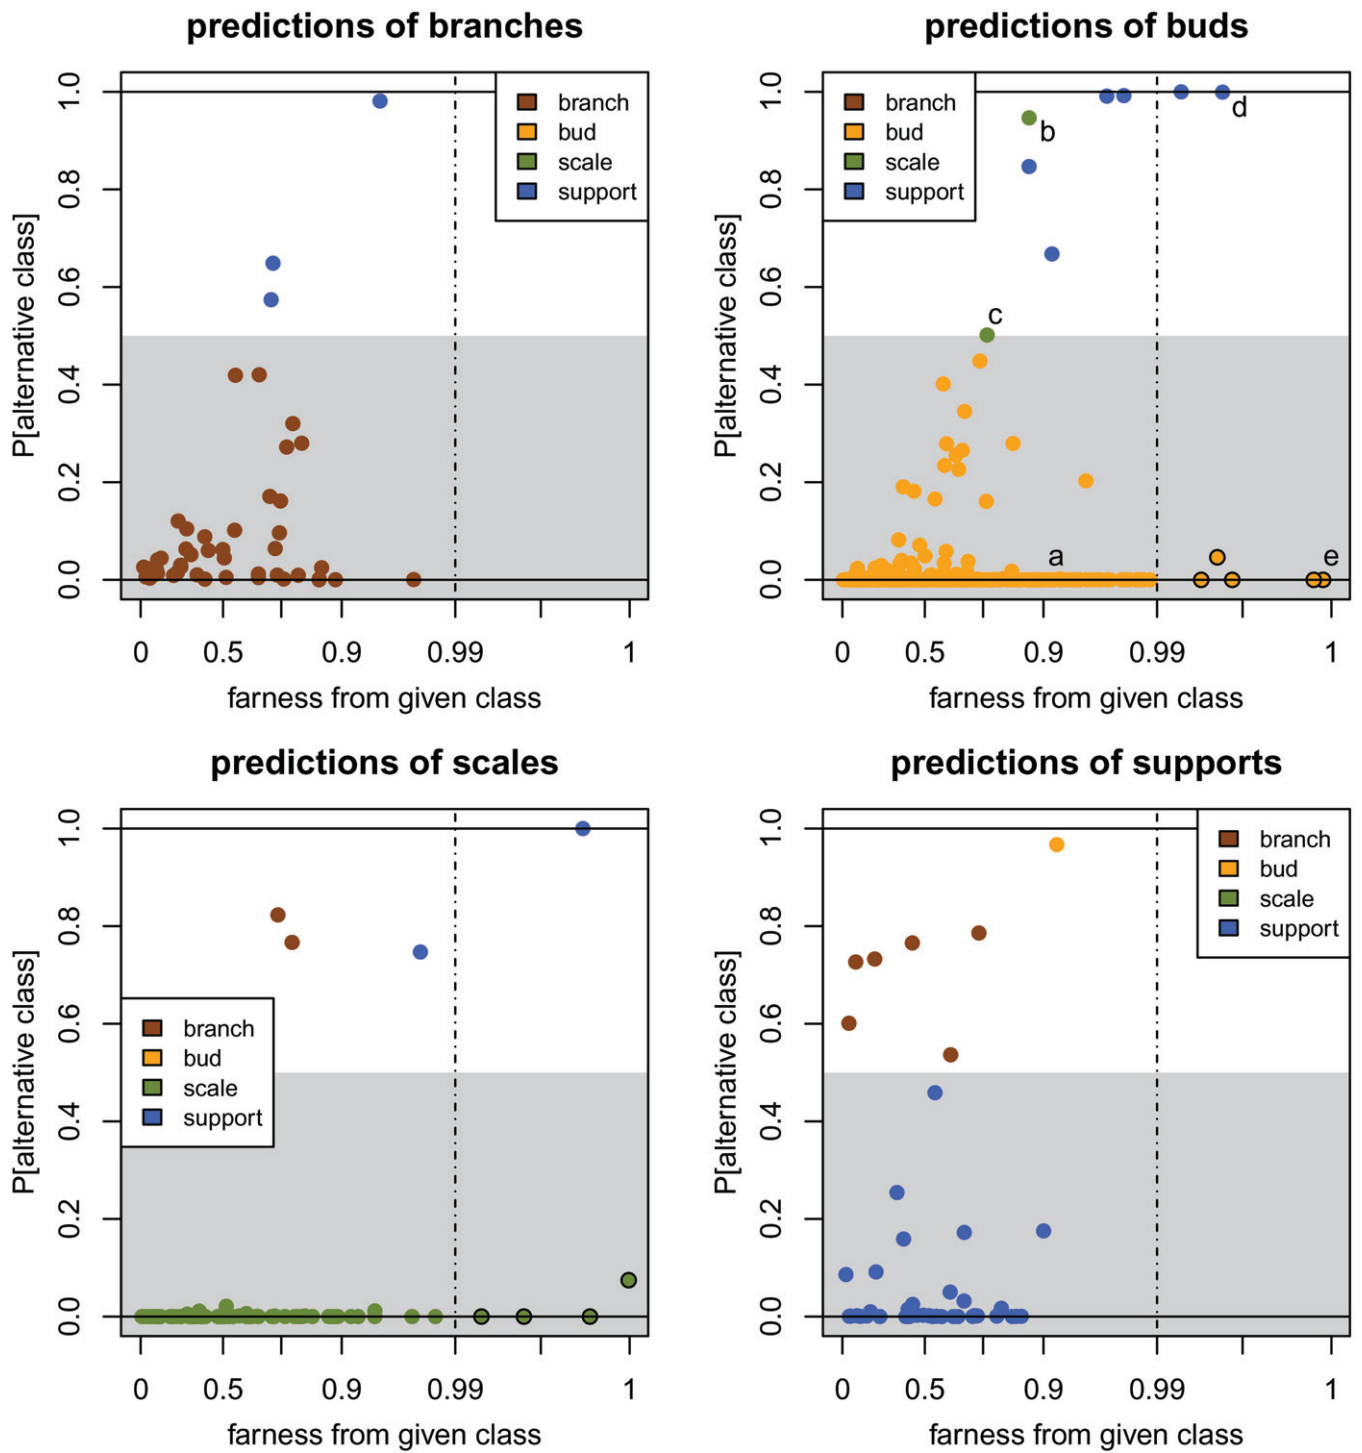


Figure 3. Floral buds data: maps of all four classes.



Figure 4. Top row: randomly sampled digits; bottom row: averaged images per class.




Figure 5. MNIST data: stacked mosaic plot where the objects flagged as outliers are shown in dark gray, as an extra class at the top.

