

# A Two-Stage Bayesian Semiparametric Model for Novelty Detection with Robust Prior Information\*

Francesco Denti  
University of California Irvine  
fdenti@uci.edu

Andrea Cappozzo  
Politecnico di Milano  
andrea.cappozzo@polimi.it

Francesca Greselin  
University of Milano-Bicocca  
francesca.greselin@unimib.it

Novelty detection methods aim at partitioning the test units into already observed and previously unseen patterns. However, two significant issues arise: there may be considerable interest in identifying specific structures within the novelty, and contamination in the known classes could completely blur the actual separation between manifest and new groups. Motivated by these problems, we propose a two-stage Bayesian semiparametric novelty detector, building upon prior information robustly extracted from a set of complete learning units. We devise a general-purpose multivariate methodology that we also extend to handle functional data objects. We provide insights on the model behavior by investigating the theoretical properties of the associated semiparametric prior. From the computational point of view, we propose a suitable  $\xi$ -sequence to construct an independent slice-efficient sampler that takes into account the difference between manifest and novelty components. We showcase our model performance through an extensive simulation study and applications on both multivariate and functional datasets, in which diverse and distinctive unknown patterns are discovered.

**Keywords:** Bayesian mixture model, Bayesian nonparametrics, Minimum Regularized Covariance Determinant, Novelty detection, Slice sampler.

---

\*<https://doi.org/10.1007/s11222-021-10017-7>

# 1 Introduction

Supervised classification techniques aim at predicting a qualitative output for a test set by learning a classifier on a fully-labeled training set. To this extent, classical methods assume that the labeled units are realizations from each and every sub-groups in the target population. However, many real datasets contradict these assumptions. As an example, one may think about an evolving ecosystem in which novel species are likely to appear over time. In other words, basic classifiers cannot handle the presence of previously unobserved - or hidden - classes in the test set. Novelty detection methods, also known as *adaptive* methods, address this issue by modeling the presence of classes in the test set that have not been previously observed in the training. Relevant examples of this type of data analysis include, but are not limited to, radar target detection (Carpenter et al., 1997), detection of masses in mammograms (Tarassenko et al., 1995), handwritten digit recognition (Tax and Duin, 1998) and e-commerce (Manikopoulos and Papavassiliou, 2002), for which labeled observations may not be available for every group.

Within the model-based family of classifiers, adaptive methods recently appeared in the literature. Miller and Browning, 2003 pioneer a mixture model methodology for class discovery. Bouveyron, 2014 introduces an adaptive classifier in which two algorithms, based respectively on transductive and inductive learning, are devised for inference. More recently, Fop et al., 2021 extend the original work of Bouveyron, 2014 by accounting for unobserved classes and extra variables in high-dimensional discriminant analysis.

Classical model-based classifiers are not robust, as they lack the capability of handling outlying observations in the training and in the test set. On the one hand, the presence of outliers in the training set can significantly alter the learning phase, resulting in poorly representative classes and therefore jeopardizing the entire classification process. In the training set, we identify as outliers units with implausible labels and/or values. Cappozzo et al., 2020 extend the work of Bouveyron, 2014 addressing this problem by using a robust estimator that relies on impartial trimming (Gordaliza, 1991). In short, the most unlikely data points under the currently estimated model are discarded. On the other hand, dealing with outliers in the test set is a more delicate task. Ideally, we would like to distinguish between *novelties*, i.e., test observations displaying a common, specific pattern, and *anomalies*, i.e., test observations that can be regarded as noise. While the distinction between novel and anomalous entities is most often apparent in practice, there exist some circumstances under which such separation is vague and somewhat philosophical. Let us go back to the aforementioned evolving ecosystem example. It may happen that at an early instant, a real novelty is mistaken to be mere noise due to its embryonic stage. Contrarily, if we fitted the same model at a later time point, the increased sample size could be sufficient to acknowledge an actual novel species.

To address the discussed challenges, we propose a two-stage Bayesian semiparametric novelty detector. We devise our model to sequentially handle the outliers in the training set and the latent classes in the test set. In the first stage, we learn the main characteristics of the known classes (for example, their mean and variance) from the labeled dataset using robust procedures. In the second phase, we fit a Bayesian semiparametric mixture

of known groups and a novelty term to the test set. We use the training insights to elicit informative priors for the known components, modeled as Gaussian distributions. The novelty term is instead captured via a flexible Dirichlet Process mixture: this modeling choice reflects the lack of knowledge about its distributional properties and overcomes the problematic and unnatural a priori specification of its number of components. We call our proposal Bayesian Robust Adaptive model for Novelty Detection, hereafter denoted as Brand. Essentially, Stage II of Brand is formed by two nested mixtures, which can provide uncertainty quantification regarding the two partitions of interest. First, Brand separates the entire test set into known components and a novelty term. Secondly, the novel data points can be a posteriori clustered into different sub-components. “True novelties” and anomalies may be distinguished, based on clusters cardinality.

The rest of the article proceeds as follows. In Section 2 we present our two-stage methodology for novelty detection. We dedicate Section 3 to the investigation of the random measures clustering properties induced by our model. In Section 4, we propose an extension of the multivariate model, delineating a novelty detection method suitable for functional data. Section 5 discusses posterior inference, while in Section 6 we present an extensive simulation study and applications to multivariate and functional data. Concluding remarks and further research directions are outlined in Section 7.

## 2 A Two-Stage Bayesian procedure for Novelty Detection

Given a classification framework, consider the complete collection of learning units  $\mathbf{X} = \{(\mathbf{x}_n, \mathbf{l}_n)\}_{n=1}^N$ , where  $\mathbf{x}_n$  denotes a  $p$ -variate observation and  $\mathbf{l}_n = j \in \{1, \dots, J\}$  its associated group label. Both terms are directly available and the distinct values in  $\mathbf{l}_n$ ,  $n = 1, \dots, N$  represent the  $J$  *observed classes* with subset sizes  $n_1, \dots, n_J$ . Correspondingly, let  $\mathbf{Y} = \{(\mathbf{y}_m, \mathbf{z}_m)\}_{m=1}^M$  be the test set where, differently from the usual setting in semisupervised learning, the unknown labels  $\mathbf{z}_m$  could belong to a set that encompasses more elements than  $\{1, \dots, J\}$ . That is, a countable number of extra classes may be “hidden” in the test with no prior information available on their magnitude or on their structure. Therefore, it is reasonable to account for the novelty term via a single flexible component from which a dedicated post-processing procedure may reveal circumstantial patterns (see Section 2.3). Both  $\mathbf{x}_n$  and  $\mathbf{y}_m$  are independent realizations of a continuous random vector (or function, see Section 4)  $\mathcal{X}$ , whose conditional distribution varies according to the associated class labels. In the upcoming Sections, we assume that each observation in class  $j$  is independent multivariate Gaussian, having density  $\phi(\cdot | \Theta_j)$  with location-scale parameter  $\Theta_j = (\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$ , where  $\boldsymbol{\mu}_j \in \mathbb{R}^p$  denotes the mean vector and  $\boldsymbol{\Sigma}_j$  the corresponding covariance matrix. This allows for the automatic implementation of standard powerful methods in the training information extraction (see Section 2.1). Notwithstanding, the proposed methodology is general enough that it can be easily extended to deal with different component distributions.

Our modeling purpose is to classify the data points of the test set either into one of the  $J$  observed classes or into the novel component. At the same time, we investigate the presence of homogeneous groups in the novelty term, discriminating between unseen

classes and outliers. To do so, we devise a two-stage strategy. The first phase, described in Section 2.1, relies on a class-wise robust procedure for extracting prior information from the training set. Then, we fit a Bayesian semiparametric mixture model to the test units. A full account of its definition is reported in Section 2.2. A diagram summarizing our modeling proposal is reported in Figure 1.

## 2.1 Stage I: Robust extraction of prior information

The first step of our procedure is designed to obtain reliable estimates  $\hat{\Theta}_j$  for the parameters of the observed class  $j$ ,  $j = 1 \dots, J$ , from the learning set. To this aim, one could employ standard methods as the maximum likelihood estimator, or different posterior estimates under the Bayesian framework. Nonetheless, these standard approaches are not robust against contamination, and the presence of only a few outlying points could entirely bias the subsequent Bayesian model, should the informative priors be improperly set. We report a direct consequence of this undesirable behavior in the simulation study of Section 6.1. Therefore, we opt for more sophisticated alternatives to learn the structure of the known classes, employing methods that can deal with outliers and label noise. Particularly, the selected methodologies involve the Minimum Covariance Determinant (MCD) estimator (Rousseeuw, 1984; Hubert et al., 2018) and, when facing high-dimensional data (as in the functional case of Section 6.3), the Minimum Regularized Covariance Determinant (MRCD) estimator (Boudt et al., 2020). Clearly, at this stage, one can use any robust estimators of multivariate scatter and location for solving this problem: see, for instance, the comparison study reported in Maronna and Yohai (2017) for a non-exhaustive list of suitable candidates.

We decide to rely on the MCD and MRCD for their well-established efficacy in the classification framework (Hubert and Van Driessen, 2004) and direct availability of fast algorithms for inference, readily implemented in the `rrcov` R package (Todorov and Filzmoser, 2009). We briefly recall the main MCD and MRCD features in the remaining part of this section. For a thorough treatment the interested reader is referred to Hubert and Debruyne, 2010 and Boudt et al., 2020, respectively.

The MCD is an affine equivariant and highly robust estimator of multivariate location and scatter, for which a fast algorithm is available (Rousseeuw and Driessen, 1999). The raw MCD estimator with parameter  $\eta^{MCD} \in [0.5, 1]$  such that  $\lfloor (n + p + 1)/2 \rfloor \leq \lfloor \eta^{MCD} N \rfloor \leq N$  defines the following location and dispersion estimates:

- $\hat{\boldsymbol{\mu}}^{MCD}$  is the mean of the  $\lfloor \eta^{MCD} N \rfloor$  observations for which the determinant of the sample covariance matrix is minimal
- $\hat{\boldsymbol{\Sigma}}^{MCD}$  is the corresponding covariance matrix, multiplied by a consistency factor  $c_0$  (Croux and Haesbroeck, 1999)

with  $\lfloor \cdot \rfloor$  denoting the floor function. The MCD is a consistent, asymptotically normal and highly robust estimator with bounded influence function and breakdown value equal to  $(1 - \lfloor \eta^{MCD} N \rfloor / N)\%$  (Butler et al., 1993; Cator and Lopuhaä, 2012). However, a major drawback is its inapplicability when the data dimension  $p$  exceeds the subset size

$[\eta^{MCD}N]$  as the covariance matrix of any  $[\eta^{MCD}N]$ -subset becomes singular. This situation appears ever so often in our context, as the MCD is group-wise applied to the observed classes in the training set, such that it is sufficient to have

$$p > \min_{n_j, j=1, \dots, J} [\eta^{MCD}n_j]$$

for the MCD solution to be ill-defined. To overcome this issue, Boudt et al., 2020 introduced the MRCD estimator. The main idea is to replace the subset-based covariance estimation with a regularized one, defined as a weighted average of the sample covariance on the  $[\eta^{MCD}N]$ -subset and a predetermined positive definite target matrix. The MRCD estimator is defined as the multivariate location and regularized scatter based on the  $[\eta^{MCD}N]$ -subset that makes its overall determinant the smallest. The MRCD preserves the good breakdown properties of its non-regularized counterpart, and besides, it is applicable in high-dimensional problems where  $[\eta^{MCD}N]$  is possibly smaller than  $p$ .

The first phase of our two-stage modeling thus works as follows: considering the available labels  $\mathbf{l}_n$ ,  $n = 1, \dots, N$  we apply the MCD (or MRCD) estimator within each class to extract  $\hat{\boldsymbol{\mu}}_j^{MCD}$  and  $\hat{\boldsymbol{\Sigma}}_j^{MCD}$ ,  $j = 1 \dots, J$ . For ease of notation, we use superscript ‘MCD’ for the robust estimates even when we consider its regularized version. In general, if the sample size is large enough, the MCD solution is preferred. There is no reason for  $\eta^{MCD}$  to be the same in all observed classes. If a group is known a priori to be particularly outliers-sensitive, one should set its associated MCD subset size to a smaller value than the remaining ones. However, since this type of information is seldom available, we subsequently let  $\eta_j^{MCD} = \eta^{MCD}$  for all classes in the learning set. This concludes the first stage: the retained estimates are then incorporated in the Bayesian model for the second stage, presented in Section 2.2. The robust knowledge extracted from  $\mathbf{X}$  is treated as a source of reliable prior information, eliciting informative hyperparameters. In this way, outliers and label noise that might be present in the labeled units will not bias the initial beliefs for the known groups in the second stage, which is the main methodological contribution of the present manuscript.

## 2.2 Stage II: BNP novelty detection in test data

We assume that each observation in the test set is generated according to a mixture of  $J+1$  elements:  $J$  multivariate Gaussians  $\phi(\cdot | \boldsymbol{\Theta}_j)$  that have been observed in the learning set, and an extra term  $f^{nov}$  called *novelty* component. In formulas:

$$\mathbf{y}_m | \boldsymbol{\pi}, \boldsymbol{\Theta}_j, f^{nov} \sim \sum_{j=1}^J \pi_j \phi(\cdot | \boldsymbol{\Theta}_j) + \pi_0 f^{nov}. \quad (1)$$

We define  $\boldsymbol{\pi} = \{\pi_j\}_{j=1}^J$ , where  $\pi_j$  denotes the prior probability of the observed class  $j$  (already present in the learning set), while  $\pi_0$  is the probability of observing some novelty. Of course,  $\sum_{j=0}^J \pi_j = 1$ . To reflect our lack of knowledge on the novelty component  $f^{nov}$ , we employ a Bayesian nonparametric specification. In particular, we

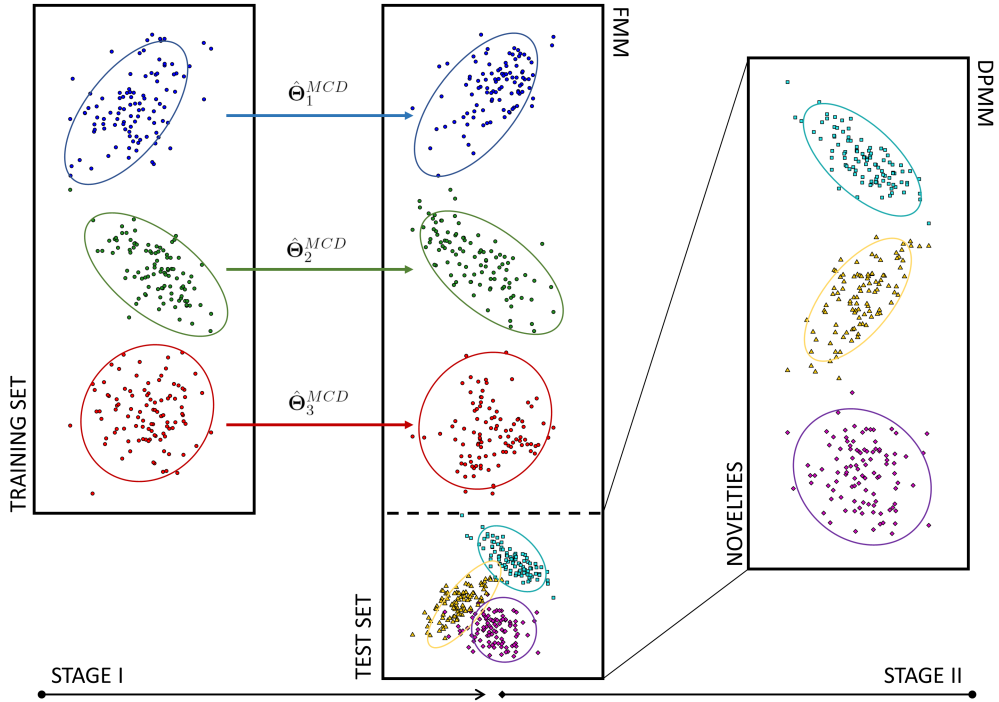


Figure 1: A diagram that summarizes Brand two-stage structure. In Stage I, robust information extraction (via the MCD estimator) is performed and subsequently used to elicit the priors for the mixture model in Stage II. In the second stage, a finite mixture model (FMM) is fitted to the data, distinguishing among known components and novelties. The novelty term is modeled with a Dirichlet Process mixture model (DPMM).

resort to the Dirichlet Process mixture model (DPMM) of Gaussians densities (Lo, 1984; Escobar and West, 1995) imposing the following structure:

$$f^{nov} = \int \phi(\cdot | \Theta^{nov}) G(d\Theta^{nov}), \quad G \sim DP(\gamma, H), \quad (2)$$

where  $DP(\gamma, H)$  denotes a Dirichlet Process with concentration parameter  $\gamma$  and base measure  $H$  (Ferguson, 1973). Adopting Sethuraman's Stick Breaking construction (Sethuraman, 1994), we can express the likelihood as follows:

$$\mathcal{L}(\mathbf{y} | \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\omega}) = \prod_{m=1}^M \left[ \sum_{j=1}^J \pi_j \phi(\mathbf{y}_m | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) + \pi_0 \sum_{h=1}^{\infty} \omega_h \phi(\mathbf{y}_m | \boldsymbol{\mu}_h^{nov}, \boldsymbol{\Sigma}_h^{nov}) \right]. \quad (3)$$

The term  $\sum_{h=1}^{\infty} \omega_h \phi(\cdot | \Theta_h^{nov})$  represents a Dirichlet Process realization convoluted with a Normal kernel, for flexibly modeling a potentially infinite number of hidden classes and/or outlying observations. The following prior probabilities for the parameters com-

plete the Bayesian model specification:

$$\begin{aligned}
\Theta_j &= (\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) \sim P_j^{Tr}, & j &= 1, \dots, J, \\
\Theta_h^{nov} &= (\boldsymbol{\mu}_h^{nov}, \boldsymbol{\Sigma}_h^{nov}) \sim H, & h &= 1, \dots, \infty, \\
\boldsymbol{\pi} &\sim Dir(a_0, a_1, \dots, a_J), & \boldsymbol{\omega} &\sim SB(\gamma).
\end{aligned} \tag{4}$$

Values  $a_1, \dots, a_J$  are the hyper-parameters of a Dirichlet distribution on the known classes. We can exploit the learning set to determine reasonable values of such hyper-parameters by setting  $a_j = n_j/N$ . The quantity  $a_0$  determines the initial prior belief on how much novelty we are expecting to discover in the test set. Generally, the parameter controlling the novelty proportion  $a_0$  is a priori considered to be small. We adopt a conjugate Normal-inverse-Wishart (NIW) prior for both the location-scale parameters of the manifest and the novel classes. For each of the known groups, we assume that

$$P_j^{Tr} \equiv NIW\left(\hat{\boldsymbol{\mu}}_j^{MCD}, \lambda^{Tr}, \nu^{Tr}, \hat{\boldsymbol{\Sigma}}_j^{MCD}\right), \quad j = 1, \dots, J$$

where  $\hat{\boldsymbol{\mu}}_j^{MCD}$  and  $\hat{\boldsymbol{\Sigma}}_j^{MCD}$  are the MCD robust estimates obtained in Stage I. At the same time, the precision parameter  $\lambda^{Tr}$  and the degrees of freedom  $\nu^{Tr}$  are treated as tuning parameters to enforce high mass around the robust estimates. By letting these two parameters go off to infinity, we can also recover the degenerate case  $P_j^{Tr} = \delta_{\hat{\Theta}_j}$  where the Dirac's delta denotes a point mass centered in  $\hat{\Theta}_j$ . That is, the prior beliefs extracted from the training set can be flexibly updated by gradually transitioning from transductive to inductive inference by increasing  $\lambda^{Tr}$  and  $\nu^{Tr}$  (Bouveyron, 2014). Similarly, we set  $H \equiv NIW(\mathbf{m}_0, \lambda_0, \nu_0, \mathbf{S}_0)$ , where the hyperparameters are chosen to induce a flat prior for the novel components. Lastly, with  $\boldsymbol{\omega} \sim SB(\gamma)$  we denote the vector of Stick-Breaking weights, composed of elements defined as

$$w_k = v_k \prod_{l < k} (1 - v_l), \quad v_k \sim Beta(1, \gamma). \tag{5}$$

It is well known that, under the DP specification, the expected number of clusters induced in the novelty term grows as  $\gamma \log M$ . We choose the DP mostly for computational convenience: if more flexibility is required, Brand can easily be adapted to accommodate different nonparametric priors, such as the Pitman-Yor process (Pitman, 1995; Pitman and Yor, 1997) or the geometric process and its extensions (De Blasi et al., 2020). To facilitate posterior inference given the specification in (4), we consider the following *complete likelihood*:

$$\begin{aligned}
\mathcal{L}(\mathbf{y}, \boldsymbol{\alpha}, \boldsymbol{\beta} | \boldsymbol{\pi}, \boldsymbol{\Theta}, \boldsymbol{\omega}) &= \prod_{m=1}^M \left[ \pi_{\alpha_m} \mathbb{1}_{\{\alpha_m > 0 \cap \beta_m = 0\}} + \right. \\
&\quad \left. + \pi_0 \mathbb{1}_{\{\alpha_m = 0 \cap \beta_m > 0\}} \omega_{\beta_m} \right] \times \\
&\quad \times \phi(\mathbf{y}_m | \boldsymbol{\Theta}_m^*), \\
\boldsymbol{\Theta}_m^* &= \boldsymbol{\Theta}_{\alpha_m} \mathbb{1}_{\{\beta_m = 0\}} + \boldsymbol{\Theta}_{\beta_m}^{nov} \mathbb{1}_{\{\alpha_m = 0\}}
\end{aligned} \tag{6}$$

where  $\alpha_m \in \{0, \dots, J\}$  and  $\beta_m \in \{0, \dots, \infty\}$  are latent variables identifying the unobserved group membership for  $\mathbf{y}_m$ ,  $m = 1, \dots, M$ . In details, the former identifies whether observation  $m$  is a novelty ( $\alpha_m = 0$ ) or not ( $\alpha_m > 0$ ), whereas the latter defines, within the novelty subset, the resulting data partition ( $\beta_m > 0$ ). To complete the specification, we set  $\omega_0 = 1$ .

