



HAL
open science

Label Shift Quantification with Robustness Guarantees via Distribution Feature Matching

Bastien Dussap, Gilles Blanchard, Badr-Eddine Chérif-Abdellatif

► To cite this version:

Bastien Dussap, Gilles Blanchard, Badr-Eddine Chérif-Abdellatif. Label Shift Quantification with Robustness Guarantees via Distribution Feature Matching. ECML-PKDD 2023, 2023, Turin (IT), Italy. pp.69-85, <10.1007/978-3-031-43424-2_5>. <hal-04122205>

HAL Id: hal-04122205

<https://hal.science/hal-04122205v1>

Submitted on 7 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-SA 4.0 - Attribution - Non-commercial use - ShareAlike - International License

Label Shift Quantification with Robustness Guarantees via Distribution Feature Matching

Bastien Dussap¹✉, Gilles Blanchard¹, Badr-Eddine Chérief-Abdellatif²
 bastien.dussap@univerite-paris-saclay.fr,
 gilles.blanchard@univerite-paris-saclay.fr, and
 badr-eddine.cherief-abdellatif@cnrs.fr

¹ Université Paris-Saclay, CNRS, Inria, Laboratoire de mathématiques d'Orsay

² CNRS, LPSM, Sorbonne Université, Université Paris Cité, France

Abstract. Quantification learning deals with the task of estimating the target label distribution under label shift. In this paper, we first present a unifying framework, distribution feature matching (DFM), that recovers as particular instances various estimators introduced in previous literature. We derive a general performance bound for DFM procedures, improving in several key aspects upon previous bounds derived in particular cases. We then extend this analysis to study robustness of DFM procedures in the misspecified setting under departure from the exact label shift hypothesis, in particular in the case of contamination of the target by an unknown distribution. These theoretical findings are confirmed by a detailed numerical study on simulated and real-world datasets. We also introduce an efficient, scalable and robust version of kernel-based DFM using Random Fourier Features.

Keywords: Learning Theory · Quantification · Kernel Mean Embedding · Label Shift · Class ratio estimation.

1 Introduction

The success of supervised learning over the last decades is mainly based on the belief that the training and test samples follow the same data generation process. However, in real-world applications, this assumption is often violated and classical learning methods are challenged. Unsupervised domain adaptation, a field of transfer learning, specifically addresses this problem by transferring knowledge from the different but related *training* or *source* domain to the *test* or *target* domain of interest [34,32].

From a formal point of view, consider a covariate space \mathcal{X} , typically a subset of \mathbb{R}^d , and a label space $\mathcal{Y} = [c] := \{1, \dots, c\}$. We define the two *source* and *target* domains as different probability distributions over the covariate-label space pair $\mathcal{X} \times \mathcal{Y}$. The target label distribution is denoted $\alpha^* = (\alpha_i^*)_{i=1}^c$ while each class- i conditional target distribution is denoted \mathbb{Q}_i . Similarly, the source label distribution is denoted $\beta^* = (\beta_i^*)_{i=1}^c$ while each class- i conditional source

distribution is denoted \mathbb{P}_i . We will consider the classical label shift hypothesis:

$$\forall i = 1, \dots, c, \quad \mathbb{P}_i = \mathbb{Q}_i. \quad (\mathcal{LS})$$

Another setting we will consider involves contamination of the target by a new class. In this *contaminated label shift* setting, we assume that the target domain is $\mathcal{X} \times \tilde{\mathcal{Y}}$, with $\tilde{\mathcal{Y}} = \{0, \dots, c\}$ and that the label shift hypothesis is still verified for the class $\{1, \dots, c\}$:

$$\mathbb{Q} = \sum_{i=1}^c \alpha_i^* \mathbb{P}_i + \alpha_0^* \mathbb{Q}_0 \quad (\mathcal{CLS})$$

$$\forall i = 1, \dots, c, \quad \mathbb{P}_i = \mathbb{Q}_i.$$

The distribution \mathbb{Q}_0 is seen as a noise or a contamination, for which we have no prior knowledge nor sample. Therefore, our objective in this contaminated scenario is to be robust to a large class of noise distributions. In Section 3.2, we will give insight on the kind of contamination we can be robust to.

In both settings, we suppose a source dataset $\{(x_j, y_j)\}_{j \in [n]} \in (\mathcal{X} \times \mathcal{Y})^n$ and a target dataset $\{x_{n+j}\}_{j \in [m]} \in \mathcal{X}^m$ are given. All data points from the source (respectively the target) dataset are independently sampled from the source (resp. the target) domain. We have access to the source labels y_j but not to the target labels which are not observed. We denote by $\hat{\mathbb{P}}_i := \sum_{j \in [n]: y_j = i} \delta_{x_j}(\cdot) / n_i$ the empirical source class- i conditional distribution, where δ_{x_j} denotes the Dirac measure at point x_j and n_i the number of instances labeled i in the source dataset. Note that $n_1 + \dots + n_c = n$. We finally denote by $\hat{\beta}$ the empirical proportions of each class in the source dataset, i.e. $\hat{\beta}_i := n_i / n$.

Several different objectives have been addressed under the label shift assumption in the literature, and can be summarised in three points: (i) *detection*, i.e. determining whether distribution shift has occurred; (ii) *correction*, i.e. fitting a classifier with high accuracy on the target distribution; and (iii) *quantification*, i.e. estimating the target label distribution [38,15,19,27,3,18]. We focus here on the last challenge, and develop a general analysis unifying several existing techniques with estimation guarantees for the target proportions α^* , as well as dealing with the contaminated label shift setting.

1.1 Related literature

The research area of quantification learning has a somewhat fragmented nature. Quantifying the target label distribution and learning a classifier are actually very closely related objectives. The most classical approach for the construction of efficient classifiers on the target domain is based on weighted empirical risk minimisation, which itself requires the estimation of the shift between the source and target distributions. Thus, while we are here interested in estimating the target proportions α_i^* , many related works are interested in estimating the weights $w_i = \alpha_i^* / \beta_i^*$. Obtaining an estimator of those weights from an estimator of the target proportions α_i^* is straightforward: simply use the labels in the

source data to form a direct estimate of the source proportions and then consider the ratio. Conversely, starting from the weights estimator, it is possible to obtain an estimate of the weights α_i^*/β_i^* by multiplying it by an estimate of the source proportions. For this reason, there have been two different literature threads that address closely related problems but apparently grew independently: the Quantification Learning literature [19,13] and the Label Shift literature [18].

Most methods dealing with label shift are expressed as variants of the so-called *Classify & Count (CC)* technique proposed in the seminal works of Forman [15,16,17]. The idea is to fit a classifier on the source dataset, e.g. a random forest [30], an SVM [4], or a Nearest-Neighbour [5], and to estimate the target class distribution using the distribution of predictions. To account for the misclassification that the underlying classifier suffers on the target set due to the label shift, Forman [17] proposed the *Adjusted Classify & Count (ACC)* method, a modification to the standard Classify & Count which simply adjusts the estimate after computing the quantifier. This approach is also popular in the label shift correction literature. Based on the same principle, more recently the *Black-Box Shift Estimation (BBSE)* algorithm introduced by [27] used the confusion matrix of the classifier to adjust the predicted label distribution, while Azizzadeneheli et al. [3] proposed to regularise the re-weighting vector in BBSE in order to make the final estimated target distribution less sensitive to the quality of the estimated confusion matrix. Another technique, using an off-the-shelf classifier, is based on the maximum likelihood principle. The most popular version of this approach is probably the *Maximum Likelihood Label Shift (MLLS)* [2] strategy which is widely used in the label shift correction community. This technique is actually a variation of the original work of Saerens et al. [38] that uses an Expectation-Maximisation (EM) algorithm to efficiently correct for the shift in class proportions between source and target distributions given estimates from a predictive model, alternately updating the target class-prior and class-posterior probabilities from some initial estimates until convergence. In fact, [18] argues that only the choice of the calibration method differs between MLLS and BBSE and that both procedures actually solve the same optimisation objective, thus explaining that the empirical advantage of MLLS over BBSE is only due to coarse calibration of the latter.

Another completely different approach consists in viewing quantification as a statistical mixture problem. Since both source and target covariates marginal distributions can be written as mixtures $\mathbb{P} = \sum_{i=1}^c \beta_i^* \mathbb{P}_i$ and $\mathbb{Q} = \sum_{i=1}^c \alpha_i^* \mathbb{Q}_i$ respectively. Under the label shift assumption, the conditional distribution of the covariates given the label is the same for both source and target data ($\mathbb{P}_i = \mathbb{Q}_i$), and can be estimated using the empirical conditional distribution $\hat{\mathbb{P}}_i := \sum_{j/y_j=i} \delta_{x_j}(\cdot)/n_i$ based on the labeled source sample. Thus, the marginal covariates distribution \mathbb{Q} can be approximated by $\sum_{i=1}^c \alpha_i^* \hat{\mathbb{P}}_i$, and the goal of quantification can be seen as finding the mixture weights $(\alpha_i)_{i=1}^c$ such that the mixture $\sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i$ resembles the most the empirical target covariates marginal distribution $\hat{\mathbb{Q}} := \sum_{j=1}^m \delta_{x_{j+n}}(\cdot)/m$ based on the target dataset, with respect to some metric. Many different statistical divergences have been considered in the

literature, such as the Hellinger distance [20], the Wasserstein distance [6], the Pearson divergence [11], the Energy distance [26] or the Maximum Mean Discrepancy (MMD) [25,42]. The last two distances operate in a Reproducing Kernel Hilbert Space (RKHS) and adapt to the label shift problem the Kernel Mean Matching (KMM) approach [22], that minimises the RKHS distance between the kernel mean embeddings of the distributions. The use of kernel methods enables here to operate in an implicit high-dimensional space without requiring computing the coordinates of the data in that space, but rather by simply computing the inner products between the images of all pairs of data in the space. However, the applicability of these methods on a larger scale still remains an important limitation due to the computation of the kernel matrix. Consequently, distribution matching has sometimes been performed on a low-dimensional function of the covariates instead of a high-dimensional kernel embedding. For instance, the Distribution γ -Similarity framework (DyS) [29] exploits the histogram of a decision function obtained by a trained binary classifier and then minimises the Topsøe distance between the two histograms, while the HDx algorithm [20] uses the Hellinger distance between histograms of the Source and Target directly. *The general formulation we will adopt in this paper is in line with this approach of minimising some distance in a feature space between feature mappings of the distributions.*

1.2 Contributions of the paper

We introduce a general framework for label shift quantification, *Distribution Feature Matching* (DFM), that generalises existing methods such as Black-Box Shift Estimation (BBSE) [27], Kernel Mean Matching (KMM) [25,42] and its variant Energy Distance Matching [26].