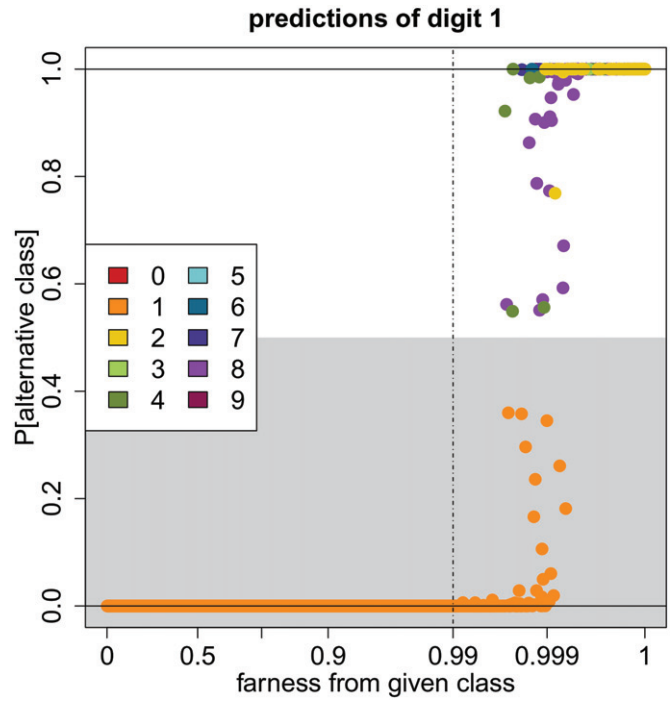


Figure 7. Class map of digit 1.

Around each object i , the kNN method determines its k -neighborhood consisting of the objects $j \neq i$ with the k smallest $d(i, j)$. Let us denote the k th such dissimilarity as d_i^* . Such a neighborhood is not always unique, as it can happen that there are other objects j' with the same dissimilarity $d(i, j') = d_i^*$. To make the neighborhood unique, a common option is to include such points j' as well, so we get neighborhoods $N(i)$ with $k(i)$ members where always $k(i) \geq k$.

The kNN method then estimates the probabilities $\hat{p}(i, g)$ of object i belonging to each of the classes g by the relative frequencies

$$\hat{p}(i, g) := n_i(g)/k(i) \quad (9)$$

where $n_i(g)$ counts how many objects in the neighborhood $N(i)$ have label g . These estimated probabilities again satisfy $\sum_{g=1}^G \hat{p}(i, g) = 1$. The kNN classifier then predicts the label of i as the class \hat{g}_i with highest $\hat{p}(i, g)$, in line with (1). Also here ties can occur. Some implementations choose randomly between tied labels. Our implementation breaks ties by assigning i to the tied label g for which the average dissimilarity between i and the members of g in $N(i)$ is lowest.

From its definition, we see that kNN makes no explicit assumptions about underlying distributions, and that it can focus on local structure (nearby objects) rather than global structure. Both aspects are quite different from DA. For a given dataset, a typical way to select an appropriate value of k is to cross validate the misclassification rate. Here we will assume that k has already been selected.

The conditional probability of the best alternative class $\text{PAC}(i)$ is obtained by applying (2) and (3) to (9). When $\text{PAC}(i) = 0$ it means that all members of $N(i)$ have the same label \hat{g}_i . At the other extreme, $\text{PAC}(i) = 1$ says that all members of $N(i)$ have the same label \hat{g}_i which differs from

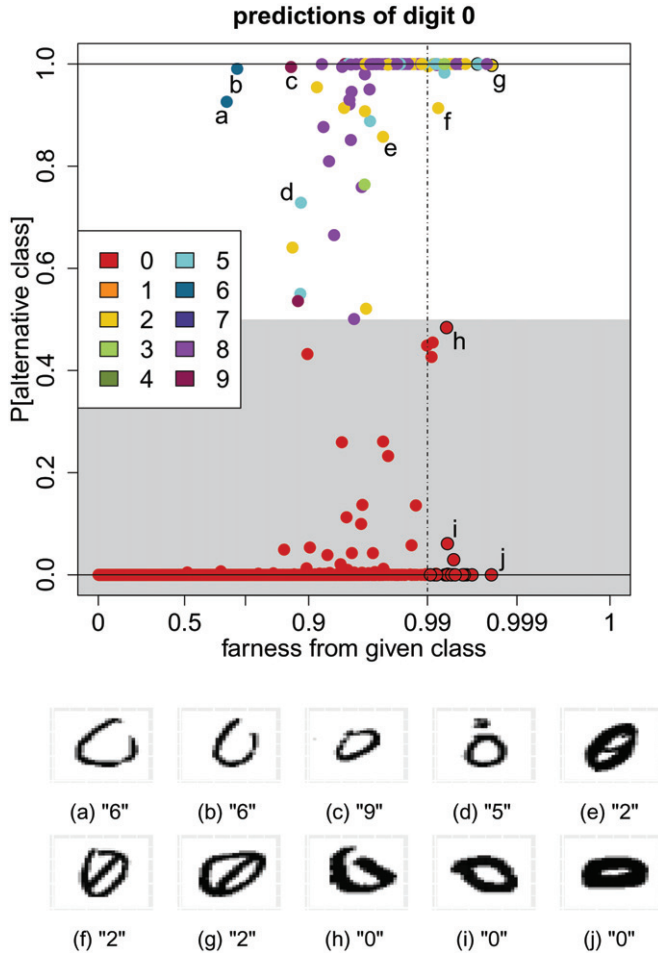


Figure 6. Class map of the digit 0, with the images corresponding to the marked points.

asymmetric binary, nominal and ordinal types, and this is implemented in the function `daisy()` of the R package `cluster` (Maechler et al. 2019).



Figure 8. Images of the digits in class 1 predicted as a 2.

g_i . The boundary is again at $\text{PAC} = 0.5$, with $\text{PAC}(i) > 0.5$ signifying that the predicted label \hat{g}_i fits better than the given label g_i . Unlike DA, here $\text{PAC}(i)$ takes discrete values, which are multiples of $1/k(i)$.

For computing fairness, we can no longer use (8) since the kNN classifier does not require coordinates, only interpoint dissimilarities. The fairness should thus be based on those dissimilarities. For each object i and class g , we compute $D(i, g)$ as the median of the k smallest dissimilarities $d(i, j)$ to all

objects j of class g . For each class g , we then divide $D(i, g)$ by $\text{median}\{D(j, g) : j \text{ belongs to class } g\}$. This makes the fairness values from all classes more comparable to each other. (This $D(i, g)$ definition did not originate from the literature.) Finally, we estimate the distribution of the $D(i, g)$ as we did for DA, and compute $\text{farness}(i)$ from (4). The overall fairness is then given by (5) as before.

The above formulas can also be used when we have a test set (validation set) with labels. For a test case i , we




Figure 9. Class map of the nonspam mails.

consider a neighborhood $N(i)$ consisting only of cases j in the training set. Maximizing (9) yields the predicted class \hat{g}_i as before. In the farness computation, we reuse the quantities $\text{median}\{D(j, g); j \text{ belongs to class } g\}$ from the training set, as well as the distribution fitted to the $D(i, g)$ of the training set.

We illustrate the resulting display on the benchmark *spam* data. These consist of 4601 emails collected by George Forman at the Hewlett-Packard Labs, and are labeled as spam (1813 mails) or non-spam (2788 mails). The data are publicly available in the R-package *kernlab* (Karatzoglou, Smola, and Hornik 2019). The mails were converted to a numerical characterization using 57 variables. Section A.4 of the supplementary material lists the variables and their interpretation. Unfortunately, the mails themselves are not available.