Lastly, we want to underline that there might be some cases where the number of novelty groups is known to be bounded and does not grow with the sample size as in the DP case. In those situations, an appealing alternative to the DPMM is the Sparse Mixture Model, studied by Rousseau and Mengersen (2011) and recently investigated in Malsiner-Walli et al. (2016).

### 2.3 Distinguishing novelties from anomalies

The advantage of employing a DPMM for the novelty part is twofold: on the one hand, all the data coming from unseen components are modeled with a unique, flexible density. On the other hand, the clustering naturally induced by the DPMM favors the separation of the novelty component into actual unseen classes and outlying units. More specifically, since the concept of an outlier does not possess a rigorous mathematical definition (Ritter, 2014), the estimated sample sizes of the discovered classes act as an appropriate feature for discriminating between scattered outlying units and actual hidden groups. That is, if a component  $\phi(\cdot | \Theta_h^{nov})$  fits only a small number of data points, we can regard those units as outliers. Similarly, we assume to have discovered an extra class whenever it possesses a substantial structure. In real applications, domain-expert supervision will always be crucial for class interpretation when extra groups are believed to have been detected. While the mixture between known and novel distributions is identifiable and not subjected to the label switching problem, the same cannot be said about the DP component modeling the novelty density. To recover a meaningful estimate for the partition of points regarded as novel ( $\beta_m > 0$ ) we first compute the pairwise coclustering matrix  $\mathcal{P} = \{p_{m,m'}\}$ , whose entry  $p_{m,m'}$  denotes the probability that  $\mathbf{y}_m$  and  $\mathbf{y}_{m'}$  belong to the same cluster. We then retrieve the best partition minimizing the Variation of Information (VI) criterion, as suggested in Wade and Ghahramani (2018). More details on how to post-process the MCMC output are given in Section 5.

## 3 Properties of the proposed semiparametric prior

We now investigate the properties of the underlying random mixing measure induced by the model specification we presented in the previous section. All the proofs are deferred to the Supplementary Material. We start by noticing that model in (3)-(4) can be generalized in the following hierarchical form, which highlights the dependence on a



discrete random measure  $\tilde{p}$ :

$$\begin{aligned}
\mathbf{y}_m | \Theta_m &\sim N(\Theta_m) \quad \Theta_m | \tilde{p} \stackrel{i.i.d.}{\sim} \tilde{p} \\
\tilde{p} &= \sum_{j=1}^J \pi_j \delta_{\Theta_j} + \pi_0 \left[ \sum_{h=1}^{+\infty} \omega_h \delta_{\Theta_h^{nov}} \right] \\
(\pi_0, \pi_1, \dots, \pi_J) &\sim Dir(a_0, a_1, \dots, a_J) \quad \omega \sim SB(\gamma) \\
\Theta_j &\sim P_j^{Tr} \quad \Theta_h^{nov} \sim H.
\end{aligned} \tag{7}$$

From (7) we can see how our model is an extension of the contaminated informative priors proposed in Scarpa and Dunson (2009), where the authors propose to juxtapose a single atom to a DP. To simplify the exposition of the results in this section, without loss of generality, we assume that both  $\Theta_j$  and  $\Theta_h^{nov}$  are univariate random variables. Consequently, we suppose that each  $P_j^{Tr}$  is a probability distribution with mean  $\mu_j$ , second moment  $\mu_{j,2}$  and variance  $\sigma_j^2$ ,  $j = 1, \dots, J$ . Similarly, let  $\mathbb{E}[\Theta_h^{nov}] = \mu_0$ ,  $\mathbb{V}[\Theta_h^{nov}] = \sigma_0^2 \forall h \geq 1$  and  $a = \sum_{j=0}^J a_j$ . For all  $m \in \{1, \dots, M\}$ , we can prove that

$$\mathbb{E}[\Theta_m] = \sum_{j=0}^J \frac{a_j}{a} \mu_j, \quad \mathbb{V}[\Theta_m] = \sum_{j=0}^J \frac{a_j}{a} \left( \mu_{j,2} - \frac{a_j}{a} \mu_j^2 \right) - 2 \sum_{l>j \geq 0} \frac{a_j a_l}{a^2} \mu_l \mu_j.$$

The overall variance can also be written in terms of variances of every observed mixture components:

$$\mathbb{V}[\Theta_m] = \sum_{j=0}^J \frac{a_j}{a} \left( \sigma_j^2 + \left( 1 - \frac{a_j}{a} \right) \mu_j^2 \right) - 2 \sum_{l>j \geq 0} \frac{a_j a_l}{a^2} \mu_l \mu_j.$$

The previous expressions are important to compute the covariance between the two random elements  $\Theta_m$  and  $\Theta_{m'}$ , which helps to understand the behavior of  $\tilde{p}$ . Consider a vector  $\boldsymbol{\varrho} = \{\varrho_j\}_{j=0}^J$ , with the first entry equal to  $\frac{1}{1+\gamma}$  and the remaining entries equal to 1. Then,

$$\begin{aligned}
Cov_\gamma(\Theta_m, \Theta_{m'}) &= \sum_{j=0}^J \left( \frac{a_j(a_j+1)}{a(a+1)} \varrho_j \mu_{j,2} - \frac{a_j^2}{a^2} \mu_j^2 \right) + \\
&\quad - \frac{2}{a^2(a+1)} \sum_{j>l \geq 0} a_j a_l \mu_j \mu_l + \\
&\quad + \frac{a_0(a_0+1)}{a(a+1)} \frac{\gamma}{1+\gamma} \mu_0^2.
\end{aligned} \tag{8}$$

The covariance is composed of three terms. In the first two, the seen and unseen components have the same influence. The last term is non-negative and entirely determined by quantities linked to the novel part of the model. Notice that if  $\gamma \rightarrow 0$  the covariance becomes

$$Cov_0(\Theta_m, \Theta_{m'}) = \sum_{j=0}^J \left( \frac{a_j(a_j+1)}{a(a+1)} \mu_{j,2} - \frac{a_j^2}{a^2} \mu_j^2 \right) - \frac{2}{a^2(a+1)} \sum_{j>l \geq 0} a_j a_l \mu_j \mu_l$$

which is the same covariance we would obtain if  $\tilde{p} = \tilde{p}_0 \equiv \sum_{j=0}^J \pi_j \delta_{\Theta_j}$ , i.e., if we were dealing with a “standard” mixture model with  $J + 1$  components. This implies that (8) can be rewritten as

$$\text{Cov}_\gamma(\Theta_m, \Theta_{m'}) = \text{Cov}_0(\Theta_m, \Theta_{m'}) - \frac{a_0(a_0 + 1)}{a(a + 1)} \frac{\gamma}{1 + \gamma} \sigma_0^2,$$

which leads to a nice interpretation. The introduction of novelty atoms decreases the “standard” covariance. This effect gets stronger as the prior weight given to the novelty component  $a_0$ , the dispersion of the base measure  $\sigma_0^2$  and/or the concentration parameter  $\gamma$  increases.

Given the discrete nature of  $\tilde{p}$ , we can expect ties between realizations sampled from this measure, say  $\Theta_m$  and  $\Theta_{m'}$ . Therefore, we can compute the probability of obtaining a tie as:

$$\mathbb{P}(\Theta_m = \Theta_{m'}) = \sum_{j=1}^J \frac{a_j(a_j + 1)}{a(a + 1)} + \frac{a_0(a_0 + 1)}{a(a + 1)} \cdot \frac{1}{1 + \gamma} = \sum_{j=0}^J \frac{a_j(a_j + 1)}{a(a + 1)} \varrho_j, \quad (9)$$

where the contribution to this probability of the novelty terms is multiplicatively reduced by a factor that depends on the inverse of the concentration parameter. If a priori we expect a large number of clusters in the novelty term (large  $\gamma$ ), the probability of a tie reduces. Indeed, some noticeable limiting cases arise:

$$\lim_{\gamma \rightarrow +\infty} \mathbb{P}(\Theta_m = \Theta_{m'}) = \sum_{j=1}^J \frac{a_j(a_j + 1)}{a(a + 1)}, \quad \lim_{\gamma \rightarrow 0} \mathbb{P}(\Theta_m = \Theta_{m'}) = \sum_{j=0}^J \frac{a_j(a_j + 1)}{a(a + 1)}.$$

If  $\gamma \rightarrow 0$  we obtain a finite mixture of  $J + 1$  components. Conversely,  $\gamma \rightarrow +\infty$  leads to the case of a DP with numerous atoms characterized by similar probability, hence annihilating the contribution to the probability of the novelty term. Moreover, suppose we rewrite the distribution of  $\pi$  as  $\text{Dir}\left(\frac{a_0}{J+1}, \frac{\tilde{a}}{J+1}, \dots, \frac{\tilde{a}}{J+1}\right)$ . In this case, the hyper-parameters relative to the observed groups are assumed equal to  $\tilde{a}$ . Then, we obtain  $a = \frac{a_0 + J\tilde{a}}{J+1}$ , and

$$\mathbb{P}(\Theta_m = \Theta_{m'}) = \frac{\frac{J\tilde{a}}{J+1} \left(\frac{\tilde{a}}{J+1} + 1\right)}{\frac{a_0 + J\tilde{a}}{J+1} \left(\frac{a_0 + J\tilde{a}}{J+1} + 1\right)} + \frac{\frac{a_0}{J+1} \left(\frac{a_0}{J+1} + 1\right)}{\frac{a_0 + J\tilde{a}}{J+1} \left(\frac{a_0 + J\tilde{a}}{J+1} + 1\right)} \cdot \frac{1}{1 + \gamma}. \quad (10)$$

As  $J$  increases, the second part of (10) vanishes. Accordingly, if we suppose an unbounded number of observed groups letting  $J \rightarrow \infty$ , then we have

$$\mathbb{P}(\Theta_m = \Theta_{m'}) = 1/(1 + \tilde{a})$$

as in the classical DP case, and the model loses its ability to detect novel instances.

## 4 Functional Novelty Detection

The modeling framework introduced in Section 2 is very general and can be easily modified to handle more complex data structures. In this section, we develop a methodology for functional classification that allows novelty functional detection, building upon model (3)-(4). We hereafter assume that our training and test instances are error-prone realizations of a univariate stochastic process  $\mathcal{X}(t)$ ,  $t \in \mathcal{T}$  with  $\mathcal{T} \subset \mathbb{R}$ .

Recently, numerous authors have contributed to the area of Bayesian nonparametric functional clustering (see, for example Bigelow and Dunson, 2009; Petrone et al., 2009; Rodriguez and Dunson, 2014; Rigon, 2019). Canale et al. (2017) propose a Pitman-Yor mixture with a spike-and-slab base measure to effectively model the daily basal body temperature in women by including the a priori known distinctive biphasic trajectory that characterizes healthy beings. Instead of modifying the base measure of the nonparametric process, Scarpa and Dunson (2009) address the same problem by contaminating a point mass with a realization from a DP. As such, part of our method can be seen as a direct extension of the latter, where  $J \geq 1$  different atoms centered in locations learned from the training set are contaminated with a DP.

Let  $\Theta_m(t) = (f_m(t), \sigma_m^2(t))$  denote the vector comprising the smooth functional mean  $f_m : \mathcal{T} \rightarrow \mathbb{R}$  and the measurement noise  $\sigma_m^2 : \mathcal{T} \rightarrow \mathbb{R}^+$  for a generic curve  $m$  in the test set, evaluated at the instant  $t$ . Then the Brand model, introduced in Section 2.2 for multivariate data, can be modified as follows:

$$\begin{aligned} y_m(t) | \Theta_m(t) &= f_m(t) + \varepsilon_m(t); \quad \varepsilon_m(t) \sim N(0, \sigma_m^2(t)) \\ \Theta_m(t) | \tilde{p} &\sim \tilde{p}, \quad \tilde{p} = \sum_{j=1}^J \pi_j \delta_{\Theta_j} + \pi_0 \left[ \sum_{h=1}^{+\infty} \omega_h \delta_{\Theta_h^{nov}} \right], \\ (\pi_0, \pi_1, \dots, \pi_J) &\sim Dir(a_0, a_1, \dots, a_J), \quad \omega \sim SB(\gamma), \\ \Theta_j &\sim P_j^{Tr}, \quad \Theta_h^{nov} \sim H, \end{aligned} \tag{11}$$

where all the distributions  $P_j^{Tr}$  and the base measure  $H$  model the functional mean and the noise independently. We propose the following informative prior for  $\Theta_j = (f_j(t), \sigma_j^2(t))$ :

$$\begin{aligned} f_j(t) &\overset{ind.}{\sim} N(\bar{f}_j(t), \varphi_j), \\ \sigma_j^2(t) &\overset{ind.}{\sim} IG \left( 2 + \frac{(\bar{\sigma}_j^2(t))^2}{v_j}, \bar{\sigma}_j^2(t) \left( 1 + \frac{(\bar{\sigma}_j^2(t))^2}{v_j} \right) \right). \end{aligned} \tag{12}$$

We denote the estimates obtained from the training set of the mean and variance functions, as  $\bar{f}_j$  and  $\bar{\sigma}_j^2$ , respectively, for each observed class  $j$ , with  $j = 1, \dots, J$ . The hyper-parameters  $\varphi_j$  define the degree of confidence we a priori assume for the information extracted from the learning set, while the Inverse Gamma (*IG*) specification ensures that  $\mathbb{E}[\sigma_j^2(t)] = \bar{\sigma}_j^2(t)$  and  $Var[\sigma_j^2(t)] = v_j$ . It remains to define how we compute  $\bar{f}_j$

and  $\bar{\sigma}_j^2$ , that is, how the robust extraction of prior information is performed in this functional extension. Applying standard procedures in Functional Data Analysis (Ramsay, James, Silverman, 2005), we first smooth each training curve  $x_n(t)$  via a weighted sum of  $B$  basis functions

$$x_n(t) \approx \sum_{b=1}^B \rho_{nb} \phi_b(t) \quad n = 1, \dots, N$$

where  $\phi_b(t)$  is the  $b$ -th basis evaluated in  $t$  and  $\rho_{nb}$  its associated coefficient. Given the acyclic nature of the functional objects treated in Section 6.3, we will subsequently employ B-spline bases (Boor, 2001). Clearly, depending on the problem at hand, other basis functions may be considered. After such representation has been performed, we are left with  $J$  matrices of coefficients each of dimension  $n_j \times B$ . By treating them as multivariate entities, as done for example in Abraham et al., 2003, we resort to the very same procedures described in Section 2.1, and we set

$$\bar{f}_j(t) = \sum_{b=1}^B \hat{\rho}_{jb}^{MCD} \phi_b(t), \quad \bar{\sigma}_j^2(t) = \frac{1}{n_j - 1} \sum_{n: \mathbf{1}_n = j \cap \mathcal{I}_{MCD}^{(j)}} (x_n(t) - \bar{f}_j(t))^2$$

where  $\hat{\rho}_{jb}^{MCD}$  is the robust location estimate on the  $n_j \times B$  matrix of coefficients, and  $\mathcal{I}_{MCD}^{(j)}$  denotes the subset of untrimmed units resulting from the MCD/MRCD procedure in group  $j$ ,  $j = 1, \dots, J$ . On the other hand, more flexibility is needed to specify the base measure  $H$  for  $\Theta_h^{nov} = (f_h^{nov}(t), \sigma_h^2{}^{nov}(t))$ . Therefore, via the same smoothing procedure considered for the training curves, we build a hierarchical specification for the quantities involved in the novelty term:

$$\begin{aligned} f_h^{nov}(t) &= \sum_{b=1}^B \rho_{hb}^{nov} \phi_b(t), & \rho_{hb}^{nov} &\sim N(\psi_h, \tau_h^2), \\ \psi_h &\sim N(0, s^2), \\ \tau_h^2 &\sim IG(a_\tau, b_\tau), & \sigma_h^2{}^{nov}(t) &\sim IG(a_H, b_H). \end{aligned} \tag{13}$$

The first line of (13) can be rewritten as

$$f_h^{nov}(t) \sim N \left( \psi_h \sum_{b=1}^B \phi_b(t), \tau_h^2 \sum_{b=1}^B \phi_b^2(t) \right).$$

We call this model functional Brand: it provides a powerful extension for functional novelty detection. A successful application is reported in Section 6.3.

## 5 Posterior Inference

The posterior distribution  $p(\boldsymbol{\pi}, \boldsymbol{\omega}, \mathbf{a}, \boldsymbol{\beta}, \boldsymbol{\Theta}, \boldsymbol{\Theta}^{nov} | \mathbf{y})$  is analytically intractable, therefore we rely upon MCMC techniques to carry out posterior inference. An easy sampling

scheme can be constructed mimicking the blocked Gibbs sampler of Ishwaran and James (2001), where the infinite series in (3) is truncated at a pre-specified level  $L < \infty$ . However, this approach leads to a non-negligible truncation error if  $L$  is too small, and to computational inefficiencies if  $L$  is set too high. Instead, we propose a modification of the  $\xi$ -sequence of the Independent Slice-efficient sampler (Kalli et al., 2011), another well known conditional algorithm to sample from the exact posterior. To adapt the algorithm to our framework, we start from the following alternative reparameterization of the model in (3)-(4):

$$\begin{aligned} \mathbf{y}_m | \tilde{\Theta}, \zeta_m &\sim N\left(\tilde{\Theta}_{\zeta_m}\right) & \zeta_m | \tilde{\boldsymbol{\pi}} &\sim \sum_{k=1}^{\infty} \tilde{\pi}_k \delta_k(\cdot) \\ \tilde{\pi}_k &= \pi_k^{\mathbb{1}_{\{0 < k \leq J\}}} \cdot (\pi_0 \cdot \omega_{k-J})^{\mathbb{1}_{\{k \geq J+1\}}} & \text{for } k \geq 1 \\ \tilde{\Theta}_k &= \Theta_k^{\mathbb{1}_{\{0 < k \leq J\}}} \cdot (\Theta_{k-J}^{nov})^{\mathbb{1}_{\{k \geq J+1\}}} & \text{for } k \geq 1 \end{aligned} \quad (14)$$

where  $\tilde{\Theta}$  is obtained by concatenating  $\Theta$  and  $\Theta^{nov}$ ,  $\delta_k$  is the usual Dirac delta function, the weights  $\boldsymbol{\pi}$  and  $\boldsymbol{\omega}$  are defined as in Equation (7), and  $\zeta_m$  is a membership label which maps each observation to its corresponding atom  $\tilde{\Theta}_{\zeta_m}$ .