The contributions of the paper are the following:

1. In Section 2, we propose a general framework of Label-shift quantification based on the minimisation of the distance between representations of distributions (in a Euclidean or Hilbert space) obtained by taking expectations of a feature mapping. We show that existing methods in the literature such as KMM or BBSE are particular instances of this general framework.
2. In Section 3, we provide a general statistical analysis of DFM under label shift. In particular, we show that our bound on the estimation error significantly improves the previous ones derived in the literature for both KMM and BBSE.
3. We also derive a novel analysis of DFM methods under departures of the label shift hypothesis, using a geometric decomposition of the problem and we show the implication of this analysis when we are in the contamination setting presented earlier. We thus show that certain DFM methods can exhibit robustness against particular types of perturbations or contaminations.
4. In Section 4, we support our theoretical results regarding the robustness of the different methods, with experiments on synthetic Gaussian mixtures and real cytometric datasets.

- Finally, we implement the KMM procedure in Python, using fast GPU-compatible code, using Random Fourier Features (RFF) to significantly reduce the computation burden while keeping the same type of theoretical guarantees. This implementation can still be used on GPU with limited memory.

2 Distribution Feature Matching

Let $\Phi : \mathcal{X} \rightarrow \mathcal{F}$ be a fixed feature mapping from \mathcal{X} into a Hilbert space \mathcal{F} (possibly $\mathcal{F} = \mathbb{R}^D$). We extend the mapping Φ to probability distributions on \mathcal{X} via taking expectation, i.e. $\Phi : \mathbb{P} \mapsto \Phi(\mathbb{P}) := \mathbb{E}_{X \sim \mathbb{P}}[\Phi(X)] \in \mathcal{F}$. Thus, it holds $\Phi(\hat{\mathbb{P}}_i) = n_i^{-1} \sum_{j \in [n]; y_j = i} \Phi(x_j)$, and similarly $\Phi(\hat{\mathbb{Q}}) = m^{-1} \sum_{j=n+1}^{n+m} \Phi(x_j)$.

We call *Distribution Feature Matching* (DFM) any estimation procedure that can be formulated as the minimiser of the following problem:

$$\hat{\alpha} = \arg \min_{\alpha \in \Delta^c} \left\| \sum_{i=1}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) - \Phi(\hat{\mathbb{Q}}) \right\|_{\mathcal{F}}^2 \quad (\mathcal{P})$$

where Δ^c is the $(c-1)$ -dimensional simplex.

In the contamination setting, we aim at finding the proportions of the non-noise classes of the target. As these proportions don't sum to one, the "hard" condition $\sum_i \alpha_i = 1$ is no longer needed. One way to overcome this is to introduce a fictitious "dummy" class in the source that formally has a vectorisation equal to 0 (note that adding a dummy class is a computational and theoretical convenience; we don't require to have a real distribution \mathbb{P}_0 that maps to 0 in the feature space for the results to hold). If we write $\Phi(\hat{\mathbb{P}}_0) := 0$ one can see that:

$$\arg \min_{\alpha \in \Delta^{c+1}} \left\| \sum_{i=0}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) - \Phi(\hat{\mathbb{Q}}) \right\|_{\mathcal{F}}^2 = \arg \min_{\alpha \in \text{int}(\Delta^c)} \left\| \sum_{i=1}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) - \Phi(\hat{\mathbb{Q}}) \right\|_{\mathcal{F}}^2, \quad (\mathcal{P}_2)$$

where $\text{int}(\Delta^c) := \{x \in \mathbb{R}^c : x \geq 0, \sum x_i \leq 1\}$. A procedure that solves \mathcal{P}_2 will be called *soft*-DFM. In Section 3, we will present theoretical results, in the classical Label Shift hypothesis (\mathcal{LS}), for DFM methods of the form (\mathcal{P}) under an identifiability and boundedness assumption.

In Section 3.2, we will show a general result for ("hard") DFM methods when the label shift hypothesis is not verified. As a corollary of these bounds we will directly obtain corresponding guarantees for the *soft*-DFM methods as well through the representation (\mathcal{P}_2) and formal inclusion of the dummy class.

In the remainder of this section, we will show the link between DFM and other classical Label Shift Quantification algorithms. However, any black-box feature mapping will be suitable for the results of Section 3.

2.1 Kernel Mean Matching

Iyer et al. [25] used Kernel Mean embedding (KME) as their mapping. We refer the reader to [31] for a survey on KME. We briefly recall that for any symmetric

and semi-definite positive kernel k defined on \mathcal{X} , one can associate a Hilbert space denoted \mathcal{H}_k , or simply \mathcal{H} when there is no ambiguity, and a "feature" mapping $\Phi : \mathcal{X} \rightarrow \mathcal{H}$ such that $\langle \Phi(x), \Phi(y) \rangle = k(x, y)$.

This mapping can be extended to the space of distributions by taking expectations as described above, which constitutes the principle of KME. We can compute scalar products between mappings using the formula $\langle \Phi(\mathbb{P}), \Phi(\mathbb{Q}) \rangle_{\mathcal{H}} = \mathbb{E}_{(X, Y) \sim \mathbb{P} \otimes \mathbb{Q}}[k(X, Y)]$, which provides a way to find an explicit solution of (\mathcal{P}) in practice. Then $D_{\Phi}(\mathbb{P}, \mathbb{Q}) = \|\Phi(\mathbb{P}) - \Phi(\mathbb{Q})\|_{\mathcal{H}}$ is a pseudo-distance on the space of measures on \mathcal{X} , called *Maximum Mean Discrepancy* (MMD) [21]. The specific use of a kernel feature mapping for class proportion estimation via (\mathcal{P}) has been called Kernel Mean Matching (KMM) by [25].

If $\mathcal{X} = \mathbb{R}^d$ with the usual Euclidean norm, Kawakubo et al. [26] proposed a particular case of KMM using the Energy kernel: $k(x, y) = \|x\| + \|y\| - \|x - y\|$, which is indeed a reproducing kernel [39].

2.2 BBSE as Distribution feature matching

Black-Box Shift Estimation is a method using the output of a black-box classifier to estimate the proportions in the target. To take into account the bias of the training data (i.e. the source) [27] used the confusion matrix.

To understand how Black-Box Shift Estimation can be cast as a Distribution Feature Matching procedure, we start from its original formulation as seeking the vector of proportions α that satisfies $Y = M\alpha$, where M is the estimated conditional confusion matrix defined as $M_{ij} = \frac{1}{n_i} \sum \mathbf{1}\{f(x_l) = i \text{ and } y_l = j\}$ and Y is the empirical mean of the observed outputs of the black-box classifier f on the target data, $Y_i = \frac{1}{n} \sum_l \mathbf{1}\{f(x_l) = i\}$. The BBSE estimate is then $\hat{\alpha} = M^{-1}Y$ (M is explicitly assumed invertible by [27]).

Proposition 1. *The BBSE estimator based on the black-box classifier f is the same as the solution of the DFM problem (\mathcal{P}) using the feature mapping $\Phi(x) = (\mathbf{1}\{f(x) = i\})_{i=1, \dots, c} \in \mathbb{R}^c$, where the positivity constraint on α is dropped.*

The proof is postponed to appendix A.4.

In the experiments to come, we will use BBSE+, our modified version of BBSE including the positivity constraint. The experimental results are slightly better for BBSE+. The reason is that in many cases, due to the presence of small classes in the source and in the target, BBSE returns negative proportions. When it does not output negative values, the two algorithms are the same.

This version of BBSE already existed in the literature, it has been used for text content analysis [24] and as a building block for classification, in a domain adaptation setting, more general than Label Shift [10].

3 Theoretical guarantees

We now provide statistical guarantees for DFM quantification procedures, All proofs can be found in appendix A.

We make the following identifiability hypothesis on the mapping Φ :

$$\sum_{i=1}^c \beta_i \Phi(\mathbb{P}_i) = 0 \iff \beta_i = 0 \quad \forall i = 1, \dots, c, \quad (\mathcal{A}_1)$$

and

$$\exists C > 0 : \quad \|\Phi(x)\|_{\mathcal{F}} \leq C \text{ for all } x. \quad (\mathcal{A}_2)$$

If we use KMM, the boundedness property is satisfied as soon as the kernel is bounded (e.g. Gaussian kernel, or any continuous kernel on a compact space). For BBSE, the boundedness is verified with $C = 1$.

Concerning Condition \mathcal{A}_1 , it is satisfied in the KMM case as long as the kernel is characteristic (e.g. Gaussian kernel) and the distributions \mathbb{P}_i are linearly independent (which is the minimal assumption for the class proportions to be identifiable). These assumption have been previously used by [25] for KMM. Similarly, for BBSE, [27] also assumed identifiability and required that the expected classifier outputs for each class be linearly independent.

We introduce the following notation and state our main theorem:

Definition 1. We denote $\hat{\mathbf{G}}$ the Gram matrix, resp. $\hat{\mathbf{M}}$ the centered Gram matrix of $\{\Phi(\hat{\mathbb{P}}_1), \dots, \Phi(\hat{\mathbb{P}}_c)\}$. That is, $\hat{\mathbf{G}}_{ij} = \langle \Phi(\hat{\mathbb{P}}_i), \Phi(\hat{\mathbb{P}}_j) \rangle$ and $\hat{\mathbf{M}}_{ij} = \langle \Phi(\hat{\mathbb{P}}_i) - \bar{\Phi}, \Phi(\hat{\mathbb{P}}_j) - \bar{\Phi} \rangle$ with $\bar{\Phi} = c^{-1} \sum_{k=1}^c \Phi(\hat{\mathbb{P}}_k)$. Furthermore, let Δ_{min} be the second smallest eigenvalue of $\hat{\mathbf{M}}$ and λ_{min} the smallest eigenvalue of $\hat{\mathbf{G}}$. In particular, it holds $\Delta_{min} \geq \lambda_{min}$.

Theorem 1. If the Label Shift hypothesis (\mathcal{LS}) holds, and if the mapping Φ verifies Assumptions (\mathcal{A}_1) and (\mathcal{A}_2) , then for any $\delta \in (0, 1)$, with probability greater than $1 - \delta$, the solution $\hat{\alpha}$ of (\mathcal{P}) satisfies:

$$\|\hat{\alpha} - \alpha^*\|_2 \leq \frac{2CR_c/\delta}{\sqrt{\Delta_{min}}} \left(\frac{\|w\|_2}{\sqrt{n}} + \frac{1}{\sqrt{m}} \right) \quad (1)$$

$$\leq \frac{2CR_c/\delta}{\sqrt{\Delta_{min}}} \left(\frac{1}{\sqrt{\min_i n_i}} + \frac{1}{\sqrt{m}} \right), \quad (2)$$

where $R_x = 2 + \sqrt{2 \log(2x)}$, $w_i = \frac{\alpha_i^*}{\beta_i}$.