The data are classified by kNN with $k = 5$, yielding an in-sample misclassification rate under 9%. Figure 9 shows the class map of the nonspam (also called ham) mails. As expected, most of the PAC values are at one of the six main levels between $\text{PAC} = 0$ and $\text{PAC} = 1$ by steps of $1/k = 0.2$. The PAC values in between these levels come from the neighborhoods $N(i)$ that contain $k(i) > k$ members due to tied dissimilarities.

Some atypical points are marked. Point a has maximal $\text{PAC} = 1$ so it is strongly predicted as spam. It corresponds to a mail containing 1506 capitalized characters, 1488 of which in a single string. Capitalization is more common in spam messages, explaining why it was predicted as spam. Point b is a mail of which 20% consists of the word “free.” This word is more common in spam mails, so its frequent occurrence makes the mail suspicious. In mail c, 7.5% of the characters are “#” which appears more often in spam mails than in non-spam, and 7.5% is the highest percentage in any non-spam mail. Note that b and c are marked as outliers, that is, they are unusual and do not lie well within either class.

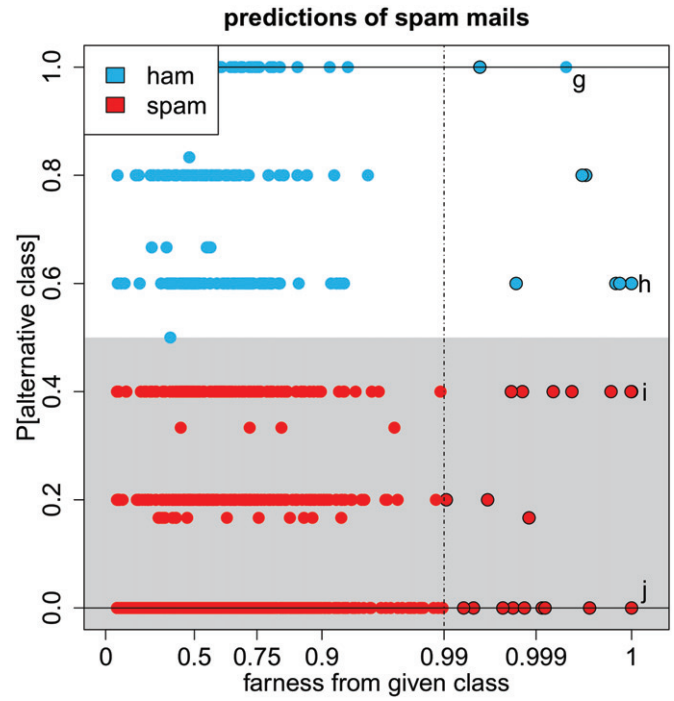


Figure 10. Class map of the spam mails.

The next points are correctly predicted as non-spam. Point d has a high farness because 30% of its characters are exclamation marks, but it also contains some nonspam features such as a high frequency of “re” (since spam mails are usually not replies). Mail e has no special characteristics except for the highest frequency of the word “report.” Mail f is classified perfectly with $\text{PAC} = 0$ but has a high farness, because it contains the number 85 much more often than any other mail in the dataset. This number is characteristic for the non-spam class, since it occurs in all the telephone and fax numbers of the HP labs, including those of the person collecting the data.

Figure 10 is the class map of the spam messages. Also here some points stand out. Mail g has $\text{PAC} = 1$ so the classifier predicted it as non-spam with high conviction. It has a very high frequency of the round bracket character, typically associated with non-spam mails. Similarly, mail h also contains many round brackets, in addition to the word “technology.” On the other hand, it also has an extreme number of “#” symbols (indicative of spam) which explains its high farness. Mail i contains a string of 9989 capitalized characters. This causes it to be correctly classified as spam, but it also gets a high farness. Similarly, mail j is perfectly (since $\text{PAC} = 0$) classified as spam, but has a high farness as well. It contains a very high frequency (almost 20%) of the word “credit.” This is a common word in spam messages, but 20% is unusually high.

5. Support Vector Machines for Two Classes

A support vector machine (SVM) is based on a kernel. Starting from a training set with n objects, the kernel matrix K is of the form $\{K(i, j); i = 1, \dots, n \text{ and } j = 1, \dots, n\}$. The values $K(i, j)$ play the role of inner products, unlike the entries of the dissimilarity matrix in the kNN method which played the role of

distances. The kernel may be derived from a coordinate dataset $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. Many kernel functions exist for that situation. The linear kernel is just $K(i, j) = \langle \mathbf{x}_i, \mathbf{x}_j \rangle$ where \langle, \rangle is the usual inner product. The polynomial kernel is $K(i, j) = (\gamma \langle \mathbf{x}_i, \mathbf{x}_j \rangle + c)^{\text{degree}}$ with tuning constants γ, c and the degree. The radial basis kernel is given by $K(i, j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$. Each of these can be seen as $K(i, j) = \langle \Phi(\mathbf{x}_i), \Phi(\mathbf{x}_j) \rangle$ where $\Phi()$ maps the original data to a feature space. For the linear kernel we can just take $\Phi(\mathbf{x}) = \mathbf{x}$ so there is no transformation. The other two kernels have a feature space of a higher dimension than the original space, in fact for the radial basis kernel that dimension is even infinite. However, the feature space is often left implicit since all computations can be carried out on the kernel matrix K itself. An advantage of kernels is the added flexibility, as it is often easier to separate classes in a higher dimensional feature space than in the original space.

The SVM applies the support vector (SV) classifier in feature space, that is, to data $\mathbf{v} = \Phi(\mathbf{x})$. The SV classifier is a method for $G = 2$ classes which looks for a linear boundary that separates the classes as well as possible. This is achieved by an optimization with a tuning constant `cost` that determines to what extent some points are allowed to be poorly classified. The value of `cost` is typically selected by cross-validation. The end result is an estimated vector $\hat{\beta}$ and intercept $\hat{\beta}_0$ yielding the prediction

$$\hat{g}(\mathbf{v}) = \begin{cases} 1 & \text{if } \hat{\beta}_0 + \langle \mathbf{v}, \hat{\beta} \rangle > 0 \\ 2 & \text{if } \hat{\beta}_0 + \langle \mathbf{v}, \hat{\beta} \rangle \leq 0. \end{cases} \quad (10)$$

One often calls $\hat{\beta}_0 + \langle \mathbf{v}, \hat{\beta} \rangle$ the decision value. Note that also classification by linear discriminant analysis (LDA) is in function of a quantity $\hat{\beta}_0 + \langle \mathbf{v}, \hat{\beta} \rangle$ but it is not the same as it derives from a different optimization.

The standard C++ library for SVM is LIBSVM by Chang and Lin (2019), and it is called by the function `svm` in the R-package `e1071` (Meyer et al. 2020). They compute posterior probabilities $\hat{p}(i, 1)$ by a form of logistic regression with the decision value as the only regressor, and put $\hat{p}(i, 2) = 1 - \hat{p}(i, 1)$. The final prediction is then given by (1). The conditional probability of the best alternative class $\text{PAC}(i)$ then follows immediately by applying (2) and (3).