Trivially, there is a one-to-one correspondence between the membership vectors  $(\alpha_m, \beta_m)$  of model (6) and  $\zeta_m$

$$\zeta_m = l \iff \alpha_m = l \cdot \mathbb{1}_{\{\zeta_m \leq J\}}, \quad \beta_m = (l - J) \cdot \mathbb{1}_{\{\zeta_m > J\}}. \quad (15)$$

However, we prefer the form of model (6) thanks to the direct interpretation of the membership latent variables  $\boldsymbol{\alpha}$  and  $\boldsymbol{\beta}$ , which associate each observation to the known or novel classes, respectively. We introduce two sequences of additional auxiliary parameters: a stochastic sequence  $\mathbf{u} = \{u_m\}_{m=1}^M$  of uniform random variables and a deterministic sequence  $\boldsymbol{\xi} = \{\xi_l\}_{l \geq 1}$ . The introduction of these two latent variables allows for a stochastic truncation at each iteration of the sampler. The stochastic threshold, called  $L$ , is given as  $L = \max L_m$  and  $L_m$  is the largest integer such that  $\xi_{L_m} > u_m$ . This threshold establishes a finite number of mixture components needed at each MCMC iteration, making computations feasible. Then, we can rewrite model (6) as

$$\mathcal{L}(\mathbf{y}, \boldsymbol{\zeta}, \mathbf{u} | \tilde{\boldsymbol{\pi}}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \prod_{m=1}^M \left[ \frac{\tilde{\pi}_{\zeta_m}}{\xi_{\zeta_m}} \mathbb{1}_{\{u_m < \xi_{\zeta_m}\}} \phi\left(\mathbf{y}_m | \tilde{\Theta}_{\zeta_m}\right) \right]. \quad (16)$$

In the definition of a dedicated deterministic sequence  $\boldsymbol{\xi}$ , it is crucial to take into account the difference between the manifest and the novel components. Usually, a very common choice is  $\xi_l = (1 - \kappa)\kappa^{l-1}$ , with  $\kappa \in (0, 1)$ . This option allows to compute each  $L_m$  analytically, being the smallest integer such that

$$L_m < 1 + \frac{\log(u_m) - \log(1 - \kappa)}{\log(\kappa)}.$$

However, the default choice of a geometrically decreasing  $\boldsymbol{\xi}$ -sequence is inappropriate in this context, since we are dealing with a mixture where not all the components are




Figure 2: Example of deterministic sequence defined according to (17), with  $\kappa = 0.25$ . The blue rectangle highlights the weights relative to the known components.

conceptually equivalent. The default  $\xi$ -sequence tends to favor components that come first in the mixture specification (in our case, the known ones). To overcome this issue, we propose the following intuitive modification. Given a value for  $\kappa \in (0, 1)$ , we equally divide the  $(1 - \kappa)\%$  of the mass into the first  $J + 1$  elements of the sequence. We then induce a geometric decay in the remaining ones to split the residual fraction  $\kappa$ . We force the element in position  $J + 1$  to have the same mass given to the manifest components, to avoid an under-representation of the novelty part. To do so, we define

$$\xi_l = \begin{cases} \frac{1-\kappa}{J+1} & \text{if } l \leq J \\ \frac{1-\kappa}{J+1} \left( \frac{(J+1)\kappa}{JK+1} \right)^{l-J-1} & \text{if } l > J + 1 \end{cases} \quad (17)$$

It is easy to prove that  $\sum_{l=1}^{+\infty} \xi_l = 1$ . According to (17), the first  $J + 1$  elements of the sequence have masses equal to  $(1 - \kappa)/(J + 1)$ . The truncation threshold  $L^*$  changes accordingly, becoming the largest integer such that

$$L^* < J + 1 + \frac{\log(\min(\mathbf{u})) - \log\left(\frac{1-\kappa}{J+1}\right)}{\log\left(\frac{(J+1)\kappa}{JK+1}\right)}. \quad (18)$$

Inequality (18) states that the truncation threshold  $L^*$  can be only greater or equal to  $J + 1$ , ensuring that the MCMC always takes into consideration the creation of at least one cluster in the novel distribution. A representation of the modified  $\xi$ -sequence is

depicted in Figure 2. We report the pseudo-code for the devised Gibbs sampler in the Appendix. The algorithm for the functional extension is not included for conciseness. However, its structure closely follows the one outlined for the multivariate case.

Once the MCMC sample is collected, we first compute the a posteriori probability of being a novelty for every test unit  $m$ ,  $PPN_m = \mathbb{P}[\mathbf{y}_m \sim f^{nov} | \mathbf{Y}]$ , that is estimated according to the ergodic mean:

$$PP\hat{N}_m = \frac{\sum_{i=1}^I \mathbb{1}_{\{a_m^{(i)}=0\}}}{I}, \quad (19)$$

where  $\alpha_m^{(i)}$  is the value assumed by the parameter  $\alpha_m$  at the  $i$ -th iteration of the MCMC chain and  $I$  is the total number of iterations. We remark that the inference on  $\alpha$  can be conducted directly, since the mixture between the  $J$  observed components and  $f^{nov}$  is not subjected to label switching. In contrast, we need to take this problem into account when dealing with  $\beta$ . To perform valid inference, one possibility is to rely on the posterior probability coclustering matrix (PPCM) as indicated in Section 2.3. Each entry of this matrix  $p_{m,m'} = \mathbb{P}[\mathbf{y}_m \text{ and } \mathbf{y}_{m'} \text{ belong to the same novelty class}]$  is estimated as

$$\hat{p}_{m,m'} = \frac{\sum_{i=1}^I \mathbb{1}_{\{\beta_m^{(i)} = \beta_{m'}^{(i)}\}}}{I}. \quad (20)$$

Once we obtain the PPCM, we employ it to estimate the best partition (BP) in the novelty subset. Indeed, one can recover the BP by minimizing a loss function defined over the space of partitions, which can be computed starting from the PPCM. A famous loss function was proposed by Binder (1978), and investigated in a BNP setting by Lau and Green (2007). However, the so-called Binder loss presents peculiar asymmetries, preferring to split clusters over merging. These asymmetries could result in a number of estimated clusters higher than needed. Therefore, we adopt the Variation of Information (VI - Meilă, 2007) as loss criterion. The associated loss function, recently proposed by Wade and Ghahramani (2018), is known to provide less fragmented results.

Finally, once the BP for the novelty component has been estimated, we can rely on a heuristic based on the cluster sizes to discriminate anomalies from actual new classes. Let us suppose that the BP consists of  $S$  novel clusters. Denote the number of instances assigned to cluster  $s \in \{1, \dots, S\}$  with  $m_s^{nov}$ . A cluster  $s$  is considered to be an agglomerate of outlying points if its cardinality  $m_s^{nov}$  is sufficiently small in comparison to the entire novelty sample size, otherwise it is regarded as a proper novel group.

## 6 Applications

### 6.1 Simulation Study

In this section, we present a simulation study aimed at highlighting the capabilities of the new semiparametric Bayesian model in performing novelty detection and we compare it with existing methodologies. We consider different scenarios varying the sample sizes

of the hidden classes and the adulteration proportions in the training set. At the same time, we evaluate the importance of the robust information extraction phase and how it affects the learning procedure.

### 6.1.1 Experimental setup

We consider a training set formed by  $J = 3$  observed classes, each distributed according to a bivariate Normal density  $\mathcal{N}_2(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$ ,  $j = 1, 2, 3$ , with the following parameters:

$$\begin{aligned} \boldsymbol{\mu}_1 &= (-5, 5)', & \boldsymbol{\mu}_2 &= (-4, -4)', & \boldsymbol{\mu}_3 &= (4, 4)' \\ \boldsymbol{\Sigma}_1 &= \begin{bmatrix} 1 & 0.9 \\ 0.9 & 1 \end{bmatrix} & \boldsymbol{\Sigma}_2 &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} & \boldsymbol{\Sigma}_3 &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \end{aligned}$$

The class sample sizes are, respectively, equal to  $n_1 = 300$ ,  $n_2 = 300$  and  $n_3 = 400$ . The same groups are also present in the test set, together with four previously unobserved classes. We generate the new classes via bivariate Normal densities with parameters:

$$\begin{aligned} \boldsymbol{\mu}_4 &= (0, 0)', & \boldsymbol{\mu}_5 &= (5, -10)', \\ \boldsymbol{\mu}_6 &= (5, -10)', & \boldsymbol{\mu}_7 &= (-10, -10)', \\ \boldsymbol{\Sigma}_4 &= \begin{bmatrix} 1 & -0.75 \\ -0.75 & 1 \end{bmatrix}, & \boldsymbol{\Sigma}_5 &= \begin{bmatrix} 1 & 0.9 \\ 0.9 & 1 \end{bmatrix}, \\ \boldsymbol{\Sigma}_6 &= \begin{bmatrix} 1 & -0.9 \\ -0.9 & 1 \end{bmatrix}, & \boldsymbol{\Sigma}_7 &= \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}. \end{aligned}$$

The test set encompasses a total of 7 components: 3 observed and 4 novelties. Starting from the above-described data generating process, we consider four different scenarios varying:

- Data contamination level
  - No contamination in the training set (**Label noise = False**)
  - 12% label noise between classes 2 and 3 (**Label noise = True**)
- Test set sample size
  - Novelty subset size equal to slightly more than 30% of the test set (**Novelty size = Not small**)

$$\begin{aligned} m_1 &= 200, & m_2 &= 200, & m_3 &= 250, & m_4 &= 90, \\ m_5 &= 100, & m_6 &= 100, & m_7 &= 10 \end{aligned}$$

- Novelty subset size equal to 15% of the test set (**Novelty size = Small**)

$$\begin{aligned} m_1 &= 350, & m_2 &= 250, & m_3 &= 250, & m_4 &= 49, \\ m_5 &= 50, & m_6 &= 50, & m_7 &= 1. \end{aligned}$$






Figure 3: Simulated data for the `Label noise = True`, `Novelty size = Not small` scenario. Classes 4, ..., 7 are not observed in the learning set.

Figure 3 exemplifies the experiment structure displaying a realization from the `Label noise = True`, `Novelty size = Not small` scenario. As it is evident from the plots, the label noise is strategically included to cause a more difficult identification of the fourth class, should the parameters of the second and third classes be non-robustly learned. Further, notice that the last group presents limited sample size and variability: it could easily be regarded as pointwise contamination (i.e., an anomaly) rather than an actual new component. Nonetheless, following the reasoning outlined in the introduction, we are interested in evaluating the ability of the nonparametric density to capture and discriminate these types of peculiar patterns as well. For each combination of contamination level and test set sample size, we simulate  $B = 100$  datasets. Results are reported in the following subsection.

### 6.1.2 Simulation results

We compare the performance of the Brand model with different hyper-parameters specifications:

- the information from the training set is either non-robustly ( $\eta_{MCD} = 1$ ) or robustly ( $\eta_{MCD} = 0.75$ ) extracted,
- the precision parameter associated with the training prior belief is either very high ( $\lambda_{Tr} = 1,000$ ) or moderately low ( $\lambda_{Tr} = 10$ ).

In addition, two model-based adaptive classifiers are considered in the comparison, namely the inductive RAEDDA model (Cappozzo et al., 2020) with labeled and unlabeled trimming levels respectively equal to 0.12 and 0.05, and the inductive AMDA model (Bouveyron, 2014). For each replication of the simulated experiment, a set of four metrics is recorded from the test set:

- *Novelty predictive value (Precision)*: the proportion of units marked as novelties by a given method truly belonging to classes 4, ..., 7,
- *Accuracy on the observed classes subset*: the classification accuracy of a given method within the subset of groups already observed in the training set,

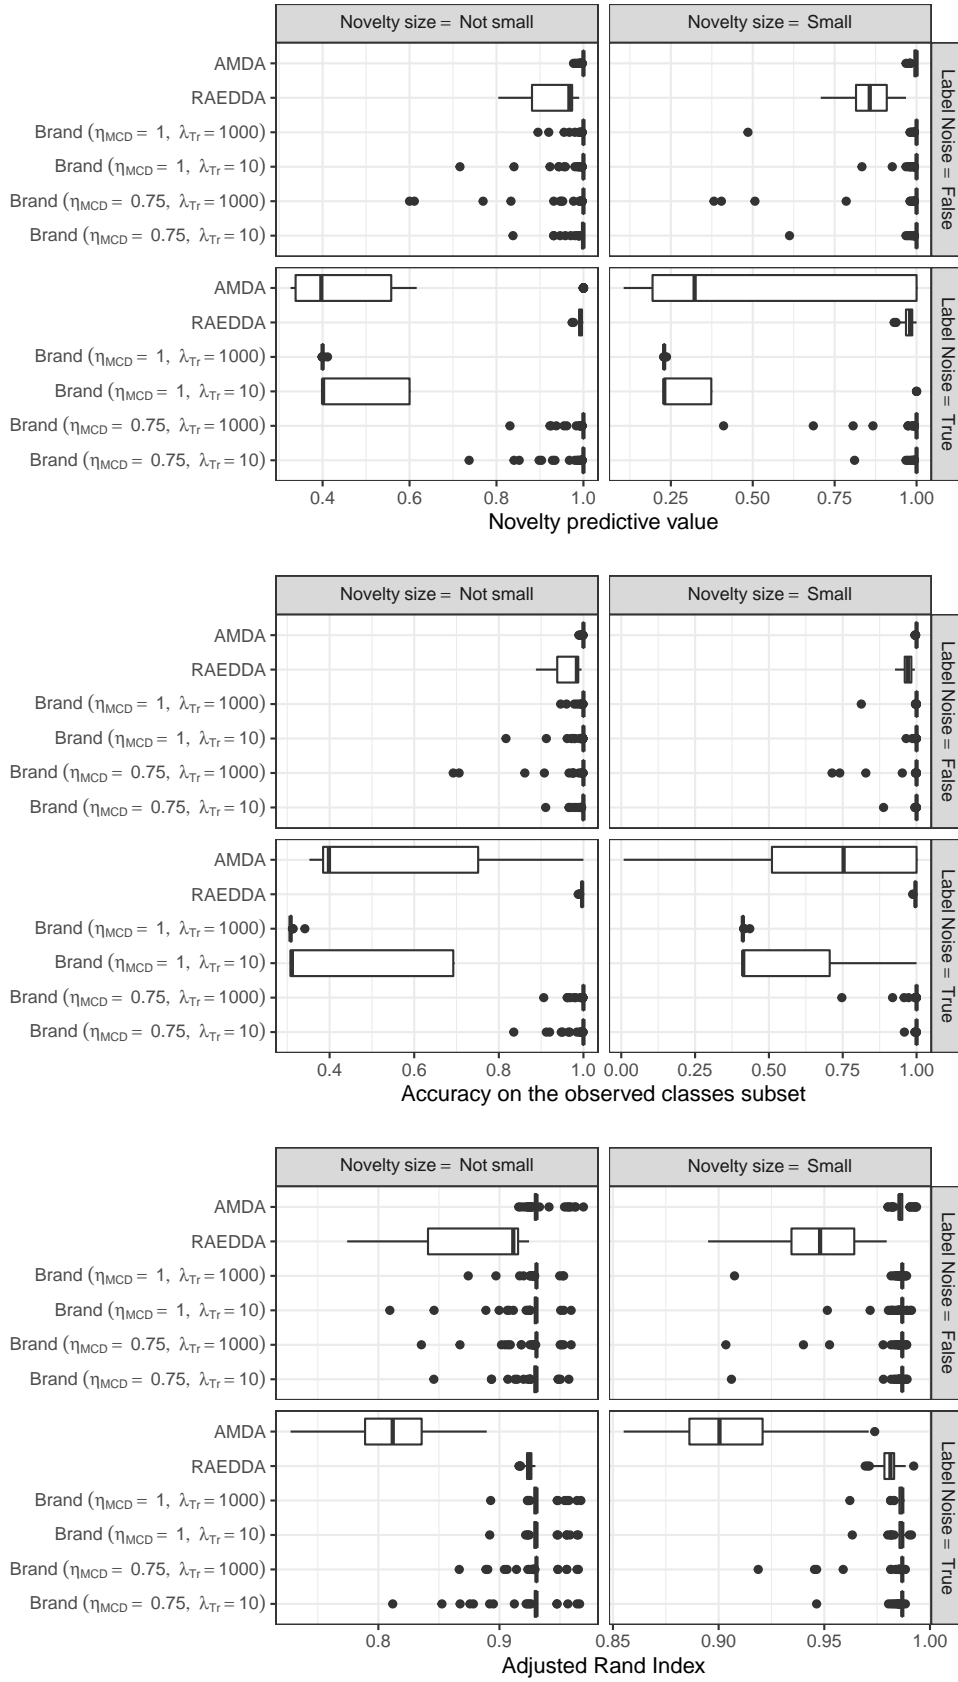


Figure 4: Box plots for (from top to bottom) novelty predictive value, accuracy on the known classes and ARI metrics for  $B = 100$  repetitions of the simulated experiment, varying data contamination level and test set sample size.

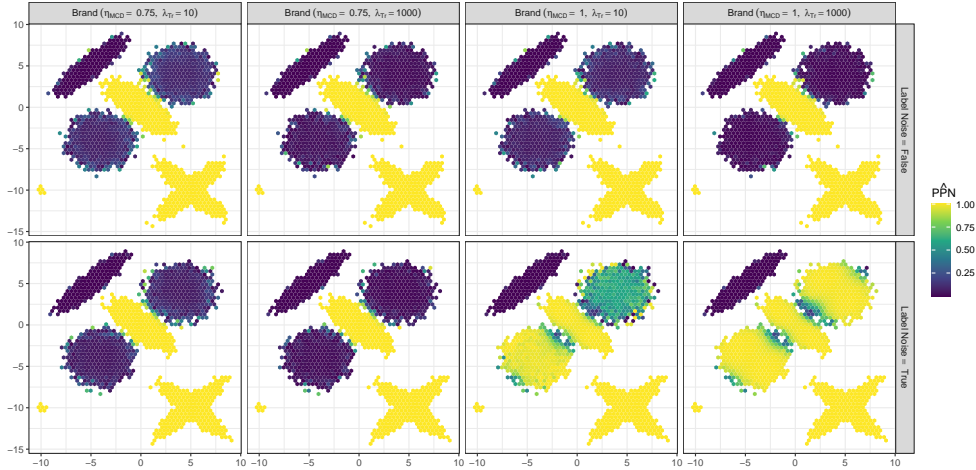


Figure 5: Hex plots of the average estimated posterior probability of being a novelty, according to formula (19), for  $B = 100$  repetitions of the simulated experiment, varying data contamination level and Brand hyper-parameters, **Not small** novelty subset size. The brighter the color the higher the probability of belonging to  $f^{nov}$ .

- *Adjusted Rand Index (ARI, Rand, 1971)*: measuring the similarity between the partition returned by a given method and the underlying true structure,
- *PPN*: a posteriori probability of being a novelty, computed according to Equation (19) (Brand only).

We run 40,000 MCMC iterations and discard the first 20,000 as a burn-in phase. Apart from the hyper-parameters for the training components, fairly uninformative priors are employed in the base measure  $H$ , with  $\mathbf{m}_0 = (0, 0)'$ ,  $\lambda_0 = 0.01$ ,  $\nu_0 = 10$  and  $\mathbf{S}_0 = 10\mathbf{I}_2$ . Lastly, a Gamma DP concentration parameter is considered with prior rate and scale hyper-parameters both equal to 1.