The same result holds when replacing α^* by the (unobserved) vector of empirical proportions $\tilde{\alpha}$ in the target sample, both on the left-hand side and in the definition of w .

Under the same assumptions, the solution $\hat{\alpha}_{\text{soft}}$ of (\mathcal{P}_2) satisfies the same bounds with Δ_{min} replaced by λ_{min} .

3.1 Comparison to related literature

We compare our result to Theorem 1 of [25] and Theorem 3 of [27], which as we have mentioned earlier hold under the same assumptions as we make here.

Concerning KMM, a comparison between our inequality (2) and Theorem 1 in [25] shows that our bound is tighter than theirs, which is of leading order

$$\frac{c}{\sqrt{m}} + \sum_i \frac{1}{\sqrt{n_i}} \quad \text{vs ours in} \quad \frac{1}{\sqrt{m}} + \max_i \frac{1}{\sqrt{n_i}}$$

up to logarithmic factors. Thus, Theorem 1 improves upon the previous upper bound by a factor of c with respect to the term in m , and reduces the sum into a maximum regarding the number of instances per class n_i , which may also decrease the order of by factor c when the classes are evenly distributed in the source dataset. Furthermore, Inequality (1) even significantly improves over both Inequality (2) and Theorem 1 in [25]. Indeed, in situations where one of the classes i on the source domain is rare, then the rate $\max_i n_i^{-1/2}$ in Inequality (2) explodes, which is not the case of the rate $\|w\|_2/\sqrt{n}$ in Inequality (1) when the source and target proportions are similar, as the weight vector w reflects the similarity between the source and target distributions. Note that we use the theoretical proportions α^* for the target in the definition of w as the empirical ones are unknown here. Hence, our theorem significantly improves the existing bound for KMM established by [25]. Similarly, our bound (1) applied to BBSE also improves Theorem 3 in [27]. In particular, when both inequalities are formulated with the same probability level (e.g. $1 - \delta$), our bound for BBSE is tighter by a factor \sqrt{c} w.r.t. the term in m than the guarantee provided by [27]. Note however that contrary to our result and to Theorem 1 in [25], the bound of [27] does not involve any empirical quantity that can be computed using the source dataset.

Another key component of the bounds is the second smallest eigenvalue Δ_{\min} of the centered Gram matrix, which replaces the minimum singular value of the Gram matrix in the case of KMM (see Theorem 1 in [25]) and the smallest eigenvalue of the confusion matrix divided by the infinite norm of the source proportions in the case of BBSE (see Inequality (3) of Theorem 3 in [27]), and improves upon both of them.

This improvement is particularly important when the two source classes are unbalanced. For instance, in a two-class setting with a black-box classifier feature map and $\beta^* = (p, 1 - p)$, the theoretical version of Δ_{\min} is equal to 1 when the classifier is perfect and replaces the factor that would be $\min\left(\frac{p}{1-p}, \frac{1-p}{p}\right) < 1$ in the bound of [27]. When the classifier is not perfect but both classes share the same classification accuracy $a \in (1/2, 1)$, then $\Delta_{\min} = 2a - 1$, which strictly improves the factor of [27] except when both classes are equally balanced, in which case both quantities are equal.

To fully understand the nature of Δ_{\min} , we can interpret DFM as the projection of the target feature mapping onto the convex hull of the source feature mappings. Our estimation is then simply the barycentric coordinate of the projection, as formalised in (QP), and Δ_{\min} is a geometric property of that convex hull. It represents the ease with which our mapping can distinguish between one class and any mixture of the other classes. For instance, for two classes and two embeddings (Φ_1, Φ_2) , one can show that $\Delta_{\min} = \frac{1}{2}\|\Phi_1 - \Phi_2\|_{\mathcal{F}}^2$. From a geometric

point of view, it is clear that the larger the convex hull (i.e. the line connecting the two features in situations where there are two classes), the less the barycentric coordinate will be affected by a small perturbation of the weights. From a statistical point of view, if our mixture is composed of two very different distributions, it will be intuitively easier to distinguish them in a new sample. The quantity Δ_{\min} (which we recall is empirical) also provides a natural criterion for the choice of the feature map hyperparameter (i.e. choice of the kernel in KMM), as the dependence in our bound only appears in Δ_{\min} which can then be maximised.

3.2 Robustness to contamination

We now introduce a novel theoretical analysis of the robustness of the method with respect to the contaminated label shift model (assumption \mathcal{CLS}). First, let us obtain a general result when Label Shift is not verified.

A naive approach would simply include the bias term $\|\Phi(\mathbb{P}_i) - \Phi(\mathbb{Q}_i)\|_{\mathcal{F}}$ in the bound. We put into light the robustness of the procedure with respect to certain types of perturbation.

Theorem 2. *Denote V the affine span of the vectors $\Phi(\mathbb{P}_i)$ and \mathcal{C} the convex hull of those same vectors. Denote Π_V and $\Pi_{\mathcal{C}}$ the orthogonal resp. convex projection onto V and \mathcal{C} .*

Suppose the same assumptions as in Theorem 1 hold, except for the exact label shift assumption \mathcal{LS} . Then, with probability greater than $1 - \delta$:

$$\|\hat{\alpha} - \alpha^*\|_2 \leq \frac{1}{\sqrt{\Delta_{\min}}} \left(3\epsilon_n + \epsilon_m + \sqrt{2\epsilon_n} B^\perp + B^\parallel \right), \quad (3)$$

with:

$$\epsilon_n = C \frac{R_{c/\delta}}{\sqrt{\min_i n_i}}; \quad \epsilon_m = C \frac{R_{1/\delta}}{\sqrt{m}}; \quad (4)$$

$$B^\perp = B^\perp(\mathbb{P}, \mathbb{Q}) = \sqrt{\|\Phi(\mathbb{Q}) - \Pi_{\mathcal{C}}(\Phi(\mathbb{Q}))\|_{\mathcal{F}}};$$

$$B^\parallel = B^\parallel(\mathbb{P}, \mathbb{Q}) = \max_i \|\Phi(\mathbb{P}_i) - \Pi_V(\Phi(\mathbb{Q}_i))\|_{\mathcal{F}}.$$

Observe that the bound (3) shows robustness of a DFM procedure against perturbations B^\perp that are "orthogonal" to the source space V in feature space. In particular, *consistency* (i.e. convergence of the bound to 0 as the sample sizes grow to infinity) is still granted if $\mathbb{Q}_i \neq \mathbb{P}_i$ but $\Pi_V(\mathbb{Q}_i) = \mathbb{P}_i$. Which type of perturbation of the distributions will result in (close to) orthogonal shifts in feature space very much depends on the feature mapping used. For BBSE, the feature space is of the same dimension as the number of sources, thus under condition (\mathcal{A}_1) , V will coincide with E_1 , the affine space of vectors summing to one. Since any distribution will be also mapped to E_1 , the orthogonal component will always be 0. Thus, we expect no particular robustness property for BBSE

methods. For more general feature maps, such as kernel methods or any other vectorisations, this orthogonality property remains to be investigated in general, but we will exhibit below a favorable scenario for KMM in the *contaminated label shift* setting \mathcal{CLS} .

We first state a corollary in the \mathcal{CLS} scenario. To do so, we recall that, in this case, we use the *soft*-DFM procedure \mathcal{P}_2 . We are now in a particular case, where the only difference between source and target is that the unknown noise class \mathbb{Q}_0 is formally replaced by the dummy class having feature map equal to 0 in the source. Introduce $\bar{V} := \text{Span}\{\Phi(\mathbb{P}_i), i \in [c]\}$ (i.e. vector span rather than affine span for V previously) and let $\Pi_{\bar{V}}$ be the orthogonal projection on \bar{V} .

Corollary 1. *Denote by $\hat{\alpha}_{\text{soft}}$ the minimiser of the soft-DFM problem \mathcal{P}_2 . Assume the contaminated Label Shift hypothesis (\mathcal{CLS}) holds, and if the mapping Φ verifies Assumptions (\mathcal{A}_1) and (\mathcal{A}_2) . Then, with probability greater than $1 - \delta$:*

$$\|\hat{\alpha}_{\text{soft}} - \alpha^*\|_2 \leq \frac{1}{\sqrt{\lambda_{\min}}} \left(3\epsilon_n + \epsilon_m + \sqrt{2\alpha_0 \epsilon_n \|\Phi(\mathbb{Q}_0)\|_{\mathcal{F}}} + \|\Pi_{\bar{V}}(\Phi(\mathbb{Q}_0))\|_{\mathcal{F}} \right),$$

with ϵ_n, ϵ_m defined as in (4).

In the particular case of KMM with a translation-invariant kernel $k(x, y) = \varphi(x - y)$, for any distributions \mathbb{P}, \mathbb{P}' it holds $\langle \Phi(\mathbb{P}), \Phi(\mathbb{P}') \rangle = \mathbb{E}_{(X, Y) \sim \mathbb{P} \otimes \mathbb{P}'}[\varphi(X - Y)]$. Thus, if φ is rapidly decaying with distance (e.g. Gaussian kernel), the feature mappings $\Phi(\mathbb{P})$ and $\Phi(\mathbb{P}')$ will be close to orthogonal (have a scalar product close to 0) whenever the distributions \mathbb{P}, \mathbb{P}' are well-separated. From this analysis, we anticipate that KMM with a Gaussian kernel will be robust against contaminations distributions \mathbb{Q}_0 whose main mass is far away from the source distributions, since its representation $\Phi(\mathbb{Q}_0)$ will then be close to orthogonal to \bar{V} in feature space.

4 Algorithm and applications

In this section, we will apply four methods on both synthetic and real datasets. We choose to test three soft-DFM methods: KMM using the Energy Kernel [26], our modified version of BBSE [27] and KMM using the Gaussian kernel [25] enhanced with Random Fourier Features, that we present in the next section. To show the benefit of the soft version, we also compare with one *hard*-DFM method: KMM enhanced with Random Fourier Features.

The main objective of the experiments is, in view of our theoretical results of Section 3.2 and particularly Corollary 1, to test robustness properties of several DFM methods against contamination of the the target dataset by different types of noise. Moreover, we want to check if the *soft* version presented in Section 2 leads to improved results in some cases, and will not hurt the results in the others.

All the code and datasets are publicly available [12]. All the computations were done on a computer equipped with a NVIDIA RTX A2000 Laptop.