Farness needs to be defined in relation to how the classifier works. The SVM attempts to linearly separate the points in the feature space. Since the kernel matrix is invariant to multiplying the \mathbf{v}_i by an orthogonal matrix, also the farness needs to be. Orthogonally equivariant linear structures suggest using principal component analysis (PCA). Since the PCA must be applied in the feature space it is, in fact, kernel PCA (KPCA), but here we will write everything in terms of the feature vectors \mathbf{v}_i .

Carrying out a PCA on the \mathbf{v}_i in class 1 yields an estimated center $\hat{\mu}_1$ of class 1, a matrix $\hat{\mathbf{U}}_1$ with the loadings in its columns, and scores $\mathbf{t}_i = (\mathbf{v}_i - \hat{\mu}_1)\hat{\mathbf{U}}_1$ for all $i = 1, \dots, n$. Here we keep all components, so the scores \mathbf{t}_i have the same dimension as the space \mathbf{V}_1 spanned by the \mathbf{v}_i in class 1. We then compute the *score distance* of any object i relative to class 1 as

$$\text{SD}(i, 1) = \sqrt{\sum_j \left(\frac{t_{ij} - \text{median}_h(t_{hj})}{\text{mad}_h(t_{hj})} \right)^2} \quad (11)$$

where `mad` is the median absolute deviation. In this formula h ranges over the members of class 1, whereas i can belong to either class. We compute $\text{SD}(i, 2)$ in the same way.

When the spaces \mathbf{V}_1 and \mathbf{V}_2 are equal, the score distances are all we need. Otherwise, \mathbf{V}_1 and/or \mathbf{V}_2 is a proper subset of the space \mathbf{V} spanned by the \mathbf{v}_i of both classes together. This often happens when using the radial basis kernel because then the dimension of \mathbf{V} can be very high (but not infinite). In such situations a point \mathbf{v}_i of class 1 may not be in \mathbf{V}_2 . We then compute how far \mathbf{v}_i is from \mathbf{V}_2 by the euclidean distance between \mathbf{v}_i and its projection on \mathbf{V}_2 given by

$$\text{OD}(i, 2) = \|\mathbf{v}_i - (\hat{\mu}_2 + (\mathbf{v}_i - \hat{\mu}_2)\hat{\mathbf{U}}_2\hat{\mathbf{U}}_2')\| \quad (12)$$

which is called the *orthogonal distance*, and $\text{OD}(i, 1)$ is computed analogously.

Next, we have to combine the score and orthogonal distances into a single farness measure. For this we first scale all $\text{SD}(i, g)$ by the median of the $\text{SD}(h, g)$ where h ranges over the members of class g . Next, we scale all $\text{OD}(i, g)$ by the median of the $\text{OD}(h^*, g)$, where h^* ranges over all cases *not* belonging to class g (since $\text{OD}(i, g) = 0$ when i belongs to g). Then the distance $D(i, g)$ of an object i to a class g is given by

$$D(i, g) = \sqrt{\text{SD}^2(i, g) + \text{OD}^2(i, g)}. \quad (13)$$

This definition is new, but when the class g has a Gaussian distribution in low dimension, $D(i, g)$ is very similar to the Mahalanobis distance (8). Finally, we estimate the distribution of D as in the previous sections, and apply (4) to obtain $\text{farness}(i)$.

Note that SVM can also be applied to $G > 2$ classes by combining pairwise classifications. Since this technique is also used with other classifiers, it will be discussed separately in Section 7.

We illustrate the class maps for SVM on a benchmark dataset in which the data do not originate from coordinates (measurements). It is one of the datasets collected and studied by Prettenhofer and Stein (2010) and consists of 4000 book reviews on Amazon. The reviews were binned into two categories: positive (with 4 or more stars out of 5) and negative (under 3 stars). The 4000 reviews were split up, the first 2000 forming the training set and the next 2000 the test set.

The data are actual texts, some of them quite long. The kernel matrix was constructed by a string kernel, in fact the function `kernelab::stringdot()` with `type="spectrum"` and `length=7`. Afterward, the SVM was trained with parameter `cost=2` using the R-package `e1071`. The combination of `length` and `cost` was selected by 10-fold cross-validation.

On the training data itself the SVM overfits, with not a single misclassified book review. This is not so surprising since kernel PCA requires as many as 1803 components to explain 95% of the variance, so we are in a truly high-dimensional setting. Since the training data have rather low PAC values, its class maps are relegated to Section A.5 of the supplementary material. On the test data, the trained SVM obtained a correct classification rate of 82%.

The class map of the negative reviews in the test data is shown in Figure 11. About 80% of these reviews were correctly classified as negative (blue). Let us look at some points that stand




Figure 11. Class map of the negative book reviews in the test data.




Figure 12. Class map of the positive book reviews in the test data.

out. Reviews a, b, and c have very high PAC. Excerpts of these reviews are in Table 1. Why was it so hard to classify them correctly? Review a has both positive and negative elements. The reviewer is positive about the author and the stories, but negative about the selection of stories presented in this particular book. Review b is similar in that it praises the author and his past works, but is negative about the current book. Point c corresponds with a review containing some grammatical errors. It also contains a substantial part that is rather neutral and the

Table 1. Excerpts of the negative reviews a to e.

Marked	Excerpts from the book reviews
a	"ptt may well be one of heinein's masterworks" "this collection looses the continuity that made the ptt a great read"
b	"i have liked his other books in the past" "this one didn't have enough insight"
c	"especially aquinas' arguments", "you will definitely not understand"
d	"not at all as interesting as i hoped it would be"
e	"i have read history books more interesting than this book." "i thought that it would be an interesting work."

Table 2. Test data: Excerpts from the positive reviews f to j.

Marked	Excerpts from the book reviews
f	"i have not found any other c++ reference that is as complete or as useful as this books" "this is not an introduction to c++ but rather a reference book"
g	"i felt that the book could have been much better" "do not expect any groundbreaking plot"
h	"this book is not for everyone", "i was disappointed in the end"
i	"i really loved the book", "i just wish the author of the book would"
j	"like the book very muc"

negative opinion in the review is rather indirectly expressed. Review d is definitely negative, but rather short (147 characters), which often leads to larger farness values due to few matches of the string kernel. Finally, review e is again negative, but relatively short and with a substantial neutral/uninformative part.

The class maps of the positive reviews in the test data are in Figure 12. Also here some points stand out from the others, and we have marked a few for illustration. Review f has the highest PAC, indicating that the classifier strongly wants to put it in the negative class. It is nevertheless a positive review, but not unequivocally so since it also indicates for which purposes you should not use this book, as seen in the excerpt in Table 2. Book reviews g and h both have high PAC values, and their farness is relatively high. These two reviews contain a mix of positive and negative comments. Book reviews i and j have high farness values. This can be explained by the fact that these reviews are quite short. Review i has 101 characters, whereas j only contains 22 characters, making it the shortest review in the test data. The mean length of the reviews is 864 characters. Review i is assigned to the negative class, due to a mix of positive and negative comments.