Figure 4 reports the results for  $B = 100$  repetitions of the experiment under the different simulated scenarios. A Table containing the values on which this plot is built is deferred to the Supplementary Material. The *Novelty predictive value* metric highlights the capability of the model to correctly recover and identify the previously unseen patterns. As expected, in the adulteration-free scenarios, all methodologies succeed well enough in separating known and hidden components. The worst performance is exhibited by the RAEDDA model for which, due to the fixed trimming level, a small part of the group-wise most extreme (but still genuine) observations is discarded, thus slightly overestimating the novelties percentage (the same happens for the ARI metric). Different results are displayed in scenarios wherein the label noise complicates the learning process. Robust procedures efficiently cope with the adulteration present in the training set, while the AMDA and the Brand methods when  $\eta_{MCD} = 1$  tend to largely overestimate the novelty component. Particularly, the harmful effect caused by

the mislabeled units is exacerbated in the Brand model that sets high confidence in the priors ( $\lambda_{Tr} = 1,000$ ), while a partial mitigation, albeit feeble, emerges when  $\lambda_{Tr}$  is set equal to 10. This consequence is even more apparent in the hex plots of Figure 5, where we see that the latter model tries to modify its prior belief to accommodate the (outlier-free) test units, while the former, forced to stick close to its prior distribution by the high value of  $\lambda_{Tr}$ , incorporates the second and third class in the novelty term. The final output, as displayed in the *Accuracy on the observed classes subset* boxplots, has an overall high misclassification error when it comes to identifying the test units belonging to the previously observed classes. Differently, setting robust informative priors prevents this undesirable behavior, as it is shown by both the high level of accuracy and the associated low posterior probability of being a novelty in the feature space wherein the observed groups lie. On the other hand, the true partition recovery, assessed by the Adjusted Rand Index, does not seem to be influenced by the label noise, with our proposal always outperforming the competing methodologies regardless of which hyperparameters are selected. As previously mentioned, for *Brand*( $\eta_{MCD} = 1, \lambda_{Tr} = 10$ ) and *Brand*( $\eta_{MCD} = 1, \lambda_{Tr} = 1,000$ ) cases the second and third classes are assimilated into the nonparametric component in the **Label Noise = True** scenario. This is due to the fact that the mislabeled units prevent Brand from correctly learning the true structures of groups two and three in Stage I. As a consequence, no correspondence between these improperly estimated classes in the training is found in the test set, so much so that the DP prior creates them anew within the novelty term. Clearly, this is a sub-optimal behavior as the separation of what is known from what is novel is completely lost, yet it may raise suspicion on dealing with a contaminated learning set, suggesting the need of a robust prior information extraction.

Additional simulated experiments, involving a high-dimensional scenario and novelty detection problem under model misspecification are included in the Supplementary Material.

## 6.2 X-ray images of wheat kernels

Sophisticated and advanced techniques like X-rays, scanning microscopy and laser technology are increasingly employed for the automatic collection and processing of images. Within the domain of computer vision studies, novelty detection is generally portrayed as a one-class classification problem. There, the aim is to separate the known patterns from the absent, poorly sampled or not well defined remainder (Khan and Madden, 2014). Thus, there is strong interest in developing methodologies that not only distinguish the already observed quantities from the new entities, but that also identify specific structures within the novelty component. The present case study involves the detection of a novel grain type by means of seven geometric parameters, recorded postprocessing X-ray photographs of kernels (Charytanowicz et al., 2010). In more detail, for the 210 samples belonging to the three different wheat varieties, high quality visualization of the internal kernel structure is detected using a soft X-ray technique and, subsequently, the image files are post-processed via a dedicated computer software package (Strumillo et al., 1999). The obtained dataset is publicly available in the University of California,

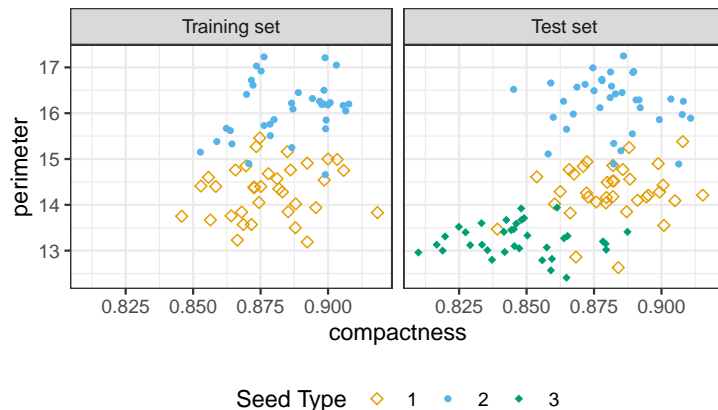


Figure 6: Learning scenario (only `compactness` and `perimeter` variables displayed) for novelty detection of an unobserved wheat variety, seed dataset.

Irvine Machine Learning data repository. This experiment involves the random selection of 70 training units from the first two cultivars, and a test set of 105 samples, including 35 grains from the third variety. The resulting learning scenario is displayed in Figure 6. The aim of the analysis is to employ Brand to detect the third unobserved variety, whilst performing classification of the known grain types with high accuracy. Firstly, the MCD estimator with hyper-parameter  $h_{MCD} = 0.95$  is adopted for robustly learning the training structure of the two observed wheat varieties. In the second stage, our model is fitted to the test set, discarding 20,000 iterations for the burn-in phase, and subsequently retaining 10,000 MCMC samples. As usual, fairly uninformative priors are employed in the base measure  $H$ , with  $\mathbf{m}_0 = \mathbf{0}$ ,  $\lambda_0 = 0.01$ ,  $\nu_0 = 10$  and  $\mathbf{S}_0 = \mathbf{I}_7$ , where  $\mathbf{0}$  denotes the 7-dimensional zero vector. For the training components, mean and covariance matrices of the Normal-inverse-Wishart priors are directly determined by the MCD output of the first stage, while  $\nu^{Tr}$  and  $\lambda^{Tr}$  are specified to be respectively equal to 250 and 1,000. The latter value indicates that after having robustly extracted information for the two observed classes, high trust is placed in the prior distributions of the known components. Model results are reported in Figure 7, where the posterior probability of being a novelty  $PPN_m = \mathbb{P}[\mathbf{y}_m \sim f^{nov} | \mathbf{Y}]$ ,  $m = 1, \dots, M$ , displayed in the plots below the main diagonal, are estimated according to the ergodic mean in (19). The a posteriori classification, computed via majority vote, is depicted in the plots above the main diagonal, where the water-green solid diamonds denote observations belonging to the novel class. The confusion matrix associated with the estimated group assignments is reported in Table 1, where the third group variety is effectively captured by the flexible process modeling the novel component.

All in all, the promising results obtained with this multivariate dataset may foster the employment of our methodology in automatic image classification procedures that supersede the one-class classification paradigm, allowing for a much more flexible anomaly and novelty detector in computer vision applications.

Table 1: Confusion matrix for the semiparametric Bayesian classifier on the test set, seeds dataset. The label “New” indicates observations that are estimated to have arisen from the novelty component.

Classification	Truth		
	1	2	3
1	30	0	7
2	2	35	0
New	3	0	28

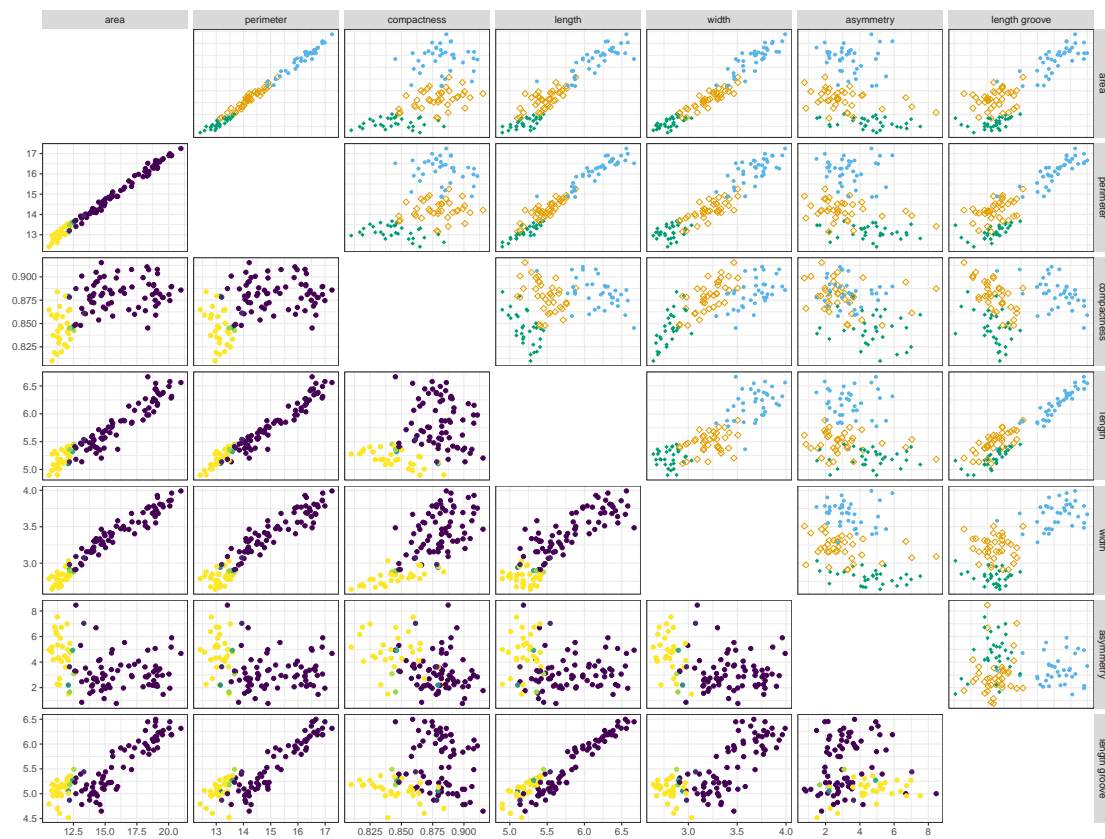


Figure 7: Test set for the considered experimental scenario, seeds dataset. Plots below the main diagonal represent the estimated posterior probability of being a novelty. The brighter the color the higher the probability of belonging to  $f^{nov}$ . Plots above the main diagonal display the associated group assignments, where the water-green solid diamonds denote observations classified as novelties.

### 6.3 Functional novelty detection of meat variety

In recent years, machine learning methodologies have experienced an ever-growing interest in countless fields, including food authentication research (Singh and Domijan, 2019). An authenticity study aims to characterize unknown food samples, correctly identifying their type and/or provenance. Clearly, no observation is to be trusted in a context wherein the final purpose is to detect potentially adulterated units, in which, for example, an entire subsample may belong to a previously unseen pattern. Motivated by a dataset of Near Infrared Spectra (NIR) of meat varieties, we employ the functional model introduced in Section 4 to perform classification and novelty detection when having a hidden class and four manually adulterated units in the test set. The considered data report the electromagnetic spectrum for a total of 231 homogenized meat samples, recorded from 400 – 2498  $nm$  at intervals of 2  $nm$  (McElhinney et al., 1999). The units belong to five different meat types, with 32 beef, 55 chicken, 34 lamb, 55 pork, and 55 turkey records. The amount of light absorbed at a given wavelength is recorded for each meat sample:  $A = \log_{10}(1/R)$  where  $R$  is the reflectance value. The visible part of the electromagnetic spectrum (400 – 780  $nm$ ) accounts for color differences in the meat types, while their chemical composition is recorded further along the spectrum. NIR data can be interpreted as a discrete evaluation of a continuous function in a bounded domain. Therefore, the procedure described in Section 4 is a sensible methodological tool for modeling this type of data objects (Barati et al., 2013). We randomly partition the recorded units into labeled and unlabeled sets. The former includes 28 chicken, 17 lamb, 28 pork, and 28 turkey samples. The latter contains the same proportion of these four meat types with an additional 32 beef units. The last class is not observed in the test set and needs to be discovered. Also, four validation units are manually adulterated and added to the test set as follows:

- a shifted version of a pork sample, achieved by removing the first 15 data points and appending the last 15 group-mean absorbance values at the end of the spectrum;
- a noisy version of a pork sample, generated by adding Gaussian white noise to the original spectrum;
- a modified version of a turkey sample, obtained by abnormally increasing the absorbance value in a single specific wavelength to simulate a spike;
- a pork sample with an added slope, produced by multiplying the original spectrum by a positive constant.

These modifications mimic the ones considered in the “Chimiométrie 2005” chemometric contest, where participants were tasked to perform discrimination and outlier detection of mid-infrared spectra of four different starches types (Fernández Pierna and Dardenne, 2007). In our context, both the beef subpopulation and the adulterated units are previously unseen patterns that shall be captured by the novelty component.

Firstly, we extract robust prior information from the learning set. Given the spectra non-cyclical nature, we approximate each training unit via a linear combination of

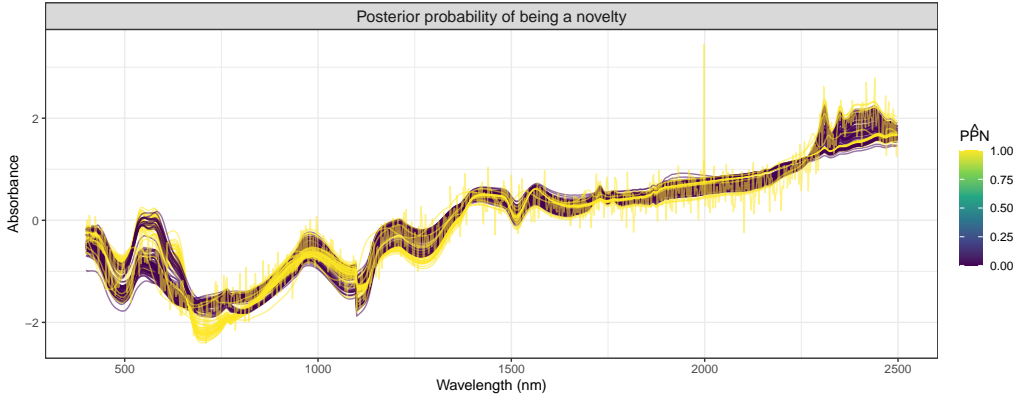


Figure 8: Estimated posterior probability of being a novelty, according to formula (19), the brighter the color the higher the probability of belonging to  $f^{nov}$ .

$B = 100$  B-spline bases, and their associated coefficients are retrieved. Given the high-dimensional nature of the smoothing process, the MRCD is employed to obtain robust group-wise estimates for the splines coefficients. These quantities, which are linearly combined with the B-spline bases, account for the training atoms  $\Theta_j$ ,  $j = 1, \dots, 4$  specified in Equation (11). We adopt a value of  $\eta_{MCD} = 0.75$  in the first stage, providing functional atoms robust against contamination that may arise in the training set. In this experiment, an inductive approach is considered, for which the training estimates will be kept fixed throughout the subsequent Bayesian learning phase. We further set  $a_\tau = 3$ ,  $b_\tau = 1$ ,  $s^2 = 1$ ,  $a_H = 5$ , and  $b_H = 1$ , inducing low variability on the noise parameters as much as not to compromise the hierarchical structure between known and novelty components. A more detailed discussion on the hyperparameters choice is deferred to the Supplementary Material, where we evaluate alternative effects for different prior settings within a controlled experiment. Once  $\Theta_j$ ,  $j = 1, \dots, 4$  are retained, the Bayesian model of Section 4 is applied to the test units running a total of 20,000 iterations and discarding the first 10,000 as warm-up. Figure 8 summarizes the results of the fitted model. Each spectrum is colored according to its a posteriori probability of being a novelty, computed as in (19). The resulting confusion matrix is reported in Table 2, where it is apparent that the previously unseen class, as well as the adulterated units (labeled as “Outliers” in the table), are successfully captured by the novelty component. The obtained classification accuracy is in agreement with the ones produced by state-of-the-art classifiers in a fully-supervised scenario (see, for example, Murphy et al., 2010; Gutiérrez et al., 2014). That is, our proposal is capable of detecting previously unseen classes and outlying units, whilst maintaining competitive predictive power.

Looking at the classification performance, we observe that Brand can correctly recover the underlying data partition, except for the turkey subgroup. Specifically, the model struggles to separate the turkey units from the chicken ones. Figure 9 provides an explanation for this issue. The left panel shows the robust functional means extracted from the training set. The right panel shows the functional test objects containing the



Classification	Truth					
	Beef	Chicken	Lamb	Pork	Turkey	Outliers
Novelty	32	0	0	0	2	4
Chicken	0	21	0	1	12	0
Lamb	0	0	17	0	0	0
Pork	0	4	0	20	3	0
Turkey	0	2	0	3	9	0

Table 2: Confusion matrix for the semiparametric Bayesian classifier on the test set, meat dataset. The label “Novelty” indicates observations that are estimated to have arisen from the  $f^{nov}$ .

two types of poultry. The overlapping is evident in both cases and it is the main reason why Brand merges the two different sets.

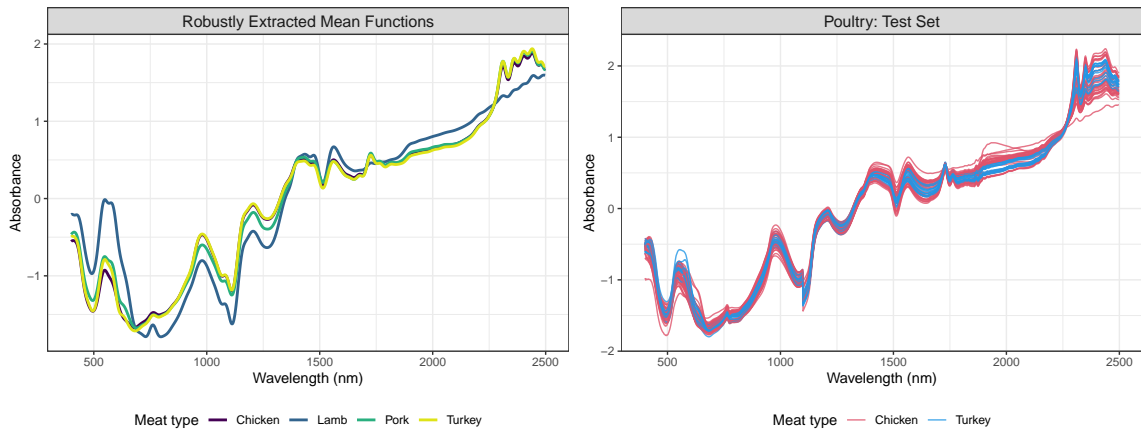


Figure 9: The left panel shows the robust functional means extracted from the training set. The right panel shows the test functional objects containing a type of poultry.

Focusing on the novelty component, the model entirely captures the beef hidden class and the adulterated units, yet two turkey samples are also incorrectly assigned. The obtained classification for the curves identified to be novelties, resulting by VI minimization, is displayed in the left panel of Figure 10, where two distinct clusters are detected. Interestingly, Brand separates the 32 beef samples (blue dashed lines) from the two turkeys (solid red lines) and classifies three of the four manually adulterated units to the outlying cluster. In contrast, the remaining one is assigned to the beef class, because of its peculiar shape, as it is shown in Figure 13 of the Supplementary Material.