4.1 Optimisation problem

Whatever the chosen mapping, solving (\mathcal{P}) amounts to solving a Quadratic Programming (QP) in dimension c . Indeed, we can rewrite the problem as:

$$\begin{aligned} & \text{minimise } \frac{1}{2}\alpha^T \hat{\mathbf{G}}\alpha + q^T \alpha & (\text{QP}) \\ & \text{subject to } \alpha \succeq 0_c \text{ and } \mathbf{1}_c^T \alpha = 1, \end{aligned}$$

with $q = \left(\langle \Phi(\hat{\mathbb{P}}_i), \Phi(\hat{\mathbb{Q}}) \rangle \right)_{i=1}^c$. This is a c -dimensional QP problem, which can be solved efficiently.

The computational bottleneck is the computation of the Gram matrix $\hat{\mathbf{G}}$ and the vector q . Using KMM directly leads to a complexity of $O(n(n+m))$, as computing q requires evaluating the kernel for every pair of points from the source and the target and computing $\hat{\mathbf{G}}$ requires evaluating the kernel for every pair of points between the source classes. Moreover, one needs to have permanent access to the source distributions, as computing q requires both the source and target raw dataset.

Due to such issues, kernel matrix approximations are often used in order to reduce the computational cost of kernel methods [8,36]. In our case we use the well-established principle of Random Fourier Features (RFF) approximation [35]. RFF allows to obtain an approximation of a translation invariant kernel k , i.e. $k(x, y) = \varphi(x - y)$, of the form: $\tilde{k}(x, y) = \tilde{\Phi}(x)^T \tilde{\Phi}(y)$, with $\tilde{\Phi}(x) \in \mathbb{R}^D$, which is itself a positive semi-definite kernel. For a theoretical analysis of the uniform approximation quality of k by \tilde{k} , see e.g. [37,41]. See the appendix C for more detail on Random Fourier Features.

Relying on RFF with D Fourier features induces a complexity of $O(D(n+m))$ since we only have to compute $\tilde{\Phi}(\hat{\mathbb{P}}_i)$ and $\tilde{\Phi}(\hat{\mathbb{Q}})$. Computing $\tilde{\Phi}(\hat{\mathbb{P}})$ reduces to a matrix multiplication, for which GPU are well suited. To deal with memory overflow on GPU, we rely on the Python package *PyKeops* [9]. With this implementation can solve (\mathcal{P}) for high-dimensional data with a very large number of points in sub-second times. For example, for two datasets containing 5×10^6 points in dimension 5, (\mathcal{P}) is solved in less than a second, while almost 2 minutes are needed when we use the exact KMM.

In the experiments, we will design by Random Fourier Features Matching (RFFM) the DFM method that uses $\tilde{\Phi}$ as a feature mapping. RFFM can be used with any translation invariant kernel, but we choose to stick to the classical Gaussian kernel: $k(x, y) = \exp(-\|x - y\|^2 / (2\sigma^2))$ where the parameter σ is optimised using the criterion derived from (1). Note that RFFM is only used as a way to speed up the computation, and hence we would obtain similar results with a classical KMM using the Gaussian kernel.

4.2 Experiments

We want to test the robustness of the DFM methods in the contaminated scenario \mathcal{CLS} . We will compare 4 methods: RFFM, softRFFM, softEnergy and

softBBSE+. RFFM is the method introduced in the previous section while soft-RFFM is RFFM when we use the soft procedure introduced in Section 2. Soft-Energy is the kernel mean matching method when we used the Energy kernel [26] and softBBSE+ is our version of BBSE [27].

We will test the methods on both synthetic data and real datasets.

Gaussian Mixture

In this setting, the source is a list of c Gaussian distributions: $\mathbb{P}_1, \dots, \mathbb{P}_c$. Our objective is to estimate α^* for different values of the contamination level $\epsilon = \alpha_0^*$ ranging from 0 to 0.3 and different kinds of noise distributions \mathbb{Q}_0 . We will test three kinds of noise (see Figure 3):

1. \mathbb{Q}_0 is uniformly distributed over the data range ("Background noise").
2. \mathbb{Q}_0 is Gaussian with a mean distant from the other means.
3. \mathbb{Q}_0 is Gaussian with a similar mean to the others.

Throughout the experiments, we fix the number of classes in the source to $c = 5$ and the number of non-contaminated points to 10000. The dimension varies from 2 to 10. For each contamination level ϵ and each possible dimension, we perform 20 repetitions with different Gaussian distributions. Results for the three experiments can be found in Figure 1.

In the absence of noise contamination in the target, all methods give excellent results because the source distributions are easy to distinguish. Obviously, the results deteriorate as the contamination level increases. All methods still perform well against background noise (the loss is around 0.1 in the worst case), with softRFFM being significantly better than the others. In the same fashion, when we add a Gaussian far away from the other distributions, softRFFM significantly outperforms the others. As discussed following Corollary 1, this is because when contamination puts mostly mass far from the other classes, the Gaussian KME of the noise distribution will be close to orthogonal to the source KMEs. This property does not hold when a class is added close to the others and thus can be more easily confounded. Thus, the results align well with our theoretical analysis.

While Theorem 2 and Corollary 1 hold for the Energy kernel as well (assuming bounded data) or BBSE+, we don't observe robustness against noise. Again, this is in line with the theoretical study for BBSE+ for which we expected no robustness. Concerning the Energy kernel, we surmise that the lack of robustness comes from the fact that $k(x, y)$ can take large values even if $\|x - y\|$ is large, hence near-orthogonality of the noise distribution to the source does not hold in the corresponding feature space, in contrast to the Gaussian kernel.

Cytometry dataset

We test the robustness of our methods on the T-cell panel of the Human Immunology Project Consortium (HIPC) [7,14]. HIPC is composed of 63 samples. Seven laboratories analysed 3 replicates from 3 different patients. The number of

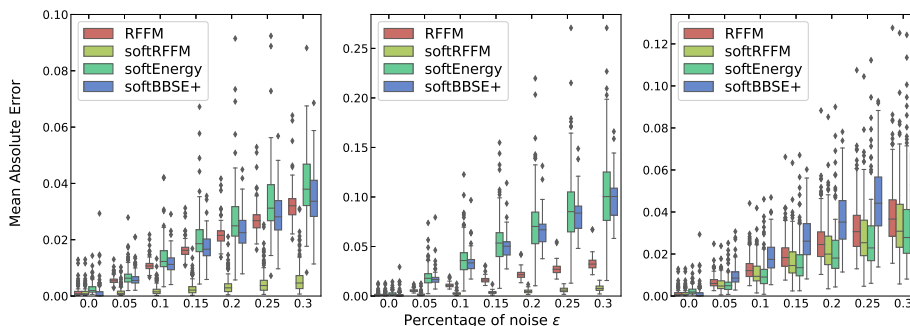


Fig. 1. Robustness of the algorithms to three types of noise. (Lower is better.) Left: background noise; middle: noise is a new class far from the others; right: noise is a new class in the middle of the others.

measurements in the samples range from 10^4 to 10^5 . The samples were manually separated into 10 categories using 7 markers.

To put ourselves in the \mathcal{CLS} setting, we choose to remove one of the class of the source, so that this class becomes the noise in the target. In detail, each of the patient replicates are joined into a single patient sample, resulting in 3 patient joined samples for each of the seven laboratories. We take each sample successively as source, and patients samples of the same laboratory as target.³ We first perform the experiment with all cell classes: in that case we can assume that we are in the vanilla Label Shift \mathcal{LS} setting. Then, we repeat the experiment 10 times, removing each time a class from the source but not from the target, thus playing the role of contaminant. The results can be found in Figure 2

We reach the same conclusion on the robustness of softRFFM compared to RFFM, softBBSE+ and softEnergy when we apply them on a noisy real-world dataset. This advantage of robustness is all the more significant as the proportion of noise is high.

5 Conclusion

We introduced Distribution Feature Matching (DFM) as a general approach for class proportion estimation (also known at label shift estimation or quantification learning), recovering methods from previous literature as special cases. We also proposed the use of Random Fourier Features to speed up the computation of kernel-based approaches and obtain an explicit finite-dimensional vectorization (or "sketch") of the distributions.

³ There can be large variability between samples coming from different laboratories, while there is homogeneity within each lab. The label shift hypothesis is therefore reasonable when keeping source and target from the same lab.

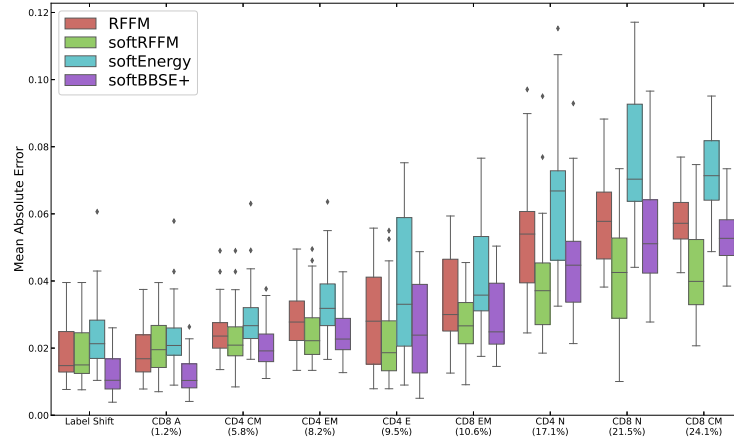


Fig. 2. Each column represents the error — computed using the ℓ_2 norm between the true proportions and the estimated proportions — obtained when some class is absent from the source but present in the target distribution. The first column gives the results when no class is discarded. The class are sorted according to the average proportions they represent in the samples (x labels mention class held out from the source and its proportion)

We provided a general theoretical analysis of DFM, improving over previously known bounds derived for specific instantiations only. Furthermore, we analysed theoretically the behavior of DFM under departures from the label shift hypothesis, a situation not studied in earlier works, and put into light a robustness against certain types of perturbations, depending on the feature mapping used. The theoretical analysis suggested a better robustness property of RFF approaches based on a rapidly decaying translation-invariant kernel, and this could be confirmed through numerical experiments on synthetic and real data.

Recent works [10,42] considered a more general situation beyond label shift. In the *Generalized Label Shift* model of [10], the condition $(\mathcal{L}\mathcal{S})$ does not hold in input space, but only after transformation through a suitable feature mapping Φ , and these authors proposed an algorithm alternating between class proportion estimation using BBSE+ and feature learning (using a separate domain adaptation algorithm, suitably adapted to handle label proportion shift). We believe that using one of the proposed methods in the present paper could be used fruitfully in such a framework to replace the BBSE+ module.

Acknowledgements

B. Dussap was supported by the program Paris Region Ph.D. of DIM Mathinnov. G. Blanchard acknowledges support of the ANR under ANR-19-CHIA-0021-01 “BiSCottE”, IDEX REC-2019-044, and of the DFG under SFB1294 - 318763901.