6. Logistic Regression

In the context of logistic regression the classes are typically denoted as 0 and 1 (instead of 1 and 2 as in the other sections), so we are predicting a binary variable y . The coordinates of the \mathbf{x}_i can be continuous and/or binary. The logistic model assumes that the response in \mathbf{x}_i has probability $p(i)$ to be 1 and $1 - p(i)$ to be 0, with $p(i) = \text{logist}(\beta_0 + \langle \mathbf{x}_i, \boldsymbol{\beta} \rangle)$. Here $\text{logist}(z) = \exp(z)/(1 + \exp(z))$ is the logistic function, and the unknown parameters β_0 and $\boldsymbol{\beta}$ are typically estimated by maximum likelihood, yielding $\hat{\beta}_0$ and $\hat{\boldsymbol{\beta}}$. Each object then obtains the posterior probabilities

$$\hat{p}(i, 1) = \text{logist}(\hat{\beta}_0 + \langle \mathbf{x}_i, \hat{\boldsymbol{\beta}} \rangle) \quad \text{and} \quad \hat{p}(i, 0) = 1 - \hat{p}(i, 1). \quad (14)$$

Unlike the SVM with two classes, here the posterior probabilities follow directly from the statistical model. Formulas (1)–(3) are applied in exactly the same way, yielding the probability of the alternative class

$$\text{PAC}(i) = |y_i - \logit(\hat{\beta}_0 + \langle \mathbf{x}_i, \hat{\boldsymbol{\beta}} \rangle)|. \quad (15)$$

The farness measure we use depends on the dimensionality of the regressors \mathbf{x}_i . If the regressors are continuous and low-dimensional we can use the farness as in Section 3 on discriminant analysis, which is based on the Mahalanobis distance. We have the choice between estimating a single $\hat{\boldsymbol{\Sigma}}$ for both classes as in linear DA, or separate $\hat{\boldsymbol{\Sigma}}_0$ and $\hat{\boldsymbol{\Sigma}}_1$ as in quadratic DA.

Note that for multiple linear regression we can replace (15) by the residual $y_i - \hat{y}_i$ and plot it versus the Mahalanobis distance relative to a robust covariance matrix of the \mathbf{x}_i , yielding the outlier map of Rousseeuw and van Zomeren (1990), later extended to multivariate regression (Rousseeuw et al. 2004).

When the \mathbf{x}_i are high dimensional, it is less likely that the covariance matrices of the classes are well-conditioned, and then one can compute farness as in Section 5 on support vector machines. Note that in that situation one may prefer to run sparse logistic regression, for instance using the R-package `glmnet` (Friedman et al. 2020), after which Equation (15) can be applied.

7. Combining Pairwise Classifications

When the data has $G > 2$ labels but the preferred classifier was designed for 2 labels, like the support vector machine, one often resorts to “one versus one” pairwise classifications. In this approach one carries out a binary classification on each pair of classes, yielding $G(G-1)/2$ comparisons.

These computations yield a matrix with $G(G-1)/2$ columns and n rows. Each column corresponds to a pair of classes (g, h) with $g < h$ and contains the estimated probabilities $\hat{p}_{(g,h)}(i, g)$ of the classification into classes g and h . In the past one typically counted how often each class came out on top, which is called majority voting. In recent times the *pairwise coupling* approach has become more popular. In this approach one estimates a set of G posterior probabilities $(\hat{p}(i, 1), \dots, \hat{p}(i, G))$ that match the pairwise probabilities in the sense that

$$\hat{p}_{(g,h)}(i, g) \approx \frac{\hat{p}(i, g)}{\hat{p}(i, g) + \hat{p}(i, h)} \quad (16)$$

for all pairs (g, h) . For this one often uses the second method of Wu, Lin, and Weng (2004) which minimizes a quadratic loss function. It is implemented in the function `couple` in the R-package `kernlab` (Karatzoglou, Smola, and Hornik 2019), with option `coupler="minpair"`.

From the resulting posterior probabilities $\hat{p}(i, g)$ one immediately obtains the multiclass prediction from (1). Applying (2) and (3) then yields $\text{PAC}(i)$.

For drawing class maps, we also need a measure of farness. For this we use one of the methods described in the earlier sections, a summary table of which is given in Section A.6. If the data have continuous coordinates and the dimension is not too high, we can compute the Mahalanobis distance as in (8). If the data have coordinates of mixed types or is given in the form of




Figure 13. Sweets data: Stacked mosaic plot of the classification. When they occur, the products flagged as outliers are shown in dark gray, as an extra class at the top.

a dissimilarity matrix we can run the same farness computation as for the k -nearest neighbor method in Section 4. And if we have high-dimensional continuous data or the input is a kernel matrix, we can apply formulas (11)–(13).

As an illustration of pairwise coupling we analyze the sweets data. This is a subset of the `nutrients_branded` dataset which is publicly available in the R-package `robCompositions` (Templ, Hron, and Filzmoser 2020). It contains data on 9 nutritional values of 804 different sweets sold in Switzerland, which are divided into four categories: “Cookies and Biscuits,” “Milk-based ice cream,” “Cakes and tarts,” and “Creams and puddings.” The nutritional variables are the contents of energy (kcal), protein, water, carbohydrates, sugars, dietary fibers, total fat, saturated fatty acids, and salt.

We fit an SVM model by the function `svm()` in the R-package `e1071` (Meyer et al. 2020), with `kernel="linear"` and `probability=TRUE`. The parameter `cost=10` was selected by 10-fold cross-validation, yielding a misclassification rate of 13%. The function `predict.svm` with again option `probability=TRUE` internally carries out pairwise coupling by the same algorithm of Wu, Lin, and Weng (2004). We used the farness measure of Section 5 on SVM. Figure 14 shows the maps of all classes, whose names were abbreviated to “biscuits,” “ice cream,” “cakes,” and “puddings.”

Figure 13 shows the stacked mosaic plot of this classification. We see at a glance that puddings (the yellow stack at the right) are often predicted as pink, the color of the stack of ice cream. Unlike Figure 1, here the option to show the outliers is switched on. They are visible as an extra predicted class in dark gray at the top for those classes where they occur.

The class maps in Figure 14 contain a lot of structure. Not being experts in nutrition, we have only marked the atypical points we can say something about. The names of the objects a to x are listed in Section A.7 of the supplementary material.

Let us start with the biscuits. A number of points have a PAC over 0.5, and their color indicates they are classified as cakes. Most of these are in fact doughs rather than finished cookies.