Finally, we investigate why two turkey units are incorrectly assigned to the novel component. A closer look at the turkey sub-population, displayed in the right panel

of Figure 10, shows how these two samples exhibit a somehow extreme pattern within their group and can, therefore, be legitimately flagged as outlying or anomalous turkeys. We report two additional figures in the Supplementary Material. The first provides a

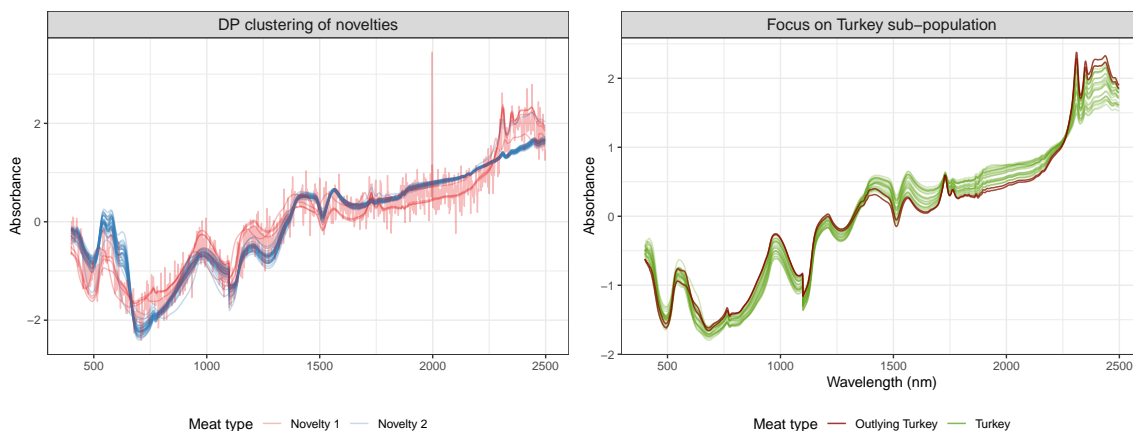


Figure 10: Left panel: best partition of the novelty component recovered by minimizing the Variation of Information loss function. The dashed blue curves are beef samples, while the solid red ones are the manually adulterated units and the two turkeys incorrectly assigned to the novel component. Right panel: true turkey sub-population in the test set, the units incorrectly assigned to the novel component are displayed with solid dark red lines.

visual summary of the estimated grouping; the second shows how the turkey test units are partitioned into different clusters.

In this section, we have shown the effectiveness of our methodology in correctly identifying a hidden group in a functional setting, while jointly achieving good classification accuracy and detection of outlying curves. The successful application of the model seems particularly desirable in fields like food authenticity, where generally there is no a priori available information on how many modifications and/or adulteration mechanisms may be present in the samples.

## 7 Conclusion and discussion

We have introduced a two-stage methodology for robust Bayesian semiparametric novelty detection. In the first stage, we robustly extract the observed group structure from the training set. In the second stage, we incorporate such prior knowledge in a contaminated mixture, wherein we have employed a nonparametric component to describe the novelty term. The latter could either correspond to anomalies or actual new groups. This distinction is made possible by retrieving the best partition within the novel subset. We have investigated the properties of the random measure underlying the model and its connections with existing methods. Subsequently, the general multivariate methodology

has been extended to handle functional data objects, resulting in a novelty detector for functional data analysis. A dedicated slice-efficient sampler, taking into account the difference between unseen and seen components, has been devised for posterior inference. An extensive simulation study and applications on multivariate and functional data have validated the effectiveness of our proposal, fostering its employment in diverse areas from image analysis to chemometrics. Brand can represent the starting point for many different research avenues. Future research directions aim at providing a Bayesian interpretation of the robust MCD estimator to propose a unified, fully Bayesian model. More versatile specifications can be adopted for the known components, weakening the Gaussianity assumption. These extensions can be obtained by adopting more flexible distributions while keeping the mean and variance of the resulting densities constrained to the findings in the training set, for example, via centered stick-breaking mixtures (Yang et al., 2010).

Similarly, functional Brand can be improved by adopting a more general prior specification via Gaussian Processes (Rasmussen and Williams, 2005). Lastly, it is of paramount interest to develop scalable algorithms, as Variational Bayes (Blei et al., 2017) and Expectation-Maximization (Dempster et al., 1977), for inference on massive datasets. Such solutions will offer both increased speed and lower computational cost, which are crucial for assuring the applicability of our proposal in the big data era.

## 8 Supporting Information

The Supplementary Material referenced throughout the article is available with this paper at the Statistics and Computing website. As supporting information, we report the proofs of the theoretical results showed in Section 3. Moreover, to complement the results presented in Section 6, we showcase the performance obtained by applying Brand to various challenging simulated data, varying the distributional assumptions and dimensionality. We also discuss an application to a higher dimensional real dataset, the popular benchmark `Wine` dataset from the UCI dataset repository, considering all its 13 features. Lastly, with the help of a controlled experiment, we guide the reader through the choice of hyperparameters and, more broadly, the whole usage of the model in the functional case. Software routines, including the implementation for both methods, the simulation study, and real data analyses of Section 6 are openly available at [github.com/AndreaCappozzo/brand-public\\_repo](https://github.com/AndreaCappozzo/brand-public_repo).

## Acknowledgement

The authors want to thank the Editor and the anonymous Reviewers for their suggestions and comments, which significantly improved the scientific value of the manuscript. During the development of this article, F. Denti was funded as a postdoctoral scholar by the NIH grant R01MH115697. Previously, he was also supported as a Ph.D. student by University of Milano - Bicocca, Milan, Italy and Università della Svizzera italiana, Lugano, Switzerland. Andrea Cappozzo and Francesca Greselin's work was supported

## Appendix

### Gibbs sampling algorithm for model (3)-(4)

---

**Algorithm 1:** Efficient Slice Sampler for the BNP-Novelty detection model

---

**Input:** Initial values for the MCMC, robust estimates from  $\mathbf{X}$ .

**Output:** Posterior MCMC sample for the parameters of interest.

**for**  $i = 1, \dots, I$  **do**

1. Sample every  $u_m$  from a uniform distribution  $\mathcal{U}(0, \xi_{\zeta_m})$ .
2. Compute the stochastic truncation term  $L^*$  according to (18).
3. Let  $m_j = \sum_{m=1}^M \mathbb{1}_{\{\alpha_m=j\}}$ , with  $j = 0, \dots, J$ . Sample  $\pi$  from a conjugate Dirichlet distribution:

$$\pi \sim \text{Dir}(a_0 + m_0, a_1 + m_1, \dots, a_J + m_J).$$

4. Sample the SB variables after integrating out  $\mathbf{u}$ :

$$v_k | \dots \sim \text{Beta}\left(1 + \sum_{m=1}^M \mathbb{1}_{\{\alpha_m=0 \cap \beta_m=k\}}, \gamma + \sum_{m=1}^M \mathbb{1}_{\{\alpha_m=0 \cap \beta_m>k\}}\right).$$

5. Compute the SB weights according to (5)
6. Compute the one-line probability weights  $\tilde{\pi}$  according to (14).
7. Sample the atoms for the observed classes  $\Theta_{j,0}$  exploiting conjugacy between the likelihood and the prior for  $j = 1, \dots, J$ .
8. Sample the atoms for the novel classes  $\Theta_{0,l}^{nov}$  exploiting conjugacy between the likelihood and the prior for  $l = 1, \dots, L^*$ .
9. Obtain  $\tilde{\Theta}$  concatenating the updated values of  $\Theta$  and  $\Theta^{nov}$ .
10. Sample each  $\zeta_m$  from the following joint discrete distribution:

$$\mathbb{P}(\zeta_m = l) \propto \frac{\tilde{\pi}_l}{\xi_l} \mathbb{1}_{\{u_m < \xi_l\}} \phi(\mathbf{y}_m | \tilde{\Theta}_l), \quad l = 1, \dots, L^*,$$

$$\mathbb{P}(\text{otherwise}) \propto 0.$$

11. Recover the values for the membership vectors  $\alpha$  and  $\beta$  using (15).  
Divide the elements in  $\tilde{\Theta}$  into  $\Theta$  and  $\Theta^{nov}$ .

**end**

---

## References

- Abraham, C., P. A. Cornillon, E. Matzner-Løber, and N. Molinari (2003). “Unsupervised curve clustering using B-splines”. In: *Scandinavian Journal of Statistics* 30.3, pp. 581–595.
- Aeberhard, S., D. Coomans, and O. D. Vel (1993). “Improvements to the classification performance of RDA”. In: *Journal of Chemometrics* 7.2, pp. 99–115.
- Akhanli, S. E. and C. Hennig (2020). “Comparing clusterings and numbers of clusters by aggregation of calibrated clustering validity indexes”. In: *Statistics and Computing* 30.5, pp. 1523–1544.
- Barati, Z., I. Zakeri, and K. Pourrezaei (2013). “Functional data analysis view of functional near infrared spectroscopy data”. In: *Journal of Biomedical Optics* 18.11, p. 117007.
- Bigelow, J. L. and D. B. Dunson (2009). “Bayesian semiparametric joint models for functional predictors”. In: *Journal of the American Statistical Association* 104.485, pp. 26–36.
- Binder, D. A. (1978). “Bayesian Cluster Analysis”. In: *Biometrika* 65.1, p. 31.
- Blei, D. M., A. Kucukelbir, and J. D. McAuliffe (2017). “Variational Inference: A Review for Statisticians”. In: *Journal of the American Statistical Association* 112.518, pp. 859–877.
- Boor, C. de (2001). *A Practical Guide to Splines - Revised Edition*.
- Boudt, K., P. J. Rousseeuw, S. Vanduffel, and T. Verdonck (2020). “The minimum regularized covariance determinant estimator”. In: *Statistics and Computing* 30.1, pp. 113–128.
- Bouveyron, C. (2014). “Adaptive mixture discriminant analysis for supervised learning with unobserved classes”. In: *Journal of Classification* 31.1, pp. 49–84.
- Butler, R. W., P. L. Davies, and M. Jhun (1993). “Asymptotics for the Minimum Covariance Determinant Estimator”. In: *The Annals of Statistics* 21.3, pp. 1385–1400.
- Canale, A., A. Lijoi, B. Nipoti, and I. Prünster (2017). “On the Pitman-Yor process with spike and slab base measure”. In: *Biometrika* 104.3, pp. 681–697.
- Cappozzo, A., F. Greselin, and T. B. Murphy (2020). “Anomaly and Novelty detection for robust semi-supervised learning”. In: *Statistics and Computing* 30.5, pp. 1545–1571.
- Carpenter, G. A., M. A. Rubin, and W. W. Streilein (1997). “ARTMAP-FD: familiarity discrimination applied to radar target recognition”. In: *Proceedings of International Conference on Neural Networks (ICNN'97)*. Vol. 3. IEEE, pp. 1459–1464.
- Cator, E. A. and H. P. Lopuhaä (2012). “Central limit theorem and influence function for the MCD estimators at general multivariate distributions”. In: *Bernoulli* 18.2, pp. 520–551.
- Charytanowicz, M., J. Niewczas, P. Kulczycki, P. A. Kowalski, S. Łukasik, and S. Zak (2010). “Complete gradient clustering algorithm for features analysis of X-ray images”. In: *Advances in Intelligent and Soft Computing* 69, pp. 15–24.

- Croux, C. and G. Haesbroeck (1999). “Influence Function and Efficiency of the Minimum Covariance Determinant Scatter Matrix Estimator”. In: *Journal of Multivariate Analysis* 71.2, pp. 161–190.
- De Blasi, P., A. F. Martínez, R. H. Mena, and I. Prünster (2020). “On the inferential implications of decreasing weight structures in mixture models”. In: *Computational Statistics and Data Analysis* 147.
- Dempster, A. N. Laird, and D. Rubin (1977). “Maximum likelihood from incomplete data via the EM algorithm”. In: *Journal of the Royal Statistical Society* 39.1, pp. 1–38.
- Escobar, M. D. and M. West (1995). “Bayesian Density Estimation and Inference Using Mixtures”. In: *Journal of the American Statistical Association* 90.430, pp. 577–588.
- Ferguson, T. S. (1973). “A Bayesian Analysis of Some Nonparametric Problems”. In: *The Annals of Statistics* 1.2, pp. 209–230.
- Fernández Pierna, J. A. and P. Dardenne (2007). “Chemometric contest at ‘Chimométrie 2005’: A discrimination study”. In: *Chemometrics and Intelligent Laboratory Systems* 86.2, pp. 219–223.
- Fop, M., P.-A. Mattei, C. Bouveyron, and T. B. Murphy (2021). “Unobserved classes and extra variables in high-dimensional discriminant analysis”. In.
- Forina, M., C. Armanino, M. Castino, and M. Ubigli (1986). “Multivariate data analysis as a discriminating method of the origin of wines”. In: *Vitis* 25.3, pp. 189–201.
- Gordaliza, A. (1991). “Best approximations to random variables based on trimming procedures”. In: *Journal of Approximation Theory* 64.2, pp. 162–180.
- Gutiérrez, L., E. Gutiérrez-Peña, and R. H. Mena (2014). “Bayesian nonparametric classification for spectroscopy data”. In: *Computational Statistics and Data Analysis* 78, pp. 56–68.
- Hennig, C. (2015). “What are the true clusters?” In: *Pattern Recognition Letters* 64, pp. 53–62.
- Hubert, M. and M. Debruyne (2010). “Minimum covariance determinant”. In: *Wiley interdisciplinary reviews: Computational statistics* 2.1, pp. 36–43.
- Hubert, M., M. Debruyne, and P. J. Rousseeuw (2018). “Minimum covariance determinant and extensions”. In: *Wiley Interdisciplinary Reviews: Computational Statistics* 10.3, pp. 1–11.
- Hubert, M. and K. Van Driessen (2004). “Fast and robust discriminant analysis”. In: *Computational Statistics & Data Analysis* 45.2, pp. 301–320.
- Ishwaran, H. and L. F. James (2001). “Gibbs Sampling Methods for Stick-Breaking Priors”. In: *Journal of the American Statistical Association* 96.453, pp. 161–173.
- Kalli, M., J. E. Griffin, and S. G. Walker (2011). “Slice sampling mixture models”. In: *Statistics and Computing* 21.1, pp. 93–105.
- Khan, S. S. and M. G. Madden (2014). “One-class classification: taxonomy of study and review of techniques”. In: *The Knowledge Engineering Review* 29.3, pp. 345–374.
- Lau, J. W. and P. J. Green (2007). “Bayesian model-based clustering procedures”. In: *Journal of Computational and Graphical Statistics* 16.3, pp. 526–558.
- Lo, A. Y. (1984). “On a Class of Bayesian Nonparametric Estimates: I. Density Estimates”. In: *The Annals of Statistics* 12.1, pp. 351–357.

- Malsiner-Walli, G., S. Frühwirth-Schnatter, and B. Grün (2016). “Model-based clustering based on sparse finite Gaussian mixtures”. In: *Statistics and Computing* 26.1-2, pp. 303–324.
- Manikopoulos, C. and S. Papavassiliou (2002). “Network intrusion and fault detection: a statistical anomaly approach”. In: *IEEE Communications Magazine* 40.10, pp. 76–82.
- Maronna, R. A. and V. J. Yohai (2017). “Robust and efficient estimation of multivariate scatter and location”. In: *Computational Statistics and Data Analysis* 109, pp. 64–75.
- McElhinney, J., G. Downey, and T. Fearn (1999). “Chemometric processing of visible and near infrared reflectance spectra for species identification in selected raw homogenised meats”. In: *Journal of Near Infrared Spectroscopy* 7.3, pp. 145–154.
- Meilä, M. (2007). “Comparing clusterings—an information based distance”. In: *Journal of Multivariate Analysis* 98.5, pp. 873–895.
- Miller, D. and J. Browning (2003). “A mixture model and EM algorithm for robust classification, outlier rejection, and class discovery”. In: *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03)*. Vol. 2. 11. IEEE, pp. II–809–12.
- Murphy, T. B., N. Dean, and A. E. Raftery (2010). “Variable selection and updating in model-based discriminant analysis for high dimensional data with food authenticity applications”. In: *The Annals of Applied Statistics* 4.1, pp. 396–421.
- Petrone, S., M. Guindani, and A. E. Gelfand (2009). “Hybrid dirichlet mixture models for functional data”. In: *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 71.4, pp. 755–782.
- Pitman, J. (1995). “Exchangeable and partially exchangeable random partitions”. In: *Probability Theory and Related Fields* 102.2, pp. 145–158.
- Pitman, J. and M. Yor (1997). “The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator”. In: *Annals of Probability* 25.2, pp. 855–900.
- Ramsay, James, Silverman, B. W (2005). *Functional Data Analysis*. Springer Series in Statistics. New York: Springer-Verlag.
- Rand, W. M. (1971). “Objective criteria for the evaluation of clustering methods”. In: *Journal of the American Statistical Association* 66.336, p. 846.
- Rasmussen, C. E. and C. K. I. Williams (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Rigon, T. (2019). “An enriched mixture model for functional clustering”. In: *arXiv*.
- Ritter, G. (2014). *Robust Cluster Analysis and Variable Selection*. Chapman and Hall/CRC.
- Rodriguez, A. and D. B. Dunson (2014). “Functional clustering in nested designs: Modeling variability in reproductive epidemiology studies”. In: *Annals of Applied Statistics* 8.3, pp. 1416–1442.
- Rousseau, J. and K. Mengersen (2011). “Asymptotic behaviour of the posterior distribution in overfitted mixture models”. In: *Journal of the Royal Statistical Society. Series B: Statistical Methodology* 73.5, pp. 689–710.
- Rousseeuw, P. J. (1984). “Least median of squares regression”. In: *Journal of the American statistical association* 79.388, pp. 871–880.

- Rousseeuw, P. J. and K. V. Driessen (1999). “A fast algorithm for the minimum covariance determinant estimator”. In: *Technometrics* 41.3, pp. 212–223.
- Scarpa, B. and D. B. Dunson (2009). “Bayesian hierarchical functional data analysis via contaminated informative priors”. In: *Biometrics* 65.3, pp. 772–780.
- Sethuraman, J. (1994). “A constructive definition of Dirichlet Process prior”. In: *Statistica Sinica* 4.2, pp. 639–650.
- Singh, M. and K. Domijan (2019). “Comparison of Machine Learning Models in Food Authentication Studies”. In: *2019 30th Irish Signals and Systems Conference (ISSC)*. IEEE, pp. 1–6.
- Strumiłło, A., J. Niewczas, P. Szczypiński, P. Makowski, and W. Woźniak (1999). *Computer system for analysis of x-ray images of wheat grains (a preliminary announcement)*.
- Tarassenko, L., P. Hayton, N. Cerneaz, and M. Brady (1995). “Novelty detection for the identification of masses in mammograms”. In: *IEE Conference Publication* 409, pp. 442–447.
- Tax, D. M. and R. P. Duin (1998). “Outlier detection using classifier instability”. In: *Joint IAPR international workshops on statistical techniques in pattern recognition (SPR) and structural and syntactic pattern recognition (SSPR)*. Springer, pp. 593–601.
- Todorov, V. and P. Filzmoser (2009). “An Object-Oriented Framework for Robust Multivariate Analysis”. In: *Journal of Statistical Software* 32.3, pp. 1–47.
- Wade, S. and Z. Ghahramani (2018). “Bayesian Cluster Analysis: Point estimation and credible balls (with Discussion)”. In: *Bayesian Analysis* 13.2, pp. 559–626.
- Yang, M., D. B. Dunson, and D. Baird (2010). “Semiparametric Bayes hierarchical models with mean and variance constraints”. In: *Computational Statistics and Data Analysis* 54.9, pp. 2172–2186.



# Supplementary Material

In this Supplementary Material, we report proofs for the theoretical results reported in the main paper and some additional numerical experiments, for both the multivariate Brand and its functional extension.

## 9 Proofs

First, we derive the moments, the variance, and the covariance for the simpler case (A):  $\tilde{p}_0 = \sum_{j=0}^J \pi_j \delta_{\Theta_j}$ . Then, we derive the same quantities starting from the discrete random measure that underlies Brand (B):  $\tilde{p} = \sum_{j=1}^J \pi_j \delta_{\Theta_j} + \pi_0 \left( \sum_{l \geq 1} \omega_l \delta_{\Theta_l^{ov}} \right)$ .