Ethical statement

Label shift quantification has uses in a number of application domains; the results in this paper are chiefly oriented towards theory and general methodology so that we don't discuss an application in particular. We only mention that the flow cytometry data used for proof-of-concept experiments is publicly available from a reputable scientific consortium and has been up to our knowledge collected following all established ethical standards.

The original *Classify & Count* [15] method for label shift quantification is known to inherit the potential biases of the classification method it is based on (i.e. the misclassification errors can be very unevenly distributed across classes and "favor" majority classes). The *Adjusted Classify & Count (ACC)* approach and related methods [17,27] aim at rectifying this bias. In the present paper, we aim at going one step further and analyze certain robustness properties of the proposed label shift quantification methods, and introduce the contaminated label shift (*CLS*) setting with the goal of investigating trustworthiness of such methods under mild violations of the standard Label Shift model. Certainly the robustness property is desirable for improved reliability in practice, but does not mean immunity against biases; additionally, the user should always be wary of stronger model violations between reference and test data, in particular class-conditional distribution shifts. We therefore recommend established good practice of regularly checking on control data possible biases or drifts from the model, in particular for sensitive applications.

References

1. Alber, Y.I., Notik, A.: On some estimates for projection operator in Banach space. arXiv preprint [funct-an/9311003](https://arxiv.org/abs/funct-an/9311003) (1993)
2. Alexandari, A., Kundaje, A., Shrikumar, A.: Maximum likelihood with bias-corrected calibration is hard-to-beat at label shift adaptation. In: International Conference on Machine Learning. pp. 222–232. PMLR (2020)
3. Azizzadenesheli, K., Liu, A., Yang, F., Anandkumar, A.: Regularized learning for domain adaptation under label shifts. arXiv preprint [arXiv:1903.09734](https://arxiv.org/abs/1903.09734) (2019)
4. Barranquero, J., Díez, J., del Coz, J.J.: Quantification-oriented learning based on reliable classifiers. *Pattern Recognition* **48**(2), 591–604 (2015)
5. Barranquero, J., González, P., Díez, J., Del Coz, J.J.: On the study of nearest neighbor algorithms for prevalence estimation in binary problems. *Pattern Recognition* **46**(2), 472–482 (2013)
6. Bigot, J., Freulon, P., Hejblum, B.P., Leclaire, A.: On the potential benefits of entropic regularization for smoothing Wasserstein estimators. arXiv preprint [arXiv:2210.06934](https://arxiv.org/abs/2210.06934) (2022)
7. Brusci, V., Gottardo, R., Kleinstein, S.H., Davis, M.M.: Computational resources for high-dimensional immune analysis from the Human Immunology Project Consortium. *Nature Biotechnology* **32**, 146–148 (2014)
8. Camoriano, R., Angles, T., Rudi, A., Rosasco, L.: Nytro: When subsampling meets early stopping. In: Artificial Intelligence and Statistics. pp. 1403–1411. PMLR (2016)
9. Charlier, B., Feydy, J., Glaunès, J.A., Collin, F.D., Durif, G.: Kernel operations on the GPU, with autodiff, without memory overflows. *Journal of Machine Learning Research* **22**(74), 1–6 (2021), <https://www.kernel-operations.io/keops/index.html>
10. Tachet des Combes, R., Zhao, H., Wang, Y.X., Gordon, G.J.: Domain adaptation with conditional distribution matching and generalized label shift. *Advances in Neural Information Processing Systems* **33**, 19276–19289 (2020)
11. Du Plessis, M.C., Sugiyama, M.: Semi-supervised learning of class balance under class-prior change by distribution matching. *Neural Networks* **50**, 110–119 (2014)
12. Dussap, B.: Distribution Feature Matching for Label Shift (2023), <https://plmlab.math.cnrs.fr/dussap/Label-shift-DFM>, (Software)
13. Esuli, A., Fabris, A., Moreo, A., Sebastiani, F.: Learning to quantify (2023)
14. Finak, G., Langweiler, M., Jaimes, M., Malek, M., Taghiyar, J., Korin, Y., Rad-dassi, K., Devine, L., Obermoser, G., Pekalski, M.L., et al.: Standardizing flow cytometry immunophenotyping analysis from the human immunophenotyping consortium. *Scientific reports* **6**(1), 1–11 (2016)
15. Forman, G.: Counting positives accurately despite inaccurate classification. In: European conference on machine learning. pp. 564–575. Springer (2005)
16. Forman, G.: Quantifying trends accurately despite classifier error and class imbalance. In: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 157–166 (2006)
17. Forman, G.: Quantifying counts and costs via classification. *Data Mining and Knowledge Discovery* **17**(2), 164–206 (2008)
18. Garg, S., Wu, Y., Balakrishnan, S., Lipton, Z.C.: A unified view of label shift estimation. arXiv preprint [arXiv:2003.07554](https://arxiv.org/abs/2003.07554) (2020)
19. González, P., Castaño, A., Chawla, N.V., Coz, J.J.D.: A review on quantification learning. *ACM Computing Surveys (CSUR)* **50**(5), 1–40 (2017)

20. González-Castro, V., Alaiz-Rodríguez, R., Alegre, E.: Class distribution estimation based on the hellinger distance. *Information Sciences* **218**, 146–164 (2013)
21. Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *The Journal of Machine Learning Research* **13**(1), 723–773 (2012)
22. Gretton, A., Smola, A., Huang, J., Schmittfull, M., Borgwardt, K., Schölkopf, B.: Covariate shift by kernel mean matching. *Dataset shift in machine learning* **3**(4), 5 (2009)
23. Gundersen, G.: Random Fourier Features. <https://gregorygundersen.com/blog/2019/12/23/random-fourier-fe> (2019)
24. Hopkins, D.J., King, G.: A method of automated nonparametric content analysis for social science. *American Journal of Political Science* **54**(1), 229–247 (2010)
25. Iyer, A., Nath, S., Sarawagi, S.: Maximum mean discrepancy for class ratio estimation: Convergence bounds and kernel selection. In: *International Conference on Machine Learning*. pp. 530–538. PMLR (2014)
26. Kawakubo, H., Du Plessis, M.C., Sugiyama, M.: Computationally efficient class-prior estimation under class balance change using energy distance. *IEICE Transactions on Information and Systems* **99**(1), 176–186 (2016)
27. Lipton, Z., Wang, Y.X., Smola, A.: Detecting and correcting for label shift with black box predictors. In: *International conference on machine learning*. pp. 3122–3130. PMLR (2018)
28. Lopez-Paz, D., Muandet, K., Schölkopf, B., Tolstikhin, I.: Towards a learning theory of cause-effect inference. In: *International Conference on Machine Learning*. pp. 1452–1461. PMLR (2015)
29. Maletzke, A., dos Reis, D., Cherman, E., Batista, G.: Dys: a framework for mixture models in quantification. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 4552–4560 (2019)
30. Milli, L., Monreale, A., Rossetti, G., Giannotti, F., Pedreschi, D., Sebastiani, F.: Quantification trees. In: *2013 IEEE 13th International Conference on Data Mining*. pp. 528–536. IEEE (2013)
31. Muandet, K., Fukumizu, K., Sriperumbudur, B., Schölkopf, B., et al.: Kernel mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine Learning* **10**(1-2), 1–141 (2017)
32. Patel, V.M., Gopalan, R., Li, R., Chellappa, R.: Visual domain adaptation: A survey of recent advances. *IEEE signal processing magazine* **32**(3), 53–69 (2015)
33. Pinelis, I.: An approach to inequalities for the distributions of infinite-dimensional martingales. In: *Probability in Banach Spaces, 8: Proceedings of the Eighth International Conference*. pp. 128–134. Springer (1992)
34. Quinonero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: *Dataset shift in machine learning*. Mit Press (2008)
35. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. *Advances in neural information processing systems* **20** (2007)
36. Rudi, A., Camoriano, R., Rosasco, L.: Less is more: Nyström computational regularization. *Advances in Neural Information Processing Systems* **28** (2015)
37. Rudi, A., Rosasco, L.: Generalization properties of learning with random features. *Advances in neural information processing systems* **30** (2017)
38. Saerens, M., Latinne, P., Decaestecker, C.: Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation* **14**(1), 21–41 (2002)
39. Sejdinovic, D., Sriperumbudur, B., Gretton, A., Fukumizu, K.: Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *The annals of statistics* pp. 2263–2291 (2013)

40. Smola, A., Gretton, A., Song, L., Schölkopf, B.: A hilbert space embedding for distributions. In: International Conference on Algorithmic Learning Theory. pp. 13–31. Springer (2007)
41. Sutherland, D.J., Schneider, J.: On the error of random Fourier features. In: Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence. pp. 862–871 (2015)
42. Zhang, K., Schölkopf, B., Muandet, K., Wang, Z.: Domain adaptation under target and conditional shift. In: International conference on machine learning. pp. 819–827. PMLR (2013)

A Proofs

A.1 Proof of Theorem 1.

To prove Theorem 1, we will require a result on the approximation error of the vector $\Phi(\mathbb{P})$ by an n -sample $\Phi(\hat{\mathbb{P}})$. For this, we use the vector Hoeffding's inequality (for a Euclidean or Hilbert norm), a result which has appeared in different versions in the literature (one of the earliest version seems to be from [33]); see also e.g. [28] or [40] in the KME setting. While this result is not new, we give here a short self-contained proof for completeness.

Theorem A.1. *Let Z_1, \dots, Z_n be independent (not necessarily identically distributed) random variables taking values in a Hilbert space \mathcal{H} (possibly $\mathcal{H} = \mathbb{R}^D$), such that $\forall i \in [n] : \|Z_i\| \leq C$ a.s. Then with probability greater than $1 - \delta$:*

$$\left\| \frac{1}{n} \sum_{i=1}^n (Z_i - \mathbb{E}[Z_i]) \right\| \leq C \frac{(2 + \sqrt{2 \log(1/\delta)})}{\sqrt{n}},$$

Proof. Since $\|Z_i\| \leq C$ we can assume without loss of generality that the random variables Z_i take values in the ball $\mathcal{B}_C := \{z \in \mathcal{H} : \|z\| \leq C\}$. Define the function $F : (\mathcal{B}_C)^n \rightarrow \mathbb{R}$ as

$$F(z_1, \dots, z_n) := \left\| \frac{1}{n} \sum_{i=1}^n (z_i - \mathbb{E}[Z_i]) \right\|.$$

Straightforward computations show that the function F satisfies the bounded difference condition. Namely, let us fix all the values z_1, \dots, z_n in \mathcal{B}_C except for the z_j which will be set to \bar{z}_j ; then

$$|F(z_1, \dots, z_n) - F(z_1, \dots, \bar{z}_j, \dots, z_n)| = \frac{1}{n} \|z_j - \bar{z}_j\| \leq \frac{2C}{n}.$$

Using McDiarmid's inequality, since Z_i are realisations of independent random variables taking values in \mathcal{B}_C , it holds with probability greater than $1 - \delta$:

$$\left\| \frac{1}{n} \sum_{i=1}^n (Z_i - \mathbb{E}[Z_i]) \right\| \leq \mathbb{E} \left[\left\| \frac{1}{n} \sum_{i=1}^n (Z_i - \mathbb{E}[Z_i]) \right\| \right] + C \sqrt{\frac{2 \log(1/\delta)}{n}}.$$

Let us write $\tilde{Z}_i = Z_i - \mathbb{E}[Z_i]$; \tilde{Z}_i are also independent random variables, bounded in norm by $2C$, and are centered. We only have to bound the expectation in the right-hand side above.