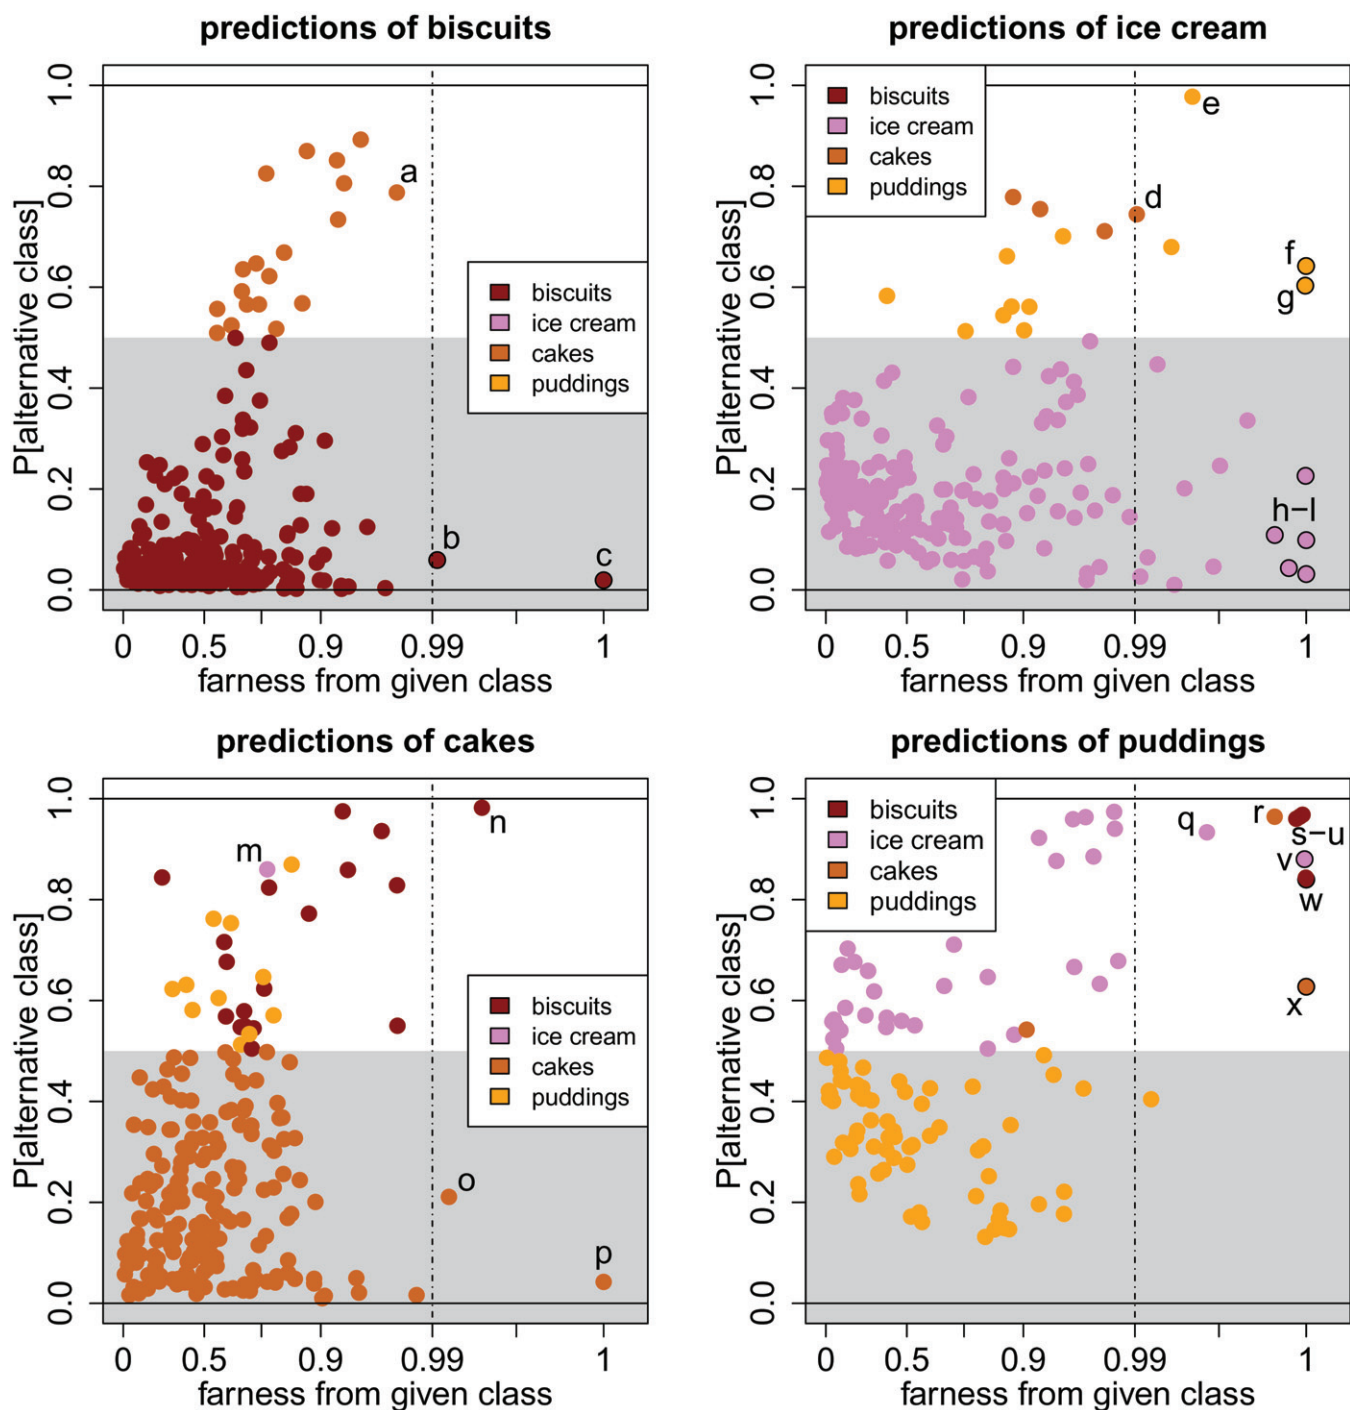


Figure 14. Class maps of all four classes in the sweets data.

They all have a significantly higher water content than most of the biscuits in the data. This results in them being classified as cakes, as cakes are similar to cookies but with a higher water content. The point marked a is a brownie produced by Burger King. Point b has an extreme level of dietary fibers, actually the highest in the whole dataset. This causes it to be flagged as an outlier, as it does not really fit well within any of the classes. Point c is even further away. It corresponds to a biscuit with relatively high salt and very low water content.

Among the ice creams, the products with $PAC > 0.5$ are of two types: one is classified as puddings, the other as cakes. The

points classified as cakes correspond to ice creams that contain a substantial percentage of cookie. Point d is a mini ice cream cone with relatively little ice cream compared with the amount of cookie. The ice creams classified as puddings are nearly all Burger King products, with a much higher salt content than average for ice cream. Point e is an example of such a product, a strawberry sundae. There are also several outliers, as indicated by their black border. Products f and g have $PAC > 0.5$ and are predicted as puddings. Products h to l are classified correctly as ice cream. It turns out that all of these outliers, from f to l, are dietary products. They have very low sugar, very low fat and high