**9.1 Case A:**  $\tilde{p}_0 = \sum_{j=0}^J \pi_j \delta_{\Theta_j}$

$$\mathbb{E} [\Theta_m] = \mathbb{E} [\mathbb{E} [\Theta_m | \tilde{p}_0]] = \mathbb{E} \left[ \sum_{j=0}^J \pi_j \Theta_j \right] = \sum_{j=0}^J \mathbb{E} [\pi_j \Theta_j] = \sum_{j=0}^J \mathbb{E} [\pi_j] \mathbb{E} [\Theta_j] = \sum_{j=0}^J \frac{a_j}{a} \mu_j.$$

$$\mathbb{E} [\Theta_m]^2 = \sum_{j=0}^J \frac{a_j^2}{a^2} \mu_j^2 + 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a^2} \mu_j \mu_l.$$

$$\mathbb{E} [\Theta_m^2] = \sum_{j=0}^J \frac{a_j}{a} \mathbb{E} [\Theta_j^2] = \mathbb{E} \left[ \sum_{j=0}^J \pi_j \Theta_j^2 \right] = \sum_{j=0}^J \frac{a_j}{a} \mu_{j,2}.$$

$$\begin{aligned} \mathbb{E} [\Theta_m \Theta_{m'}] &= \mathbb{E} [\mathbb{E} [\Theta_m \Theta_{m'} | \tilde{p}]] = \mathbb{E} [\mathbb{E} [\Theta_m | \tilde{p}] \mathbb{E} [\Theta_{m'} | \tilde{p}]] = \mathbb{E} \left[ \left( \sum_{j=0}^J \pi_j \Theta_j \right)^2 \right] \\ &= \sum_{j=0}^J \mathbb{E} [\pi_j^2] \mathbb{E} [\Theta_j^2] + 2 \sum_{j>l \geq 0} \mathbb{E} [\pi_j \pi_l] \mathbb{E} [\Theta_j] \mathbb{E} [\Theta_l] \\ &= \sum_{j=0}^J \frac{a_j (a_j + 1)}{a(a+1)} \mu_{j,2} + 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a(a+1)} \mu_j \mu_l. \end{aligned}$$

$$\begin{aligned}
\mathbb{V}[\Theta_m] &= \sum_{j=0}^J \frac{a_j}{a} \mu_{j,2} - \sum_{j=0}^J \frac{a_j^2}{a^2} \mu_j^2 - 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a^2} \mu_j \mu_l \\
&= \sum_{j=0}^J \frac{a_j}{a} \left( \mu_{j,2} - \frac{a_j}{a} \mu_j^2 \right) - 2 \sum_{l>j \geq 0} \frac{a_j a_l}{a^2} \mu_l \mu_j.
\end{aligned}$$

$$\begin{aligned}
Cov_0(\Theta_m, \Theta_{m'}) &= \mathbb{E}[\Theta_m \Theta_{m'}] - \mathbb{E}[\Theta_m] \mathbb{E}[\Theta_{m'}] \\
&= \sum_{j=0}^J \frac{a_j(a_j+1)}{a(a+1)} \mu_{j,2} + 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a(a+1)} \mu_j \mu_l - \sum_{j=0}^J \frac{a_j^2}{a^2} \mu_j^2 - 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a^2} \mu_j \mu_l \\
&= \sum_{j=0}^J \left( \frac{a_j(a_j+1)}{a(a+1)} \mu_{j,2} - \frac{a_j^2}{a^2} \mu_j^2 \right) - 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a^2(a+1)} \mu_j \mu_l.
\end{aligned}$$

**9.2 Case B:**  $\tilde{p} = \sum_{j=0}^J \pi_j \delta_{\Theta_j} + \pi_0 \left( \sum_{l \geq 1} \omega_l \delta_{\Theta_l^{nov}} \right)$

$$\mathbb{E}[\Theta_m] = \sum_{j=0}^J \frac{a_j}{a} \mu_j.$$

$$\mathbb{E}[\Theta_m]^2 = \sum_{j=0}^J \frac{a_j^2}{a^2} \mu_j^2 + 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a^2} \mu_j \mu_l.$$

$$\mathbb{E}[\Theta_m^2] = \mathbb{E} \left[ \sum_{j=0}^J \pi_j \Theta_j^2 \right] = \sum_{j=0}^J \frac{a_j}{a} \mathbb{E}[\Theta_j^2] = \sum_{j=0}^J \frac{a_j}{a} \mu_{j,2}.$$

$$\begin{aligned}
\mathbb{E}[\Theta_m \Theta_{m'}] &= \mathbb{E}[\mathbb{E}[\Theta_m \Theta_{m'} | \tilde{p}]] = \mathbb{E}[\mathbb{E}[\Theta_m | \tilde{p}] \mathbb{E}[\Theta_{m'} | \tilde{p}]] = \mathbb{E} \left[ \left( \sum_{j=0}^J \pi_j \Theta_j + \pi_0 \sum_{l \geq 1} \omega_l \Theta_l^{nov} \right)^2 \right] \\
&= \mathbb{E} \left[ \left( \sum_{j=0}^J \pi_j \Theta_j \right)^2 + \left( \pi_0 \sum_{l \geq 1} \omega_l \Theta_l^{nov} \right)^2 + 2 \left( \sum_{j=0}^J \pi_j \Theta_j \right) \left( \pi_0 \sum_{l \geq 1} \omega_l \Theta_l^{nov} \right) \right] \\
&= \sum_{j=1}^J \frac{a_j(a_j+1)}{a(a+1)} \mu_{j,2} + 2 \sum_{j>l \geq 1} \frac{a_j a_l}{a(a+1)} \mu_j \mu_l \\
&\quad + \frac{a_0(a_0+1)}{a(a+1)} \left[ \frac{\mu_{0,2}}{1+\gamma} + \frac{\gamma \mu_0^2}{1+\gamma} \right] + 2\mu_0 \sum_{j=1}^J \frac{a_0 a_j}{a(a+1)} \mu_j \\
&= \sum_{j=0}^J \frac{a_j(a_j+1)}{a(a+1)} \mu_{j,2} + 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a(a+1)} \mu_j \mu_l - \frac{a_0(a_0+1)}{a(a+1)} \frac{\gamma}{1+\gamma} (\mu_{0,2} - \mu_0)^2.
\end{aligned}$$

$$\begin{aligned}
\mathbb{V}[\Theta_m] &= \sum_{j=0}^J \frac{a_j}{a} \mu_{j,2} - \sum_{j=0}^J \frac{a_j^2}{a^2} \mu_j^2 - 2 \sum_{j>l} \frac{a_j a_l}{a^2} \mu_j \mu_l \\
&= \sum_{j=0}^J \frac{a_j}{a} \left( \mu_{j,2} - \frac{a_j}{a} \mu_j^2 \right) - 2 \sum_{l>j \geq 0} \frac{a_j a_l}{a^2} \mu_l \mu_j.
\end{aligned}$$

$$\begin{aligned}
\text{Cov}_\gamma(\Theta_m, \Theta_{m'}) &= \mathbb{E}[\Theta_m \Theta_{m'}] - \mathbb{E}[\Theta_m] \mathbb{E}[\Theta_{m'}] \\
&= \sum_{j=0}^J \frac{a_j(a_j+1)}{a(a+1)} \mu_{j,2} + 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a(a+1)} \mu_j \mu_l - \frac{a_0(a_0+1)}{a(a+1)} \frac{\gamma}{1+\gamma} (\mu_{0,2} - \mu_0^2) \\
&\quad - \sum_{j=0}^J \frac{a_j^2}{a^2} \mu_j^2 - 2 \sum_{j>l \geq 0} \frac{a_j a_l}{a^2} \mu_j \mu_l \\
&= \text{Cov}_0(\Theta_m, \Theta_{m'}) - \frac{a_0(a_0+1)}{a(a+1)} \frac{\gamma}{1+\gamma} \sigma_0^2.
\end{aligned}$$

Let us represent all the possible values of  $\Theta$  in one vector  $\tilde{\Theta} = (\Theta_1, \dots, \Theta_J, \Theta_1^{nov}, \dots)$

$$\begin{aligned}
\mathbb{P}(\Theta_m = \Theta_{m'}) &= \mathbb{E}[\mathbb{P}(\Theta_m = \Theta_{m'} | \tilde{p})] \\
&= \sum_{j \geq 0} \mathbb{E} \left[ \mathbb{P}(\Theta_m = \tilde{\Theta}_j | \tilde{p}) \cdot \mathbb{P}(\Theta_{m'} = \tilde{\Theta}_j | \tilde{p}) \right] \\
&= \sum_{j=1}^J \mathbb{E} \left[ \mathbb{P}(\Theta_m = \tilde{\Theta}_j | \tilde{p}) \cdot \mathbb{P}(\Theta_{m'} = \tilde{\Theta}_j | \tilde{p}) \right] \\
&\quad + \sum_{j \geq J+1} \mathbb{E} \left[ \mathbb{P}(\Theta_m = \tilde{\Theta}_j | \tilde{p}) \cdot \mathbb{P}(\Theta_{m'} = \tilde{\Theta}_j | \tilde{p}) \right] \\
&= \sum_{j=1}^J \mathbb{E}[\pi_j^2] + \sum_{j \geq J+1} \mathbb{E}[\pi_0^2 \omega_j^2] \\
&= \sum_{j=1}^J \mathbb{E}[\pi_j^2] + \sum_{j \geq J+1} \mathbb{E}[\pi_0^2] \mathbb{E}[\omega_j^2] \\
&= \sum_{j=1}^J \frac{a_j(a_j+1)}{a(a+1)} + \frac{a_0(a_0+1)}{a(a+1)} \cdot \frac{1}{1+\gamma}.
\end{aligned} \tag{21}$$

## 10 Multivariate Brand - Additional experiments

The present section integrates and extends the simulation study reported in Section 6.1 of the manuscript. Particularly, the multivariate Brand is applied to a real 13-dimensional dataset (Section 10.1), and to two additional simulation studies: the former dealing with non-Gaussian shaped components (Section 10.2) and the latter encompassing higher-dimensional, sparse covariance structures (Section 10.3).

### 10.1 Wine dataset

The dataset, publicly available in the University of California Irvine Machine Learning repository, comprises 13 chemical measurements from 178 wine samples from the Piedmont region, Italy (Forina et al., 1986). The samples arise from three different cultivars: Barolo, Grignolino, and Barbera. We randomly select 66 samples from the first two varieties to build the training set, whereas the remaining 112 samples, including 48 wines from the third cultivar, defines the test set. The multivariate Brand methodology is fitted to the datasets to detect the third unobserved wine type. The model hyperparameters are set as follows. First,  $h_{MCD} = 0.95$  induces robust priors elicitation for the two observed classes. Fairly uninformative priors are used for the base measure  $H$ , namely  $\mathbf{m}_0 = \mathbf{0}$ ,  $\lambda_0 = 0.01$ ,  $\nu_0 = 15$  and  $\mathbf{S}_0 = \mathbf{I}_{13}$ , where with  $\mathbf{0}$  we denote the 13-dimensional zero vector. Such priors agree with the ones employed in Section 6.2.1 in the main paper for the seed dataset, underlying their general applicability in scenarios where no initial information is available. After initiating the MCMC with 20,000 iterations for the burn-in phase, 20,000 dependent samples are retained from the target posterior distribution. The resulting classification is reported in Table 3: our semi-parametric model almost perfectly recovers the underlying partition, efficiently identifying the novel wine type. The derived accuracy is particularly high, with performance comparable to those obtained in fully-supervised experiments (Aeberhard et al., 1993).

Table 3: Confusion matrix for multivariate Brand on the test set, wine dataset. The label “New” indicates observations that are estimated to have arisen from the novelty component.

Classification	Truth		
	Barolo	Grignolino	Barbera
Barolo	29	0	0
Grignolino	1	32	2
New	0	0	48

### 10.2 Non-Gaussian shaped classes

In this section, we employ the same hyperparameters configuration considered for the experimental setup in Section 6.1.1 of the main paper to generate samples from elliptical distributions that differ from the Gaussian, with the final aim of validating Brand




Figure 11: Synthetic data for the simulated experiments of Section 10.2. Classes 1 to 5 are generated according to a multivariate  $t$  distribution with 5 degrees of freedom, while classes 6 and 7 are realizations from a multivariate Laplace distribution.

performance under model misspecification. As discussed, the general structure reported in Equation (1) allows for any distributional kernel specification. However, departures from Gaussianity would induce the loss of the conjugacy properties, worsening the computational cost. Therefore, we now want to test how robust the Gaussian kernel is in capturing symmetric components with heavier tails. To this extent, we repeat the experiment for the `Label noise = False` and `Novelty size = Not small` scenario, with resulting sample sizes being equal to

$$n_1 = 300, n_2 = 300, n_3 = 400,$$

and

$$m_1 = 200, m_2 = 200, m_3 = 250, m_4 = 90, m_5 = 100, m_6 = 100, m_7 = 10,$$

for the training and test sets, respectively. In contrast to the simulations reported in the main paper, we generate groups 1 to 5 via a multivariate  $t$  distribution with 5 degrees of freedom, while classes 6 and 7 are realizations from a multivariate Laplace distribution. The resulting learning framework is displayed in Figure 11. The simulation results, for the same hyper-parameters specification employed in the main paper, are reported in Figure 12. We immediately notice that different prior settings do not substantially influence the overall model performance, with satisfactorily good results showcased for all the considered metrics. Notwithstanding, when we compare the simulation outcome in Figure 12 with the one displayed in Figure 4 of the main paper, it is apparent that the results are in some measure influenced by model misspecification.

In detail, the novelty term tends to absorb all those units sampled from the (heavy) tails of the known components, as they present very low Gaussian density, particularly when  $\eta_{MCD} = 0.75$ . Consequently, the clustering induced by the DPMM shows many more




Figure 12: Box plots for (from left to right) accuracy on the known classes, ARI and novelty predictive value metrics for  $B = 100$  repetitions of the simulated experiment of Section 10.2.

groups a posteriori in this scenario, with singletons trying to accommodate patterns that the main classes cannot explain. While clearly accuracy measures for assessing the goodness of a clustering procedure shall be application-dependent (Hennig, 2015), we argue here that, in principle, the outcome showcased by our method could still be relevant in contexts that violate the Gaussianity assumption. More specifically, if the sought classes are assumed to be unimodal and elliptical, both stages in the brand methodology concur to achieve this result. Stage I trims the group-wise most outlying values, and Stage II flexibly captures all the unexplained variability with Gaussian-like shapes. Once this has been accomplished, one can employ the output evaluation in terms of trimmed units and a posteriori assignment to determine which assumptions were not met by the dataset at hand. As an example, the identification of small novel clusters in the vicinity of bigger ones may be an indication that components with heavier tails are needed to properly account for the true underlying partition. For a general overview on the difficult problem of finding groups in data and associated clustering validity measures, the interested reader is referred to Akhanli and Hennig (2020).

### 10.3 High-dimensional heteroscedastic classes

For this experiment, we consider the same number of classes (known and hidden ones) and sample sizes for training and test sets introduced in the previous section. Each component is distributed according to a multivariate Normal density with mean vectors equal to:

$$\begin{aligned}
 \boldsymbol{\mu}_1 &= (-5, 5, -5, 5, -5, 5, 0, 0, 0, 0)', & \boldsymbol{\mu}_2 &= (-4, -4, -4, -4, -4, -4, 0, 0, 0, 0)', \\
 \boldsymbol{\mu}_3 &= (4, 4, 4, 4, 4, 4, 0, 0, 0, 0)', & \boldsymbol{\mu}_4 &= (0, 0, 0, 0, 0, 0, 0, 0, 0, 0)', \\
 \boldsymbol{\mu}_5 &= (5, -10, 5, -10, 5, -10, 0, 0, 0, 0)', & \boldsymbol{\mu}_6 &= (5, -10, 5, 10, -5, -10, 0, 0, 0, 0)',
 \end{aligned}$$

$$\boldsymbol{\mu}_7 = (-10, -10, -10, -10, -10, -10, 0, 0, 0, 0)',$$

and covariance matrices exhibiting different degrees of sparsity, as illustrated in Figure 13. We display the resulting learning framework in Figure 14, where the training and test




Figure 13: Covariance matrices for the simulated experiments of Section 10.3.

samples are reported in the lower and in the upper diagonal plots, respectively. Notice that the generative mechanism induces a quite challenging novelty detection problem, as the last four dimensions are irrelevant for group separation.

By again monitoring the metrics defined in the main manuscript, we aim at investigating the Brand performance when dealing with a 10-dimensional dataset with heterogeneous covariance patterns. Given the well-established efficacy of both the MCD and RMCD estimators in high-dimensional settings (Rousseeuw and Driessen, 1999; Hubert et al., 2018; Boudt et al., 2020), we focus here on the second stage of the Brand method, studying its sensitivity under different hyper-parameters specifications. In detail, a total of 8 models are fitted to  $B = 100$  simulated datasets varying:

- the concentration parameter  $\gamma$  of the Dirichlet Process prior, letting it be equal to 1 or 5,
- the degrees of freedom  $\nu_0$  associated with the novelty components, letting it be equal to 10 or 100,
- the precision parameter  $\lambda_0$  associated with the novelty components, letting it be equal to 0.01 or 1.

Simulation results are reported in Figure 15. We immediately notice that the overall performance showcased by our methodology is satisfactorily good, regardless of the hyperparameters specification. Both known and novel patterns are correctly identified as such, with results mirroring the ones obtained in the bi-dimensional experiments reported in Section 6.1.2 of the main paper. In particular, eliciting very flat priors for the base measure  $H$  ( $\lambda_0 = 0.01$ ,  $\nu_0 = 10$ ) still produces excellent outcomes. The only appreciable difference is in terms of ARI, induced by the DP concentration parameter  $\gamma$ . Specifically, increasing  $\gamma$  favors a priori the creation of more clusters. As a result, the model is more prone to accommodate components with a limited sample size, like the seventh group in the present experiment. On the other hand, when  $\gamma = 1$  the 20

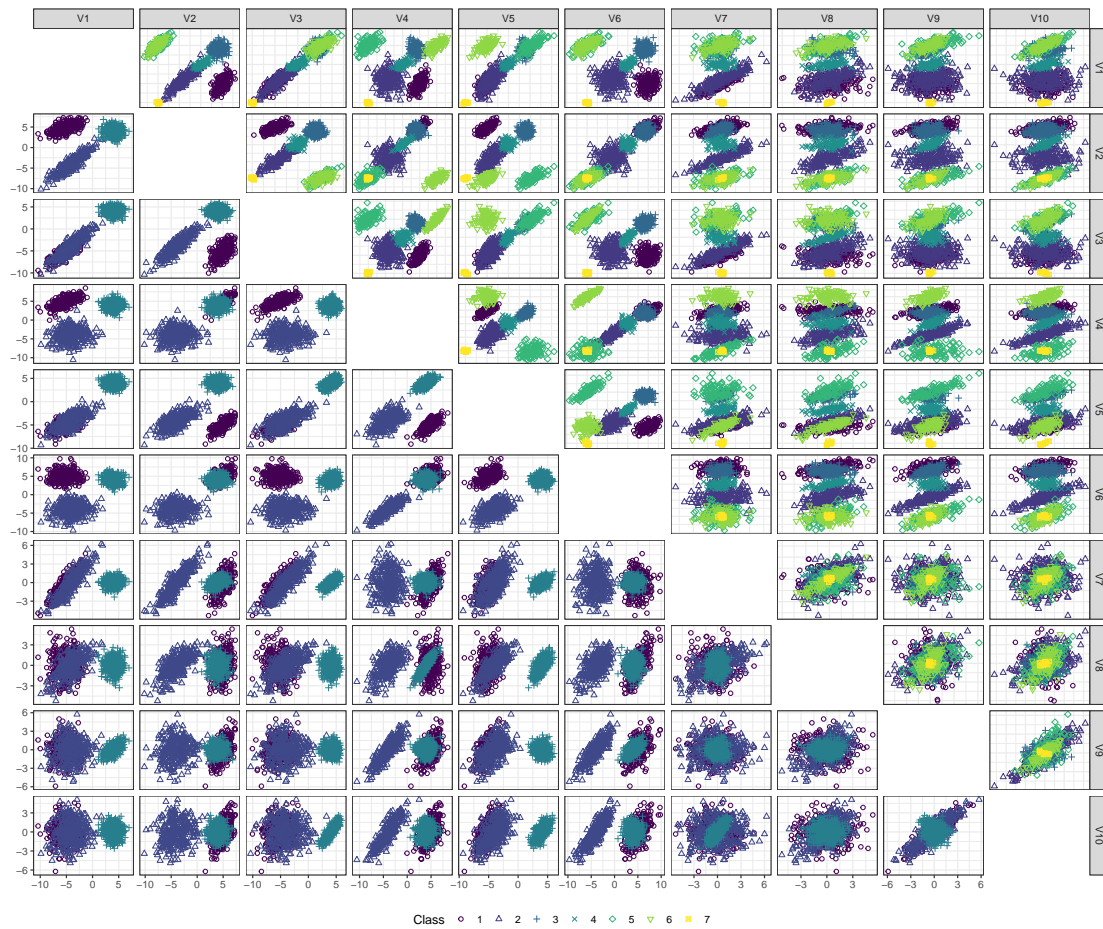


Figure 14: Learning scenario for the simulated experiment of Section 10.3. Plots below the main diagonal represent the training set. Plots above the main diagonal represent the test set.

test units arisen by the last density are usually merged within some other components, producing the slight difference in ARI visible in the middle plot of Figure 15.