By Jensen's inequality,

$$\begin{aligned} \mathbb{E} \left[\left\| \frac{1}{n} \sum \tilde{Z}_i \right\| \right] &\leq \sqrt{\mathbb{E} \left[\left\| \frac{1}{n} \sum \tilde{Z}_i \right\|^2 \right]} \\ &= \frac{1}{n} \left(\sum_{i,j=1}^n \mathbb{E}[\langle \tilde{Z}_i, \tilde{Z}_j \rangle] \right)^{\frac{1}{2}} \\ &= \frac{1}{n} \left(\sum_{i=1}^n \mathbb{E}[\|\tilde{Z}_i\|^2] \right)^{\frac{1}{2}} \leq \frac{2C}{\sqrt{n}}, \end{aligned}$$

where we have used that since for $i \neq j$ the variables \tilde{Z}_i, \tilde{Z}_j are independent and centered it holds $\mathbb{E}[\langle \tilde{Z}_i, \tilde{Z}_j \rangle] = 0$. \square

We now turn to the proof of Theorem 1.

Proof. Let $\alpha \in \Delta^c$ be fixed. Later we will consider the choices $\alpha = \alpha^*$ or $\alpha = \tilde{\alpha}$ (the population resp. empirical class proportions in the target domain), and will specify this when needed, but a large part of the argument holds for any α .

Given a feature map Φ , let us use the notation introduced in definition 1. Let us note D_Φ the function defined by $D_\Phi(\mathbb{P}_1, \mathbb{P}_2) = \|\Phi(\mathbb{P}_1) - \Phi(\mathbb{P}_2)\|$. Note that D_Φ will be a distance if and only if the mapping Φ is injective and will be a pseudo-distance otherwise.

It holds

$$\begin{aligned} D_\Phi \left(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i \right)^2 &= \left\| \sum_{i=1}^c \hat{\alpha}_i \Phi(\hat{\mathbb{P}}_i) - \sum_{i=1}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) \right\|^2 \\ &= \left\| \sum_{i=1}^c (\hat{\alpha}_i - \alpha_i) \Phi(\hat{\mathbb{P}}_i) \right\|^2 \\ &= (\hat{\alpha} - \alpha)^T \hat{\mathbf{G}} (\hat{\alpha} - \alpha) \\ &\geq \left(\min_{\substack{\|u\|=1 \\ \mathbf{1}^T u = 0}} u^T \hat{\mathbf{G}} u \right) \|\hat{\alpha} - \alpha\|^2 \\ &\stackrel{(\dagger)}{=} \lambda_{\min}(\hat{\mathbf{G}}) \|\hat{\alpha} - \alpha\|^2, \end{aligned}$$

with equality (\dagger) proven in Theorem B.1.

Thus in order to bound $\|\hat{\alpha} - \alpha\|$ we have to upper-bound: $D_\Phi(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i)$. By the triangle inequality, this is upper-bounded by $D_\Phi(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \hat{\mathbb{Q}}) + D_\Phi(\hat{\mathbb{Q}}, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i)$. By definition of $\hat{\alpha}$, we have $D_\Phi(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \hat{\mathbb{Q}}) \leq D_\Phi(\sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i, \hat{\mathbb{Q}})$. Hence, we can upper-bound the quantity by $2D_\Phi(\sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i, \hat{\mathbb{Q}})$.

Using the triangle inequality once again, we can upper bound the previous quantity by $D_{\Phi} \left(\sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \mathbb{P}_i \right) + D_{\Phi} \left(\sum_{i=1}^c \alpha_i \mathbb{P}_i, \hat{\mathbb{Q}} \right)$.

Using Theorem A.1 and the union bound, it holds with probability greater than $1 - \delta/2$:

$$\begin{aligned} D_{\Phi} \left(\sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \mathbb{P}_i \right) &\leq \sum_{i=1}^c \alpha_i D_{\Phi}(\hat{\mathbb{P}}_i, \mathbb{P}_i) \\ &\leq \sum_{i=1}^c \alpha_i C \frac{2 + \sqrt{2 \log 2c/\delta}}{\sqrt{n_i}} \\ &= CR_{c/\delta} \sum_{i=1}^c \frac{\alpha_i}{\sqrt{n_i}}, \end{aligned} \quad (5)$$

with $R_x = 2 + \sqrt{2 \log 2x}$.

Since $n_i = n\beta_i$, it holds $\sum_{i=1}^c \frac{\alpha_i}{\sqrt{n_i}} = \frac{1}{\sqrt{n}} \sum_{i=1}^c \frac{\alpha_i}{\sqrt{\beta_i}}$. If we write $w_i = \frac{\alpha_i}{\beta_i}$, then by Hölder's inequality:

$$\sum_{i=1}^c \frac{\alpha_i}{\sqrt{\beta_i}} = \sum_{i=1}^c \left(w_i \sqrt{\beta_i} \right) \leq \|w\|_2 \underbrace{\sqrt{\sum_{i=1}^c \beta_i}}_{=1}. \quad (6)$$

We finally turn to bounding the term $D_{\Phi} \left(\sum_{i=1}^c \alpha_i \mathbb{P}_i, \hat{\mathbb{Q}} \right)$. This is the only point of the proof where we need to use the label shift assumption and specify α . First, under (\mathcal{LS}) this is equal to $D_{\Phi} \left(\sum_{i=1}^c \alpha_i \mathbb{Q}_i, \hat{\mathbb{Q}} \right)$. We now distinguish between two possibilities:

- If $\alpha = \alpha^*$, then $\sum_{i=1}^c \alpha_i^* \mathbb{Q}_i = \mathbb{Q}$, so that using Theorem A.1, it holds with probability greater than $1 - \delta/2$:

$$D_{\Phi} \left(\sum_{i=1}^c \alpha_i^* \mathbb{Q}_i, \hat{\mathbb{Q}} \right) = D_{\Phi} \left(\mathbb{Q}, \hat{\mathbb{Q}} \right) \leq C \frac{R_{1/\delta}}{\sqrt{m}}. \quad (7)$$

- For $\alpha = \tilde{\alpha}$, it holds $\tilde{\alpha}_i = m_i/m$, where m_i is the number of target sample points of class i . We then get

$$\begin{aligned} D_{\Phi} \left(\sum_{i=1}^c \tilde{\alpha}_i \mathbb{Q}_i, \hat{\mathbb{Q}} \right) &= \left\| \frac{1}{m} \sum_{i=1}^c m_i \Phi(\mathbb{Q}_i) - \frac{1}{m} \sum_{j=n+1}^{n+m} \Phi(X_j) \right\| \\ &= \left\| \frac{1}{m} \sum_{j=n+1}^{n+m} \Phi(\mathbb{Q}_{Y_j}) - \frac{1}{m} \sum_{j=n+1}^{n+m} \Phi(X_j) \right\| \\ &= \left\| \frac{1}{m} \sum_{j=n+1}^{n+m} (\Phi(X_j) - \Phi(\mathbb{Q}_{Y_j})) \right\|. \end{aligned}$$

Now, notice that conditionally to the labels $(Y_j)_{j=n+1}^{n+m}$, the target sample points X_j are independent, not identically distributed but with respective class conditional distribution \mathbb{Q}_{Y_j} . We can therefore still appeal to Theorem A.1, and conclude that it holds with probability greater than $1 - \delta/2$:

$$\left\| \frac{1}{m} \sum_{j=n+1}^{n+m} (\Phi(X_j) - \Phi(\mathbb{Q}_{Y_j})) \right\| \leq C \frac{R_{1/\delta}}{\sqrt{m}}. \quad (8)$$

In both cases we therefore get the same bound for this last term.

Putting together (5), (6) and either (7) (for $\alpha = \alpha^*$) or (8) (for $\alpha = \tilde{\alpha}$) and $R_{1/\delta} \leq R_{c/\delta}$, gives the first claimed inequality of the theorem.

To obtain the second claimed inequality, we can see that the worst case scenario for w is obtain when $\alpha_i = 1$ for the smallest $\tilde{\beta}_i$ and 0 for the others. Hence, $\frac{1}{\sqrt{n}} \|w\|_2 \leq \max_i \frac{1}{\sqrt{n_i}}$. \square

Remark A.1. The best case scenario for $\|w\|_2$ is obtain when either α^* (or $\tilde{\alpha}$) equals $\tilde{\beta}$ or when $\beta_i = \frac{1}{n}$. In both cases, $\|w\|_2 = \sqrt{c}$.

Corollary A.1. *The result of Theorem 1 still holds, with Δ_{min} replace by λ_{min} , if we add a "dummy" class whose mapping is equals to 0, in other words, if we solve the soft-DFM version.*

Proof. Using the same notation introduced in the proof of Theorem 1, we have:

$$\begin{aligned} D_{\Phi}^2 \left(\sum_{i=0}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \sum_{i=0}^c \alpha_i \hat{\mathbb{P}}_i \right) &= (\hat{\alpha} - \alpha)^T \hat{\mathbf{G}} (\hat{\alpha} - \alpha) \\ &\geq \left(\min_{\|u\|=1} u^T \hat{\mathbf{G}} u \right) \|\hat{\alpha} - \alpha\|^2 \\ &= \lambda_{\min} \|\hat{\alpha} - \alpha\|^2 \end{aligned}$$

with $\alpha_0 := 0$ and $\min_{\|u\|=1} u^T \hat{\mathbf{G}} u$ the smallest eigenvalue of the Gram matrix $\hat{\mathbf{G}}$.

From this point on, we can follow the rest of the proof of Theorem 1. \square

A.2 Theorem 2

Let us first prove a lemma.