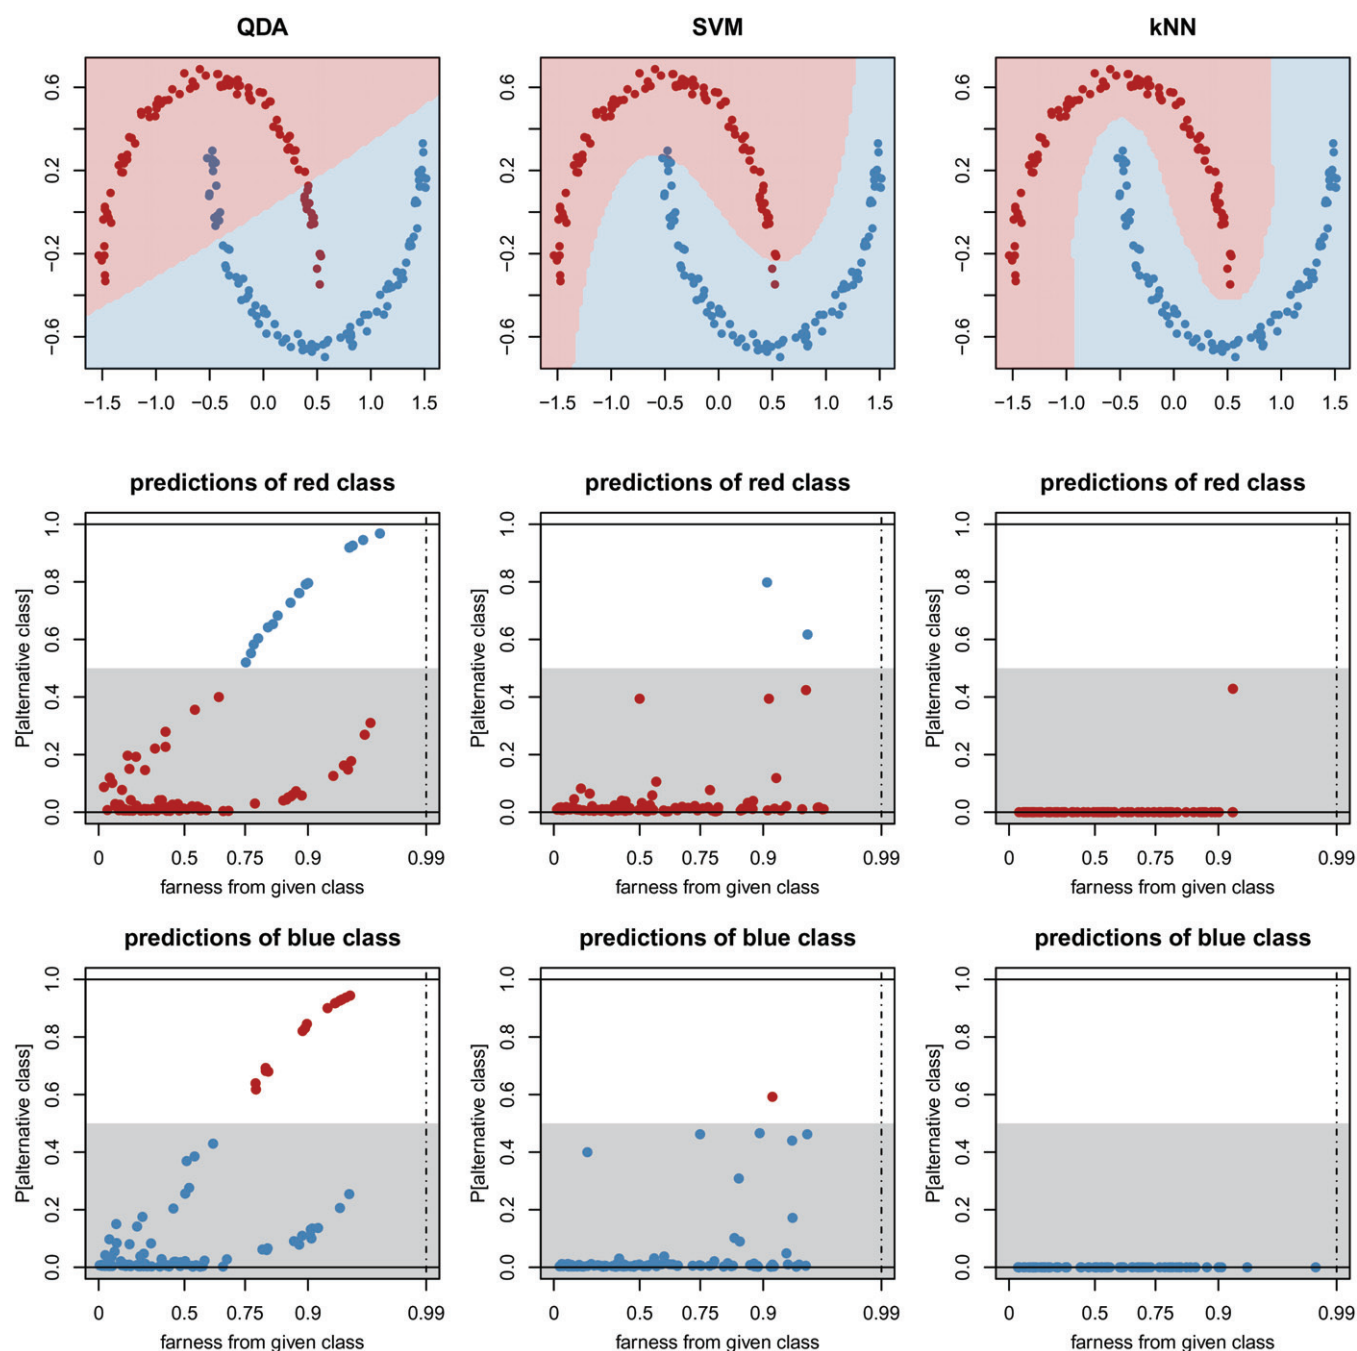


Figure 15. Example where three classifiers are applied to the same dataset. Each column shows the corresponding classification regions, followed by the maps of both classes.

dietary fiber contents. The points \mathfrak{f} and \mathfrak{g} also have unusually high salt levels, making them more similar to puddings.

In the class map of cakes, object \mathfrak{m} is a strawberry tart, with high water content and low sugar relative to most of the cakes. Due to its high water content it is misclassified as ice cream. There is also a group of cakes classified as biscuits. Most of them have very low water contents, that are more common for biscuits than for cakes. These points correspond to either cake mixes or fairly dry cakes based on nuts. The carrot cake \mathfrak{n} is an extreme example of such a product, as it also has very high levels of protein and dietary fibers. The lentil tart \mathfrak{o} is a dietary product. It has many characteristics of a cake, but its salt levels are rather

low and its dietary fiber is very high. Finally, point \mathfrak{p} is a dough, rather than a baked cake.

In the class of puddings, the PAC values are higher on average than in the other classes, so this class is the least well-separated. Most of the misclassified puddings were assigned to the ice cream class. The classifier has a hard time distinguishing between pudding and ice cream. (There is no variable indicating the product's storage temperature.) Most of the puddings misclassified as ice cream are milk-based, causing relatively high fat levels. For instance, product \mathfrak{q} is a chocolate mousse. Point \mathfrak{r} has both a high PAC and a high farness, and is classified as a cake. Nearly all of its nutritional values are exceptional for

pudding but normal for cake. On the store website it looks like a small chocolate cake, so it may actually have been mislabeled. The “puddings” marked s to u are in fact dry pudding mixes rather than finished products. They have extremely low water and extremely high carbohydrate contents, causing them to be misclassified as biscuits, which are drier than puddings. Point v is an outlier, which leans most toward ice cream. It is a dietary product, with extremely low sugars and extremely high dietary fibers. The outlier w is again a dry pudding mix, this time with the highest sugar content in the whole dataset. Finally, pudding x has the highest salt level of any product in the data set. This causes it to be flagged as an outlier and classified as a cake.