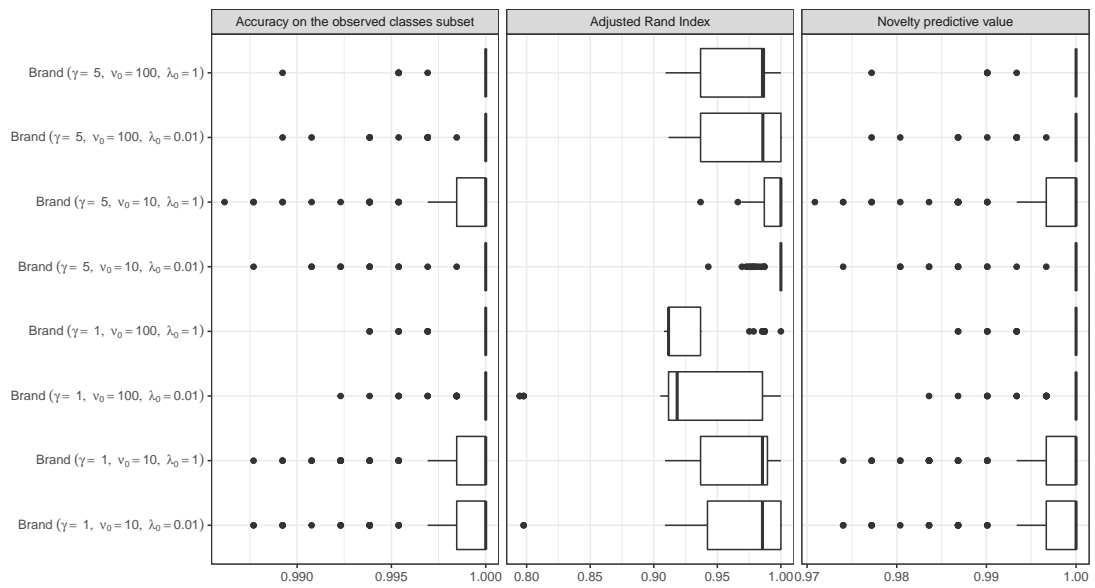


Figure 15: Box plots for (from left to right) accuracy on the known classes, ARI, and novelty predictive value metrics for  $B = 100$  repetitions of the simulated experiment of Section 10.3.

## 11 Functional Brand - Controlled experiment

This section investigates how functional Brand behaves according to various prior specifications and different levels of noise in the data. We aim to provide the researchers and practitioners with guidelines for the hyperparameter specification to exploit the flexibility of our model at its fullest. In the next experiments, we will consider the following six functions, evaluated on the interval  $t \in [0, 6]$ :

$$\begin{aligned} f_1(t) &= 5 \cos(\exp(\sin(t))), & f_2(t) &= 3 \log(\sin(t^{1.5}) + 1), \\ f_3(t) &= 2t \cos(t - 2.5), & f_4(t) &= -3|t - 1| \sin(t), \\ f_5(t) &= |t - 2| \cos(t), & f_6(t) &= |t - 1|^2 \sin(t). \end{aligned}$$

Figure 16 provides a visual representation of the six functions. Each of these functions presents its distinctive peculiarities. Simultaneously, they overlap, especially in the left half of the support, which could make the classification more difficult conditioning on the considered noise level. These six functions constitute the functional means of different groups we are going to study.




Figure 16: The six different functions considered in this experiment.

We suppose that every function is observed on a discretized collection of 100 time points between 0 and 6. In all our simulation settings, we fix the number of functions sampled from each group to 50. We assume that samples from the first three groups are contained in both the training set  $\mathbf{X}$  and in the test set  $\mathbf{Y}$ , while samples from the last three groups represent the novelties we want to detect, and therefore they appear only

in  $\mathbf{Y}$ . We are left with 150 data objects in the training sets and 300 in the test sets. More formally, for each group in the training set contained in both the training set  $\mathbf{X}$ , we have

$$x_j(t) = f_j(t) + \epsilon_j(t) \quad \text{with } j = 1, \dots, 3, \quad \text{and } \epsilon_j(t) \sim N(0, \sigma_\epsilon^2) \quad \forall t = 1, \dots, 100.$$

For each group in the test set  $\mathbf{Y}$ , we have

$$y_j(t) = f_j(t) + \epsilon_j(t) \quad \text{with } j = 1, \dots, 6, \quad \text{and } \epsilon_j(t) \sim N(0, \sigma_\epsilon^2) \quad \forall t = 1, \dots, 100.$$

Finally, we assume three different levels of noise in the data generating process, considering  $\sigma_\epsilon \in \{0.25, 0.50, 0.75\}$ , yielding three simulated cases of training and test sets  $(\mathbf{X}_1, \mathbf{Y}_1)$ ,  $(\mathbf{X}_2, \mathbf{Y}_2)$  and  $(\mathbf{X}_3, \mathbf{Y}_3)$ .

Figure 17 shows the three generated test sets, stratified by the increasing noise level. As already mentioned, a higher level of noise induces a more evident overlap among the functional objects. We expect this overlap to make the classification more challenging.




Figure 17: The test sets considered in the analysis stratified by the noise level used to generate the functional data. Left panel:  $\mathbf{Y}_1$ ; central panel:  $\mathbf{Y}_2$ ; right panel:  $\mathbf{Y}_3$ .

### 11.1 The robust estimation of the mean function

We focus on Stage I, and we investigate the robust extraction of prior information in the functional case. In the main paper, we discuss how we smooth the training functions using B-splines bases. Similarly, for this study, we use 100 bases of order 5, which provide

accurate results for a reasonable computational cost.

Our functional robust estimation works as follows. We first collect the spline coefficients, and we interpret them as multivariate objects. We then apply the MCD estimator on the spline coefficients. We therefore recover, for each group, the robust mean coefficients that, once convoluted with the bases, will produce our mean functions. A crucial quantity is the percentage of extreme coefficients to remove in our robust estimation. This percentage significantly affects also the recovery of the training noise  $\bar{\sigma}_j^2(t)$ . In the following, we demonstrate how the ratio of trimmed values  $\eta_{MCD}$  affects Stage I in the functional case.

To showcase the effect of the MCD estimator, we contaminate the known classes in the training sets with time-point specific outliers and label noise. In detail, we add noise sampled from  $\varepsilon(t) \sim N(0, 25)$  to fifteen randomly selected functions at fifteen random timestamps. We contaminate five functions from the first group, four in the second, and six in the third. Then, we shuffle 10% of the functions across groups to generate label noise. To illustrate this step, we report in Figure 18 a visual depiction of the contaminated version of the training dataset  $\mathbf{X}_3$ , stratified by class.




Figure 18: Functional objects contained in  $\mathbf{X}_3$  - generated specifying  $\sigma_\varepsilon = 0.75$ . Each known class is characterized by label noise and time-specific outliers (highlighted by red dots).

We consider three different specifications of  $\eta_{MCD}$ :

- **no trimming** by setting  $\eta_{MCD} = 1$ : all the estimated spline coefficients are considered;

- **mild trimming**,  $\eta_{MCD} = 0.95$ : only 5% of the estimated spline coefficients are removed for the estimation of the group centroids and functional noise;
- **strong trimming**,  $\eta_{MCD} = 0.75$ : a fourth of the estimated spline coefficients is removed before computing the group centroids and functional noise.

We report the extracted mean functions and the estimated noise in the panels of Figure 19. Specifically, the group centroids are depicted in red, superimposed onto the observations of every group. We also plot in blue the intervals of variation computed as  $\bar{f}_j(t) \pm q \cdot \bar{\sigma}_j^2(t)$  to give an idea of the estimated noise, for  $j, q = 1, 2, 3$ . Each column of the plot shows the evolution of the mean and noise functions when considering different trimming levels, stratified by group. We can observe how employing trimming benefits the estimation of the representative functions in each group.

For example, in the first row of Figure 19, we see how the red line fails to represent the group characteristics. This behavior is exacerbated close to the boundaries in all the groups. It is evident how even a small percentage of trimming yields considerable benefits. In our experience, we obtain very satisfactory results in the functional case when considering a **strong trimming**. Results in a case of **extreme trimming**, obtained by setting  $\eta_{MCD} = 0.5$  were explored but not reported, since similar to the case  $\eta_{MCD} = 0.75$ . In conclusion, if enough data are available, we suggest considering strong trimming as the safest option when dealing with the functional case. In this way, we can eradicate the influence of potential outliers on both the mean and the variance. We will use the mean and noise functions recovered with  $\eta_{MCD} = 0.75$  for the subsequent analyses.

## 11.2 Model performance according to different prior specifications

We can summarise the hyperparameters specific of functional Brand in three main subsets:

- **Subset 1**: hyperparameters controlling the dispersion around the extracted mean functions of the known groups:  $\varphi_j$  and  $\nu_j$ ;
- **Subset 2**: hyperparameters controlling the dispersion around the novel spline coefficients, representing the unknown mean functions:  $a_\tau$ ,  $b_\tau$ , and  $s^2$ ;
- **Subset 3**: hyperparameters controlling the dispersion of the noise around the novel mean functions:  $a_H$  and  $b_H$ .

We investigate the sensitivity of functional Brand to various hyperparameter specifications while considering the different noise levels in the observed data, as previously discussed. As a first consideration, we point out that setting the hyperparameters controlling the variance around the known classes and/or the novel ones requires the most care. A poorly elicited prior for the error terms could lead to losing Brand hierarchical structure, letting the novelty components take over the known groups in lieu of their flexibility. The opposite issue may also arise, since larger mixture weights are assigned to the known components. In the following, we will provide some considerations regarding the

tuning of the hyperparameters listed in **Subset 1** and **Subset 2**. For the sake of brevity, we omit the results obtained tweaking the parameters that belong to **Subset 3**, which control the noise level in the novelty term  $\sigma_h^{2nov}(t)$ ,  $\forall h$ . Given the insights collected in several applications, we suggest eliciting an Inverse Gamma prior that favors low noise values, promoting more sensitive estimation of the random functions  $f_h^{nov}(t)$ ,  $\forall h$ . Thus, we fix  $a_H = 5$  and  $b_H = 1$  in the following, as well as in the functional application in the main paper.

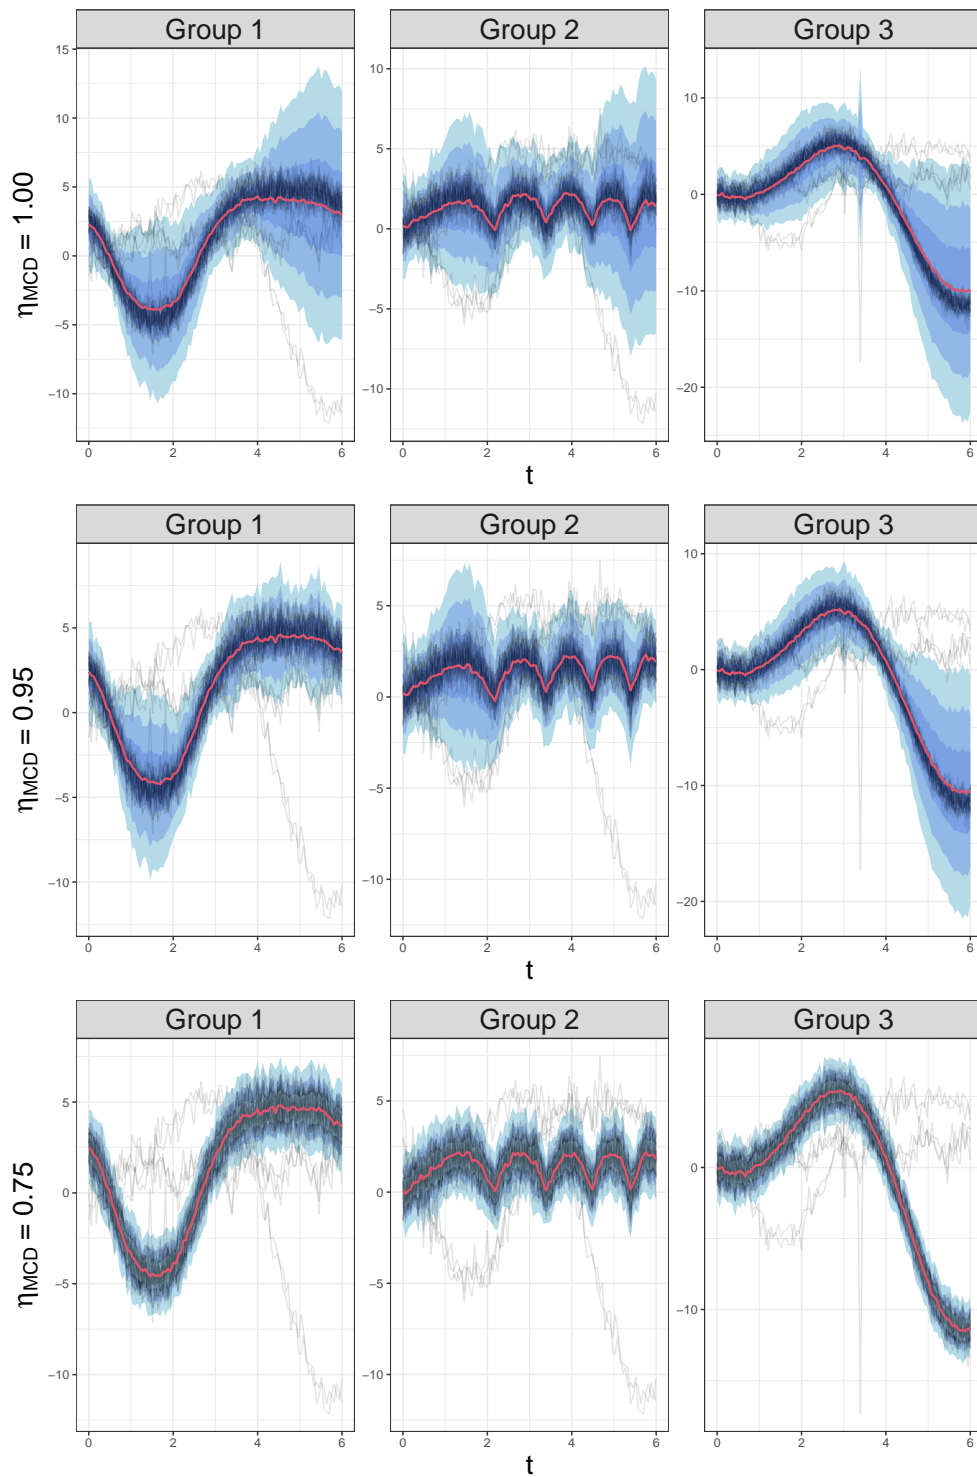


Figure 19: Estimated representative function (in red) for each of the known groups, characterized by the presence of label noise and time-specific outliers. Each row displays a different level of trimming, while groups are reported column-wise.

Let us first discuss the hyperparameter specification that we adopted in the main text. We are going to assume these values as our default setting. The values are

$$\varphi_j \approx 0, \quad \nu_j \approx 0, \quad a_\tau = 3, \quad b_\tau = 1, \quad s^2 = 1 \quad a_H = 5, \quad b_H = 1.$$

First, let us describe and justify our choice of hyperparameters. The parameters in **Subset 1** are set to negligible values. In this way, we are implicitly enforcing an **inductive** learning approach: the information extracted from the training set for the known groups is trustworthy and need not be updated. We calibrate the parameters in **Subset 2** to induce a reasonably tight prior over the novel splines coefficients to help the model identifiability. One can claim that an uninformative specification would also be reasonable to reflect our ignorance about the novel groups. However, as we will show, letting the parameters to vary too freely may negatively affect the results. Finally, as already discussed, the informative specification of the parameters in **Subset 3** bounds the variance of the errors around the novel functions: this will avoid limiting cases where all the information contained in the data is explained by the noise term, helping the identification of the latent mean functions.

In the following, we tweak one subset of parameters at a time to investigate its effect, starting from this prior configuration. In all the subsequent experiments, we run 7,500 MCMC iterations after a burn-in period of the same length.

### 11.2.1 Tuning hyperparameters: Subset 1

For this first experiment, we consider the three hyperprior configurations devised to increase the uncertainty around the functional means and variance extracted from the training set. In other words, we want to understand how, *ceteris paribus*, the concentration of the priors for  $f_j(t)$  and  $\sigma_j^2(t)$ ,  $j = 1, \dots, J$  affects the estimation of the mean functions of *the known groups*.

As already mentioned, a proper elicitation of the prior noises  $\sigma_j^2(t)$  for the known classes is crucial for carrying out sensible inference. Throughout many simulation studies, we observed that  $\sigma_j^2(t)$  plays a fundamental role in preserving the specific hierarchy that characterizes Brand. In fact, in scenarios where the prior variance of the measurement errors around the training mean is too wide, the known components lose their “priority” over the novel ones. In other words, the robust estimates extracted from the training set are deemed as less representative of the truth. Simultaneously, the flexibility of the novelty terms leaves the new estimated functions free to adapt to the dataset, ending up absorbing the known classes. In short, if the variances of the known components are not carefully tuned, Brand loses the ability to distinguish between known groups and novelties, hindering the resulting classification. To provide an idea on how different choices of  $\varphi_j$ , and  $\nu_j$  affect the results, we consider:

HPC 1 - Inductive setting:  $\varphi_j \approx 0, \quad \nu_j \approx 0,$

HPC 2 - Informative specification:  $\varphi_j = 0.01, \quad \nu_j = 1/100000,$



HPC 3 - “Uninformative” specification:  $\varphi_j = 1$ ,  $\nu_j = 1/100$ . The variance around the mean function is set to 1.

For each configuration, we estimate functional Brand on the three different test sets  $\mathbf{Y}_1$ ,  $\mathbf{Y}_2$ , and  $\mathbf{Y}_3$ . We show the confusion matrices displaying the assignation of the functional objects contained in the test set into known classes or novelties in Table 4. Moreover, to better exemplify the model’s behavior, we display the functional estimates of the known mean functions for  $\mathbf{Y}_3$  in Figure 20. For each known group, we collect the robustly extracted mean functions (in red), the simulated posterior mean functions (computed pointwise, in blue) on top of the MCMC iterations (in black).

We can detect a recurring pattern across different datasets. On the one hand, the first and second prior configurations (HPC 1 and HPC 2) lead to perfect classification of the functional data into known components (Classes 1, 2, and 3 into Cluster 1, 2, and 3, respectively) and novelties (Classes 4, 5, and 6 into Cluster 0) regardless the level of noise in the dataset. Looking at the different panels of Figure 20, we see how HPC 2 allows more uncertainty in the mean estimates (represented by a wider range of variation of the MCMC simulations) and a difference between the starting robust means (in red) and the MCMC posterior means (in blue). As expected, in the `inductive` setting, the red and blue lines coincide.

On the other hand, we can see what happens when tight constraints are not placed on the variances  $\sigma_j^2(t)$  by focusing on the last block of columns in Table 4 and the last row of panels in Figure 20. In all the three cases  $\mathbf{Y}_1$ ,  $\mathbf{Y}_2$ , and  $\mathbf{Y}_3$  we see how the novelty component takes over and ruins the classification process. For example, in the bottom row of Figure 20 we see that in the three panels the estimated posterior mean function coincides with the prior since no data have been assigned to those components. In the first two datasets, we see how the higher variances lead Brand to mislead a novel group as a known one and viceversa.