Lemma A.1. *Let \mathcal{H} be a Hilbert space, \mathcal{C} be a closed convex subset and V an affine subspace of \mathcal{H} such that $\mathcal{C} \subset V \subset \mathcal{H}$. For every $x \in \mathcal{H}$ we have*

$$\Pi_{\mathcal{C}}(x) = \Pi_{\mathcal{C}}(\Pi_V(x)),$$

with $\Pi_{\mathcal{C}}$ and Π_V the minimum distance projection functions onto \mathcal{C} and V .

Proof. Let us take $x \in \mathcal{H}$ and $c \in \mathcal{C}$. Note $p = \Pi_V(x)$, since $c \in V$ we can use Pythagoras' theorem :

$$\|x - c\|^2 = \|x - p\|^2 + \|p - c\|^2.$$

The point c that minimises $\|x - c\|^2$, i.e $\Pi_{\mathcal{C}}(x)$, is the same point c that minimises $\|p - c\|^2$, i.e $\Pi_{\mathcal{C}}(\Pi_V(x))$. \square

Let us now prove Theorem 2.

Proof. Below, we use the notation α for α^* .

Let us use the notation $D_{\Phi}(\mathbb{P}_1, \mathbb{P}_2) = \|\Phi(\mathbb{P}_1) - \Phi(\mathbb{P}_2)\|$.

We write $\mathcal{C}_n = \text{ConvHull}\langle \Phi(\hat{\mathbb{P}}_1), \dots, \Phi(\hat{\mathbb{P}}_c) \rangle$, the convex hull of the mapped empirical distributions $\hat{\mathbb{P}}_i$, and $\Pi_{\mathcal{C}_n}$ the projection onto this convex. In the same fashion let us write $\mathcal{C} = \text{ConvHull}\langle \Phi(\mathbb{P}_1), \dots, \Phi(\mathbb{P}_c) \rangle$ and the associated projection $\Pi_{\mathcal{C}}$, and finally V the affine subspace generated by the $\Phi(\mathbb{P}_i)$, namely $V = \{\sum_{i=1}^c \lambda_i \Phi(\mathbb{P}_i) \mid \sum_{i=1}^c \lambda_i = 1\}$. We will sometimes write $\Pi_{\mathcal{C}}(\mathbb{Q})$ as a shorthand for $\Pi_{\mathcal{C}}(\Phi(\mathbb{Q}))$.

With the same computation as in the proof of Theorem 1, we have:

$$\|\hat{\alpha} - \alpha\| \leq \frac{1}{\sqrt{\Delta_{\min}}} D_{\Phi} \left(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i \right).$$

Once again we have to upper-bound $D_{\Phi} \left(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i \right)$. By the triangle inequality:

$$D_{\Phi} \left(\sum_{i=1}^c \hat{\alpha}_i \hat{\mathbb{P}}_i, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i \right) \leq \underbrace{\left\| \sum_{i=1}^c \hat{\alpha}_i \Phi(\hat{\mathbb{P}}_i) - \Pi_{\mathcal{C}}(\mathbb{Q}) \right\|}_{(1)} + \underbrace{\left\| \Pi_{\mathcal{C}}(\mathbb{Q}) - \sum_{i=1}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) \right\|}_{(2)}. \quad (9)$$

Let us analyse the second term first. We use the triangle inequality:

$$\left\| \Pi_{\mathcal{C}}(\mathbb{Q}) - \sum_{i=1}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) \right\| \leq \left\| \Pi_{\mathcal{C}}(\mathbb{Q}) - \sum_{i=1}^c \alpha_i \Phi(\mathbb{P}_i) \right\| + D_{\Phi} \left(\sum_{i=1}^c \alpha_i \mathbb{P}_i, \sum_{i=1}^c \alpha_i \hat{\mathbb{P}}_i \right).$$

Using 5 and 6, the second term can be bounded by $CR_{c/\delta} \frac{\|w\|_2}{\sqrt{n}}$, with probability greater than $1 - \delta/2$, as we did in the proof of Theorem 1. For the first term we will require three elements:

1. $\sum_{i=1}^c \alpha_i \Phi(\mathbb{P}_i)$ lies in \mathcal{C} .
2. For all x , $\Pi_{\mathcal{C}}(x) = \Pi_{\mathcal{C}}(\Pi_V(x))$, see Lemma A.1.
3. $\Pi_{\mathcal{C}}$ is a contraction.

With that in mind,

$$\begin{aligned} \left\| \sum_{i=1}^c \alpha_i \Phi(\mathbb{P}_i) - \Pi_{\mathcal{C}}(\Phi(\mathbb{Q})) \right\| &= \left\| \Pi_{\mathcal{C}} \left(\sum_{i=1}^c \alpha_i \Phi(\mathbb{P}_i) \right) - \Pi_{\mathcal{C}}(\Pi_V(\Phi(\mathbb{Q}))) \right\| \\ &\leq \left\| \sum_{i=1}^c \alpha_i (\Phi(\mathbb{P}_i) - \Pi_V(\Phi(\mathbb{Q}_i))) \right\| \\ &\leq \max_i \|\Phi(\mathbb{P}_i) - \Pi_V(\Phi(\mathbb{Q}_i))\|. \end{aligned}$$

Hence, we have

$$\left\| \Pi_{\mathcal{C}}(\mathbb{Q}) - \sum_{i=1}^c \alpha_i \Phi(\hat{\mathbb{P}}_i) \right\| \leq CR_{c/\delta} \frac{\|w\|_2}{\sqrt{n}} + \underbrace{\max_i \|\Phi(\mathbb{P}_i) - \Pi_V(\Phi(\mathbb{Q}_i))\|}_{B^{\parallel}(\mathbb{P}, \mathbb{Q})}.$$

Let us turn to the first term of (9). By definition of $\hat{\alpha}$, it holds $\sum_{i=1}^c \hat{\alpha}_i \Phi(\hat{\mathbb{P}}_i) = \Pi_{\mathcal{C}_n}(\Phi(\hat{\mathbb{Q}}))$. Using the triangle inequality

$$\left\| \Pi_{\mathcal{C}_n}(\hat{\mathbb{Q}}) - \Pi_{\mathcal{C}}(\mathbb{Q}) \right\| \leq \left\| \Pi_{\mathcal{C}_n}(\hat{\mathbb{Q}}) - \Pi_{\mathcal{C}_n}(\mathbb{Q}) \right\| + \left\| \Pi_{\mathcal{C}_n}(\mathbb{Q}) - \Pi_{\mathcal{C}}(\mathbb{Q}) \right\|.$$

Since $\Pi_{\mathcal{C}_n}$ is a contraction, we have $\left\| \Pi_{\mathcal{C}_n}(\hat{\mathbb{Q}}) - \Pi_{\mathcal{C}_n}(\mathbb{Q}) \right\| \leq D_{\Phi}(\mathbb{Q}, \hat{\mathbb{Q}})$.

The only thing left to bound is the term $\|\Pi_{\mathcal{C}_n}(\Phi(\mathbb{Q})) - \Pi_{\mathcal{C}}(\Phi(\mathbb{Q}))\|$. For this we use results of [1] relating the distance between convex projections onto two different convex sets in relation to their Hausdorff distance. Recall the definition of the Hausdorff distance between two sets:

Definition A.1. *Let X and Y be two non-empty subsets of a metric space (M, d) . We define their Hausdorff distance $H(X, Y)$ by:*

$$H(X, Y) = \max \left\{ \sup_{x \in X} d(x, Y), \sup_{y \in Y} d(X, y) \right\}$$

Using Theorem 3.6 and Remark 3.7 of [1], the latter quantity is smaller than $\sqrt{2H(\mathcal{C}, \mathcal{C}_n)(r+d)}$, where $r = \text{dist}(0, \Phi(\mathbb{Q}))$, $d = \max\{\text{dist}(0, \mathcal{C}), \text{dist}(0, \mathcal{C}_n)\}$ and 0 the origin. Since the problem has a geometrical nature and is invariant by translation, we can translate everything so that $\Phi(\mathbb{Q})$ is the origin of the space. Hence, the bound reads $\sqrt{2H(\mathcal{C}, \mathcal{C}_n)d}$ with $d = \max\{\text{dist}(\Phi(\mathbb{Q}), \mathcal{C}), \text{dist}(\Phi(\mathbb{Q}), \mathcal{C}_n)\}$. Let us take care of the Hausdorff distance first:

$$\begin{aligned} \sup_{x \in \mathcal{C}} d(x, \mathcal{C}_n) &= \sup_{x \in \mathcal{C}} \|x - \Pi_{\mathcal{C}_n}(x)\| \\ &= \sup_{\lambda \in \Delta^c} \inf_{\beta \in \Delta^c} \left\| \sum_i \lambda_i \Phi(\mathbb{P}_i) - \sum_i \beta_i \Phi(\hat{\mathbb{P}}_i) \right\| \\ &\leq \sup_{\lambda \in \Delta^c} \left\| \sum_i \lambda_i \Phi(\mathbb{P}_i) - \sum_i \lambda_i \Phi(\hat{\mathbb{P}}_i) \right\| \\ &\leq \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i). \end{aligned}$$

A similar argument holds for $\sup_{x \in \mathcal{C}_n} d(x, \mathcal{C})$, and hence $H(\mathcal{C}, \mathcal{C}_n) \leq \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i)$.