8. Comparing Different Classifiers on the Same Data

Even though class maps are primarily intended to learn more about the objects in a classification, it can be interesting to compare the class maps obtained by several classifiers on the same data. To illustrate this we consider a simple toy example with 200 points in two dimensions, for which the ground truth is clear. The data are shown in the top row of Figure 15. The red and blue classes cannot be separated by a straight line.

The left column of Figure 15 shows the result of the QDA classifier. The light red and light blue regions in the top plot were obtained by classifying a dense grid of new data by the trained model, so a new data point in the light blue region will be assigned to the blue class. There are about 16% of misclassifications, indicating that QDA is not a suitable classifier here. Its underlying assumption of roughly elliptical classes indeed does not hold. Below this plot is the class map of the red class. There we see that quite a few points have $PAC > 0.5$ and are assigned to the blue class. The class map has a pattern with thin threads that correspond to the shape of the class. The map of the blue class is similar.

Since the data cannot be separated linearly in the original space, a kernel transform may be helpful. Here, we applied the radial kernel with $\gamma = 0.5$, followed by the SVM with $\text{cost} = 1.0$. For the purpose of this illustration, we did not try to overfit the values of γ and cost . The results are in the middle column of Figure 15. The assignment regions now have very different shapes, and are much better adapted to the particular forms of these classes. Now only three points are misclassified, and they are clearly visible in the class maps. A few other points, that lie close to the boundary between the light blue and light red regions, have substantial $PAC(i)$ values below 0.5.

The third column in Figure 15 shows the result of k -nearest neighbors with $k = 7$. This classifier turns out to work best in this particular example. It has no misclassifications, so one would likely choose kNN here, but the class maps give a finer picture. Not only are there no points with $PAC(i) > 0.5$ (which would be misclassified), but also below 0.5 the map looks quite different from that of SVM. Here only a single point has $PAC(i) > 0$, meaning that its k -neighborhood contains points from the other class. Since it occurs in the map of the red class, it is clear which point it is.

9. Conclusions

The proposed class map reflects two basic notions, for each object in the class. In the vertical direction, it shows the posterior probability of it belonging to its the best alternative class rather than its given class. In the horizontal direction, we see how far the object is from its given class, with a separate plotting symbol when it is far from all classes.

This visualization often provides useful information about the data, as illustrated with examples of discriminant analysis, k -nearest neighbors, support vector machines, and pairwise coupling. It also allows to compare the results obtained by different classifiers.

The methods in this article have been implemented in the R package `classmap` (Raymaekers and Rousseeuw 2021b), which also contains three vignettes reproducing the examples.

Acknowledgments

This research was funded by projects of Internal Funds KU Leuven. We thank to the editor, associate editor, and two reviewers for their constructive comments.

Funding

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

Supplementary Materials

This consists of a text with additional figures and tables, which is also available at <https://arxiv.org/abs/2007.14495>.

References

- Cai, Q., Zhang, C., and Peng, C. (2009), “Analysis of Classification Margin for Classification Accuracy With Applications,” *Neurocomputing*, 72, 1960–1968. [152]
- Chang, C.-C., and Lin, C.-J. (2019), *LIBSVM: A Library for Support Vector Machines*, National Taiwan University. [159]
- Fix, E., and Hodges, J. L. (1951), “Discriminatory Analysis. Nonparametric Discrimination; Consistency Properties,” Report Number 4, Project Number 21-49-004, USAF School of Aviation Medicine, Randolph Field, TX. Reprinted in *International Statistical Review* 57, 238–247 (1989). [154]
- Friedman, J., Hastie, T., Tibshirani, R., Narasimhan, B., Tay, K., and Simon, N. (2020), *Package glmnet: Lasso and Elastic Net Regularized Generalized Linear Models* (CRAN, R package version 4.0-2). [161]
- Friendly, M. (1994), “Mosaic Displays for Multi-Way Contingency Tables,” *Journal of the American Statistical Association*, 89, 190–200. [151]
- Hartigan, J., and Kleiner, B. (1981), “Mosaics for Contingency Tables,” in *Computer Science and Statistics: Proceedings of the 13th Symposium on the Interface*, New York: Springer, pp. 268–273. [151]
- Hastie, T., Tibshirani, R., and Friedman, J. (2017), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 12th printing. Springer Series in Statistics. New York: Springer. [151]
- Karatzoglou, A., Smola, A., and Hornik, K. (2019), *kernlab, the Kernel-Based Machine Learning Lab*. CRAN (R package version 0.9-29). [158,161]
- Kaufman, L., and Rousseeuw, P. (1990), *Finding Groups in Data: An Introduction to Cluster Analysis*, Hoboken, NJ: Wiley-Interscience. [154]
- Le Cun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998), “Gradient-Based Learning Applied to Document Recognition,” in *Proceedings of the IEEE* (Vol. 86), IEEE, pp. 2278–2324. [154]
- Maechler, M., Rousseeuw, P. J., Struyf, A., and Hubert, M. (2019), *cluster*. CRAN (R package version 2.1.0). [156]

- McLachlan, G. J. (2004), *Discriminant Analysis and Statistical Pattern Recognition*, Hoboken, NJ: Wiley-Interscience. [153]
- Meyer, D., Dimitriadou, E., Hornik, K., Weingessel, A., and Leisch, F. (2020), *e1071*. CRAN (R package version 1.7-4). [159,161]
- Prettenhofer, P., and Stein, B. (2010), “Cross-Language Text Classification Using Structural Correspondence Learning,” in *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pp. 1118–1127. [159]
- Raymaekers, J., and Rousseeuw, P. J. (2020), *cellWise: Analyzing Data With Cellwise Outliers* (CRAN, R package version 2.2.3). [152]
- (2021a), “Transforming Variables to Central Normality,” *Machine Learning*, DOI: 10.1007/s10994-021-05960-5. [152]
- (2021b), *classmap: Visualizing Classification Results*. Available at <https://CRAN.Rproject.org/package=classmap>. [164]
- Rousseeuw, P. J. (1987), “Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis,” *Journal of Computational and Applied Mathematics*, 20, 53–65. [152]
- Rousseeuw, P. J., Van Aelst, S., Van Driessen, K., and Agulló, J. (2004), “Robust Multivariate Regression,” *Technometrics*, 46, 293–305. [161]
- Rousseeuw, P. J., and van Zomeren, B. (1990), “Unmasking Multivariate Outliers and Leverage Points,” *Journal of the American Statistical Association*, 85, 633–651. [161]
- Templ, M., Hron, K., and Filzmoser, P. (2020), *robCompositions: An R-package for Robust Statistical Analysis of Compositional Data* (CRAN, R package version 2.2.1). [161]
- Wouters, N., De Ketelaere, B., Deckers, T., Baerdemaeker, J. D., and Saeys, W. (2015), “Multispectral Detection of Floral Buds for Automated Thinning of Pear,” *Computers and Electronics in Agriculture*, 113, 93–103. [151]
- Wu, T. F., Lin, C. J., and Weng, R. C. (2004), “Probability Estimates for Multi-Class Classification by Pairwise Coupling,” *Journal of Machine Learning Research*, 5, 975–1005. [161]