We can conclude that, for the functional setting, care is needed when tuning the discussed hyperprior parameters. Thus, when enough data are available, we suggest either to rely on the `inductive` learning process or, otherwise, to induce small a priori variations.

### 11.2.2 Tuning hyperparameters: Subset 2

The discussion for this second set of hyperparameters is structured similarly to the one reported in the previous Subsection. We now investigate Brand’s ability to identify and correctly cluster the novelty terms, while tweaking the variability of the prior distributions of  $\psi_h$  and  $\tau_h^2$  (and, therefore, of  $f_h^{nov}(t)$ ). We consider the following three hyperparameters configurations:

HPC 1 - Adopted specification:  $a_\tau = 3$ ,  $b_\tau = 1$ ,  $s^2 = 1$

HPC 2 - Low variability:  $a_\tau = 10$ ,  $b_\tau = 1$ ,  $s^2 = 0.001$

HPC 3 - High variability:  $a_\tau = 0.1$ ,  $b_\tau = 0.1$ ,  $s^2 = 10$ .

We have also investigated a more extreme case,  $a_\tau = 0.01$ ,  $b_\tau = 0.01$ ,  $s^2 = 1000$ , for

Ground Truth	HPC 1						HPC 2						HPC 3						
	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	
$\mathbf{Y}_1$	<b>0</b>	0	0	0	50	50	50	0	0	0	50	50	50	50	50	50	0	50	50
	<b>1</b>	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	50	0	0
	<b>2</b>	0	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0	0
	<b>3</b>	0	0	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0
$\mathbf{Y}_2$	<b>0</b>	0	0	0	50	50	50	0	0	0	50	50	50	50	50	50	0	50	50
	<b>1</b>	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	50	0	0
	<b>2</b>	0	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0	0
	<b>3</b>	0	0	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0
$\mathbf{Y}_3$	<b>0</b>	0	0	0	50	50	50	0	0	0	50	50	50	50	50	50	50	50	50
	<b>1</b>	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0	0	0
	<b>2</b>	0	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0	0
	<b>3</b>	0	0	50	0	0	0	0	0	50	0	0	0	0	0	0	0	0	0

Table 4: Confusion matrices obtained in the first part of the simulation study. The different datasets are reported along the rows. The different hyperprior specifications are reported along the columns.

which we omit the results since this extremely uninformative specification resulted in numerical instabilities and poor classification performance.

The first and second hyperprior specifications lead to perfect clustering for all three datasets. We instead report in Table 5 the results obtained on  $\mathbf{Y}_3$ . A tighter prior for the novelty components leads to the misclassification of only few observations (2/300) and the correct detection of the overall cluster structure. Conversely, a wide prior can overrule the hierarchical structure of the model, absorbing known classes into the novelty term.

Overall, we can conclude that Brand is more robust to different hyperprior settings regarding the specification of the novelty components. Our specification (HPC 1) provides a nice balance between the flexibility required by the novelty term and the stability of the results.




Figure 20: Subset 1, Dataset  $\mathbf{Y}_3$ . The known groups are presented column-wise, while the various hyperprior specification are reported row-wise. The panels contain the extracted robust mean function estimates (in red), the simulated posterior mean functions (computed pointwise, in blue) on top of the MCMC iterations (in black).

$\mathbf{Y}_3$		HPC 1						HPC 2						HPC 3					
Ground Truth		<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>	<b>1</b>	<b>2</b>	<b>3</b>	<b>4</b>	<b>5</b>	<b>6</b>
CKC	<b>1</b>	50	0	0	0	0	0	50	0	0	0	0	0	50	0	0	0	0	0
	<b>2</b>	0	50	0	0	0	0	0	50	0	0	0	0	0	1	0	0	0	0
	<b>3</b>	0	0	50	0	0	0	0	0	49	0	0	0	0	47	0	0	0	0
CNC	<b>4</b>	0	0	0	50	0	0	0	0	1	50	0	0	0	1	0	50	0	0
	<b>5</b>	0	0	0	0	50	0	0	0	0	0	50	0	0	0	50	0	47	0
	<b>6</b>	0	0	0	0	0	50	0	0	0	0	0	49	0	0	0	0	0	50
ENC	<b>7</b>	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	1	0
	<b>8</b>	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0
	<b>9</b>	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	1	0

Table 5: Classification results for dataset  $\mathbf{Y}_3$  under the three different hyperparameter configurations. The table displays the ground truth across the columns (1-3: known groups, 4-6 novel groups). We divided the resulting labels from the clustering process (on the rows) into three subgroups: CKC, correct known clusters (1-3); CNC, correct novel clusters (4-6); and ENC, extra novel clusters.

## 12 Additional Figures

### 12.1 Meat variety dataset analysis

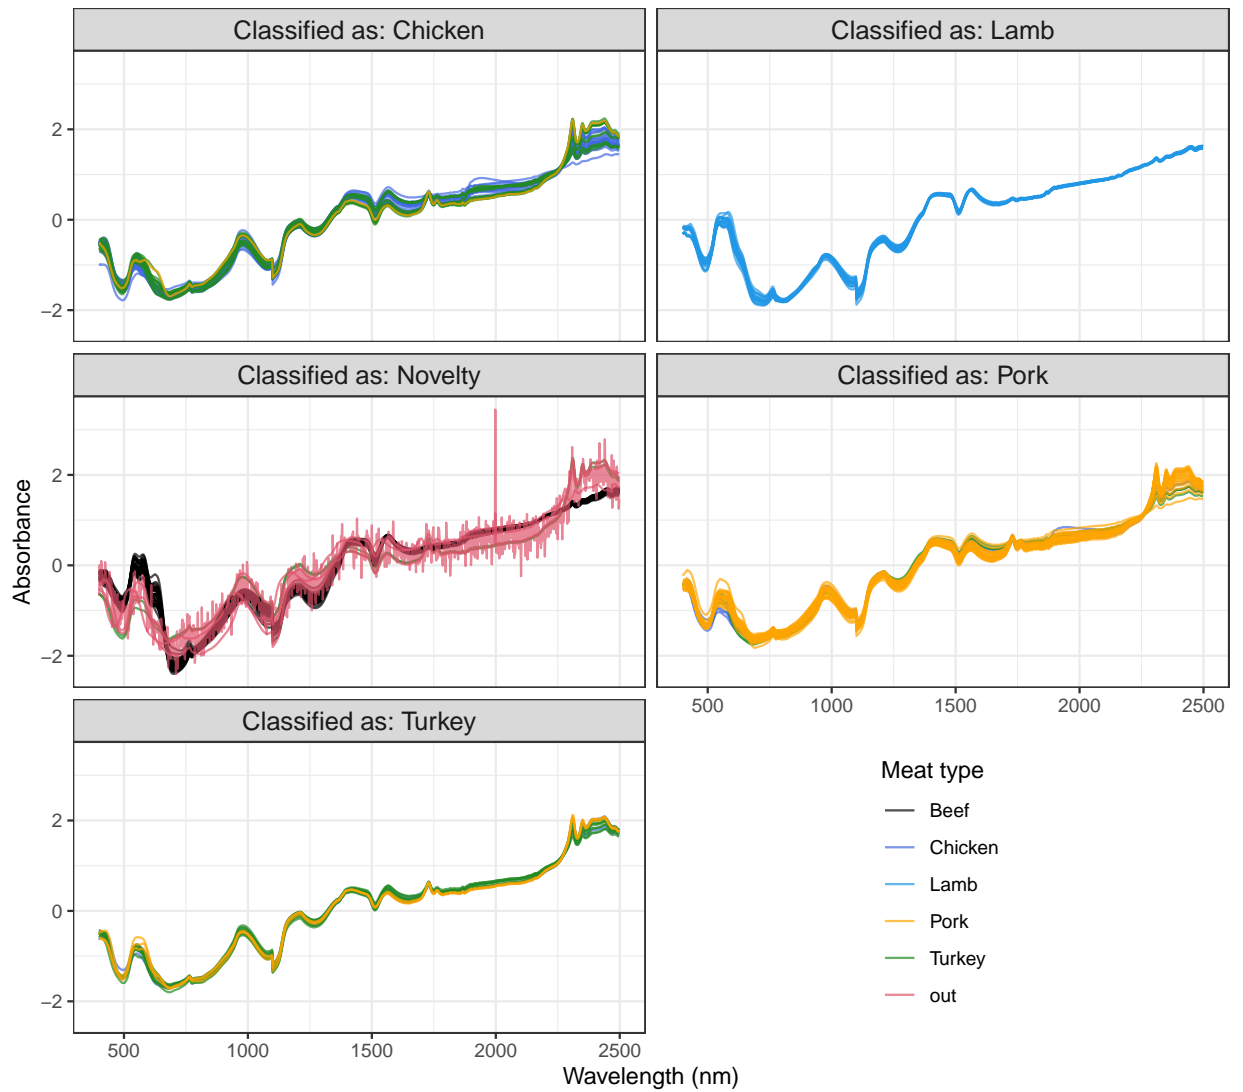


Figure 21: Summary of the resulting classification given by the model considered in Section 6.3. Each functional object is colored according to its correct data type, while each panel contains the meat spectra that Brand classifies together.

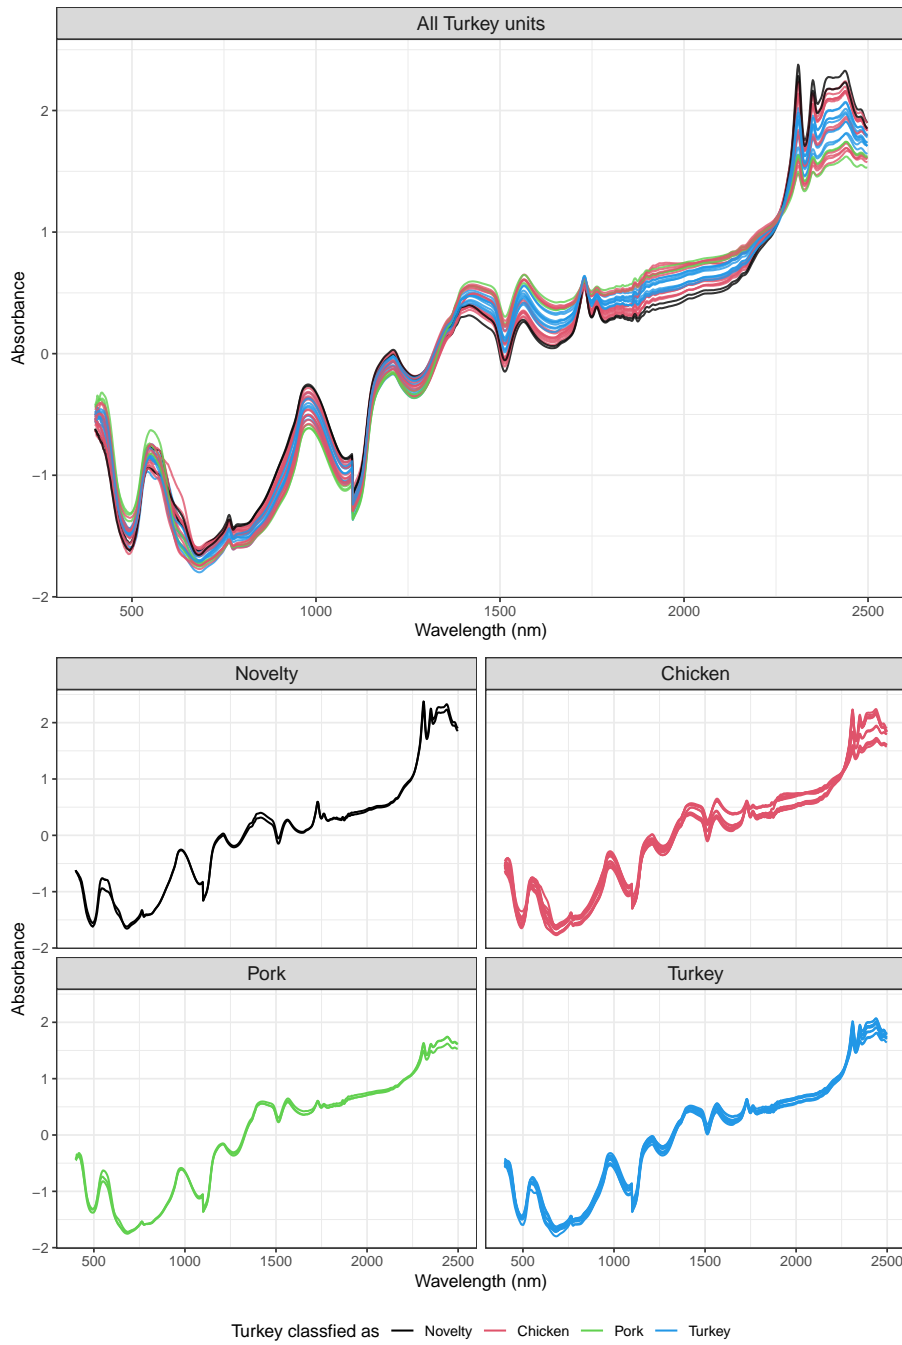


Figure 22: Classification of the turkey units according to Brand. The different colors indicate the estimated labels. The top panel shows all the functional data, while the bottom four panels break down the data objects in the four classes, highlighting their different characteristics.

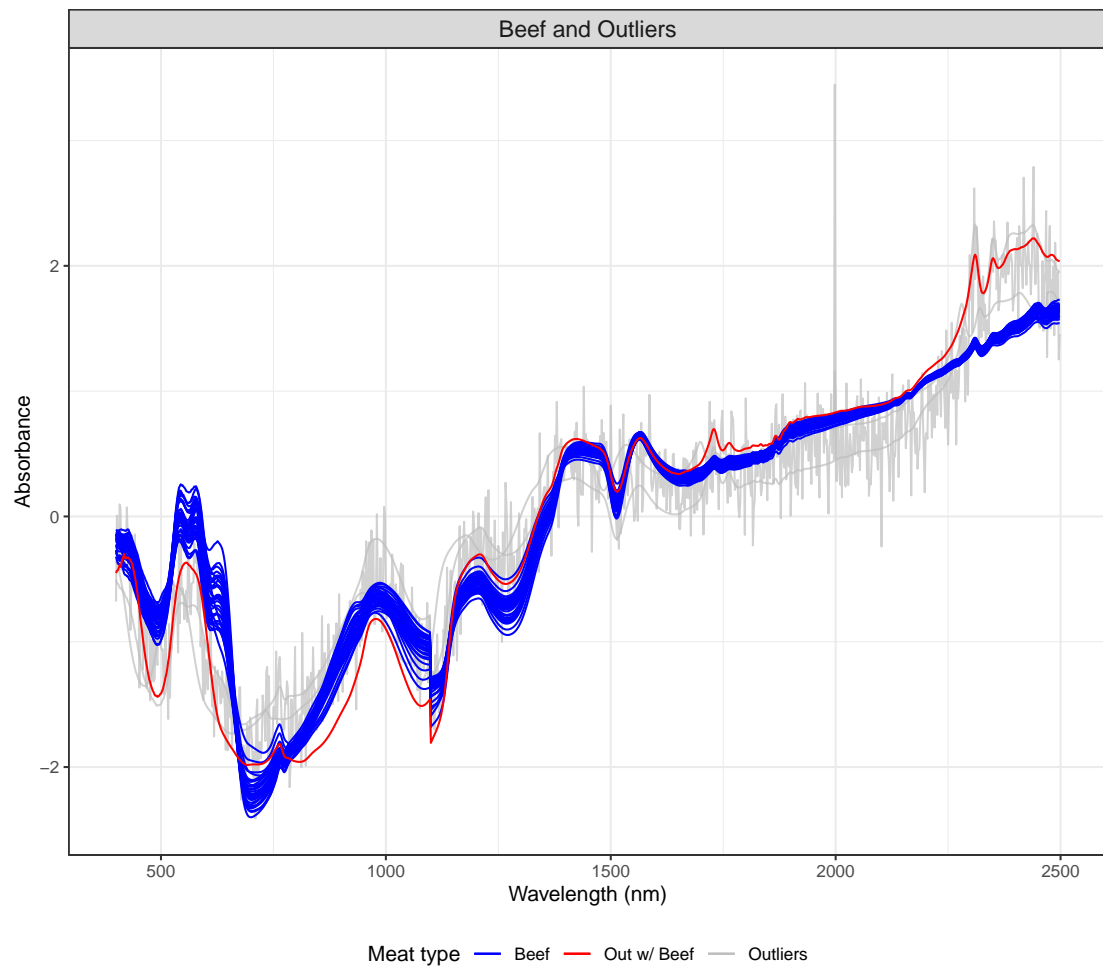


Figure 23: Functional objects classified as Novelties, colored according to their assigned cluster (Beef, in blue and Outliers, in gray). We highlighted in red the outlying functional data that is assigned to the Beef cluster.

## 13 Additional Table for the Simulation Study of Section 6.1

Table 6: Accuracy on the known classes, Adjusted Rand Index and Novelty predictive value metrics for  $B = 100$  repetitions of the simulated experiment, varying data contamination level and test set sample size. Standard errors are reported in parentheses.

	Label noise = False			Label noise = True		
	Accuracy	ARI	Precision	Accuracy	ARI	Precision
<hr/> Novelty Size = Not small <hr/>						
<i>AMDA</i>	0.999 (0.002)	0.932 (0.014)	0.998 (0.004)	0.562 (0.223)	0.814 (0.03)	0.519 (0.215)
<i>RAEDDA</i>	0.966 (0.029)	0.885 (0.041)	0.934 (0.051)	0.996 (0.003)	0.924 (0.003)	0.992 (0.005)
<i>Brand</i> ( $\eta_{MCD} = 1, \lambda_{Tr} = 1,000$ )	0.998 (0.001)	0.93 (0.005)	0.997 (0.002)	0.309 (0.001)	0.931 (0.002)	0.4 (0.001)
<i>Brand</i> ( $\eta_{MCD} = 1, \lambda_{Tr} = 10$ )	0.995 (0.001)	0.927 (0.005)	0.992 (0.002)	0.465 (0.206)	0.931 (0.01)	0.481 (0.13)
<i>Brand</i> ( $\eta_{MCD} = 0.75, \lambda_{Tr} = 1,000$ )	0.99 (0.007)	0.928 (0.007)	0.985 (0.013)	0.997 (0.004)	0.93 (0.008)	0.995 (0.009)
<i>Brand</i> ( $\eta_{MCD} = 0.75, \lambda_{Tr} = 10$ )	0.997 (0.001)	0.928 (0.005)	0.994 (0.002)	0.994 (0.001)	0.927 (0.007)	0.99 (0.001)
<hr/> Novelty Size = Small <hr/>						
<i>AMDA</i>	0.999 (0.002)	0.986 (0.004)	0.996 (0.011)	0.69 (0.29)	0.905 (0.026)	0.459 (0.353)
<i>RAEDDA</i>	0.969 (0.015)	0.947 (0.019)	0.855 (0.06)	0.996 (0.003)	0.981 (0.003)	0.977 (0.014)
<i>Brand</i> ( $\eta_{MCD} = 1, \lambda_{Tr} = 1,000$ )	0.998 (< 0.01)	0.986 (0.001)	0.994 (0.002)	0.413 (0.003)	0.986 (0.003)	0.23 (0.001)
<i>Brand</i> ( $\eta_{MCD} = 1, \lambda_{Tr} = 10$ )	0.999 (0.01)	0.986 (0.001)	0.996 (0.002)	0.504 (0.222)	0.986 (0.003)	0.289 (0.285)
<i>Brand</i> ( $\eta_{MCD} = 0.75, \lambda_{Tr} = 1,000$ )	0.992 (0.022)	0.985 (0.006)	0.979 (0.056)	0.996 (0.001)	0.985 (0.002)	0.986 (0.007)
<i>Brand</i> ( $\eta_{MCD} = 0.75, \lambda_{Tr} = 10$ )	0.999 (< 0.01)	0.986 (< 0.01)	0.995 (0.002)	0.999 (0.001)	0.986 (0.001)	0.997 (0.004)