We could simply bound d by 2 but we would obtain a loose bound. Instead, if we write $\Pi_{\mathcal{C}}(\Phi(\mathbb{Q})) = \sum_{i=1}^c \lambda_i \Phi(\mathbb{P}_i)$ then

$$\begin{aligned} d(\Phi(\mathbb{Q}), \mathcal{C}_n) &= \|\Phi(\mathbb{Q}) - \Pi_{\mathcal{C}_n}(\Phi(\mathbb{Q}))\| \\ &\leq \|\Phi(\mathbb{Q}) - \sum_{i=1}^c \lambda_i \Phi(\hat{\mathbb{P}}_i)\| \\ &\leq \|\Phi(\mathbb{Q}) - \Pi_{\mathcal{C}}(\Phi(\mathbb{Q}))\| + \left\| \sum_{i=1}^c \lambda_i \Phi(\mathbb{P}_i) - \sum_{i=1}^c \lambda_i \Phi(\hat{\mathbb{P}}_i) \right\| \\ &\leq \|\Phi(\mathbb{Q}) - \Pi_{\mathcal{C}}(\Phi(\mathbb{Q}))\| + \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i). \end{aligned}$$

Finally, using A.1, we get with probability higher than $1 - \delta/2$:

$$\begin{aligned} \sqrt{2H(\mathcal{C}, \mathcal{C}_n)d} &\leq \sqrt{2} \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i) + \sqrt{2 \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i) \|\Phi(\mathbb{Q}) - \Pi_{\mathcal{C}}(\Phi(\mathbb{Q}))\|} \\ &= \sqrt{2} \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i) + \sqrt{2 \max_i D_{\Phi}(\mathbb{P}_i, \hat{\mathbb{P}}_i) B^{\parallel}(\mathbb{P}, \mathbb{Q})} \\ &\leq \sqrt{2} \max_i \frac{CR_{c/\delta}}{\sqrt{n_i}} + \sqrt{2 \max_i \frac{CR_{c/\delta}}{\sqrt{n_i}} B^{\perp}(\mathbb{P}, \mathbb{Q})}. \end{aligned}$$

By putting everything together, with probability at least $(1 - \delta)$, we have:

$$\begin{aligned} \|\hat{\alpha} - \alpha^*\| &\leq \frac{1}{\sqrt{\Delta_{\min}}} \left(\frac{CR_{1/\delta}}{\sqrt{m}} + \sqrt{2} \max_i \frac{CR_{c/\delta}}{\sqrt{n_i}} + \sqrt{2 \max_i \frac{CR_{c/\delta}}{\sqrt{n_i}} B^{\perp}(\mathbb{P}, \mathbb{Q})} \right. \\ &\quad \left. + CR_{c/\delta} \frac{\|w\|_2}{\sqrt{n}} + B^{\parallel}(\mathbb{P}, \mathbb{Q}) \right) \\ &\leq \frac{1}{\sqrt{\Delta_{\min}}} \left(c_1 \max_i \frac{CR_{c/\delta}}{\sqrt{n_i}} + \frac{CR_{1/\delta}}{\sqrt{m}} + \sqrt{2 \max_i \frac{CR_{c/\delta}}{\sqrt{n_i}} B^{\perp} + B^{\parallel}} \right), \end{aligned}$$

with $c_1 = 1 + \sqrt{2} \leq 3$. \square

A.3 Corollary 1

We directly apply Theorem 2 with the "dummy" class $\phi(\mathbb{P}_0) := 0$. In that case, if we write $\tilde{\mathbf{G}}$ the Gram matrix of $\{\phi(\mathbb{P}_0), \phi(\hat{\mathbb{P}}_1) \cdots \phi(\hat{\mathbb{P}}_c)\}$ then, as the first column of this matrix is zero,

$$\Delta_{\min}(\tilde{\mathbf{G}}) = \min_{\substack{\|u\|=1 \\ \mathbf{1}^T u = 0}} u^T \tilde{\mathbf{G}} u = \min_{\|u\|=1} u^T \hat{\mathbf{G}} u = \lambda_{\min}.$$

In the same fashion, the affine subspace $V := \text{AffSpan}\{\Phi(\mathbb{P}_i), i \in [0, \dots, c]\}$ is equal to $\bar{V} := \text{Span}\{\Phi(\mathbb{P}_i), i \in [c]\}$. Finally, the convex hull $\text{ConvHull}\{\Phi(\mathbb{P}_i), i \in$

$[0, \dots, c]$ is equal to $\text{int}(C)$: the interior of the convex hull $\text{ConvHull}\{\Phi(\mathbb{P}_i), i \in [c]\}$

As we are in the contaminated label shift setting, $\mathbb{P}_i = \mathbb{Q}_i$ and hence

$$B^\parallel(\mathbb{P}, \mathbb{Q}) = \max_i \|\Phi(\mathbb{P}_i) - \Pi_{\bar{V}}(\Phi(\mathbb{Q}_i))\|_{\mathcal{F}} = \|\Pi_{\bar{V}}(\Phi(\mathbb{Q}_0))\|$$

The "orthogonal" term $B^\perp(\mathbb{P}, \mathbb{Q})$ can be bounded by $\sqrt{\alpha_0 \|\Phi(\mathbb{Q}_0)\|}$ as follows:

$$\begin{aligned} B^\perp(\mathbb{P}, \mathbb{Q})^2 &= \|\Phi(\mathbb{Q}) - \Pi_{\text{int}(C)}(\Phi(\mathbb{Q}))\| \\ &\stackrel{(\dagger)}{\leq} \left\| \sum_{i=1}^c \alpha_i^* \phi(\mathbb{P}_i) + \alpha_0 \phi(\mathbb{Q}_0) - \sum_{i=1}^c \alpha_i^* \phi(\mathbb{P}_i) \right\| \\ &= \alpha_0 \|\phi(\mathbb{Q}_0)\|, \end{aligned}$$

for the inequality (\dagger) , we use the fact that $\sum_{i=1}^c \alpha_i^* \phi(\mathbb{P}_i) \in \text{int}(C)$.

Remark A.2. Both theorem 2 and corollary 1 are true if we replace α^* by $\tilde{\alpha}$.

A.4 Proposition 1.

Proof. It is straightforward to check that for the mentioned feature mapping, $\Phi(\hat{\mathbb{P}}_i)$ in the DFM setting is exactly the i -th column of M in the BBSE notation, and $\Phi(\hat{\mathbb{Q}}) = Y$. Hence the DFM objective in (\mathcal{P}) rewrites to $\|\alpha^T M - Y\|^2$, and since M is assumed invertible, in that setting the unconstrained solution is $\hat{\alpha} = M^{-1}Y$. Furthermore, the sum-1 condition $\mathbf{1}^T \alpha = 1$ (where $\mathbf{1}$ denotes a vector of ones of dimension c) is automatically satisfied for the unconstrained solution: obviously it holds $\mathbf{1}^T M = \mathbf{1}^T$, hence $\mathbf{1}^T = \mathbf{1}^T M^{-1}$, and $\mathbf{1}^T Y = 1$, so that $\mathbf{1}^T M^{-1} Y = 1$.

B Properties of Δ_{\min}

Let (b_1, \dots, b_c) be a c -tuple of vectors of \mathbb{R}^D assumed to be linearly independent and \bar{b} their mean. We denote M the Gram matrix of those vectors, i.e. $M_{ij} = \langle b_i, b_j \rangle$. We also write M^c the centered Gram matrix of the vectors : $(M^c)_{i,j} = \langle b_i - \bar{b}, b_j - \bar{b} \rangle$. Finally, we denote $\mathbf{1}$ a vector of ones (of dimension c).

In this appendix we prove two claims about the quantity $\Delta_{\min} := \Delta_{\min}(b_1, \dots, b_c)$ defined in definition 1 as the second smallest eigenvalue of M^c .

Theorem B.1. *For any number of classes c , Δ_{\min} is equal to $\min_{\substack{\|u\|=1 \\ \mathbf{1}^T u=0}} u^T M u$.*

Proof. Let us take u such that $\|u\| = 1$, $\mathbf{1}^T u = 0$ and P the projection matrix on $\langle \mathbf{1} \rangle^\perp$, such that $Pu = u$. We have:

$$\begin{aligned} u^T M u &= (Pu)^T M (Pu) \\ &= u^T (PMP) u. \end{aligned}$$

Observe that PMP is a symmetric matrix of rank $c-1$, hence the eigenvectors v_i (associated to the eigenvalues $\lambda_1 \geq \lambda_2 \dots \geq \lambda_c = 0$) form an orthonormal base. In particular $v_c = c^{-1/2} \mathbf{1}$. Since $u \in \langle \mathbf{1} \rangle^\perp$, then $u \in \langle v_1, \dots, v_{c-1} \rangle$. There exist $\alpha \in \mathbb{R}^{c-1}$ such that $u = \sum_{i=1}^{c-1} \alpha_i v_i$, and since $\|u\| = 1$ then $\|\alpha\| = 1$. With that in mind :

$$\begin{aligned} u^T (PMP) u &= \left(\sum_{i=1}^{c-1} \alpha_i v_i \right)^T \left(\sum_{i=1}^{c-1} \alpha_i (PMP) v_i \right) \\ &= \left(\sum_{i=1}^{c-1} \alpha_i v_i \right)^T \left(\sum_{i=1}^{c-1} \alpha_i \lambda_i v_i \right) \\ &= \sum_{i,j=1}^{c-1} \lambda_i \alpha_i \alpha_j \langle v_i, v_j \rangle \\ &= \sum_{i=1}^{c-1} \lambda_i \alpha_i^2. \end{aligned}$$

Hence $\Delta_{\min}(M)$ is equal to $\min_{\|\alpha\|=1} \sum_{i=1}^{c-1} \lambda_i \alpha_i^2$. Using the change of variable $\beta_i = \alpha_i^2$, this is equivalent to find : $\min_{\beta \in \Delta^{c-1}} \langle \lambda, \beta \rangle$, which is equal to λ_{c-1} . All that is left to do is a straightforward computation of PMP with $P = I_c - \frac{1}{c} \mathbf{1}^T \mathbf{1}$, to find that $(PMP)_{i,j} = \langle b_i - \bar{b}, b_j - \bar{b} \rangle$.

A direct corollary of this theorem is that λ_{\min} is greater than the smallest singular value of the matrix (b_1, \dots, b_c) .

Theorem B.2. *In particular for two classes, $\lambda_{\min}(b_1, b_2) = \frac{1}{2} \|b_1 - b_2\|^2$.*

Proof. In two dimensions, the conditions $\|x\| = 1$ and $\mathbf{1}^T x = 0$ can only be verified for 2 points, $x = \left(\sqrt{\frac{1}{2}}, -\sqrt{\frac{1}{2}}\right)$ and $x = \left(-\sqrt{\frac{1}{2}}, \sqrt{\frac{1}{2}}\right)$. If we compute $x^T M x$ for these two points we obtain : $\frac{1}{2}\|b_1\|^2 - \langle b_1, b_2 \rangle + \frac{1}{2}\|b_1\|^2$ which is equal to $\frac{1}{2}\|b_1 - b_2\|^2$.

C Short review of Random Fourier Features

Random Fourier Features are based on Bochner’s Theorem:

Theorem C.1. *A continuous function φ on \mathbb{R}^D is positive definite if and only if φ is the Fourier transform of a non-negative measure.*

A direct corollary of this result is that every continuous invariant kernel k (associated to a function φ) is the Fourier transform of a non-negative measure that we denote Λ_k . One can show that $\Lambda_k(\mathbb{R}^D) = \varphi(0)$. Hence, for a normalized continuous invariant kernel, Λ_k is a distribution referred to as the *spectral distribution* of k .

The kernel function can hence be written as:

$$k(x, y) = \mathbb{E}_{\omega \sim \Lambda_k} [e^{i\omega^T(x-y)}] = \mathbb{E}_{\omega \sim \Lambda_k} [\cos(\omega^T(x-y))].$$

By the Monte-Carlo principle, using a sample $(\omega_i)_{i=1}^{D/2}$ i.i.d. from Λ_k , the feature map $\Phi: \mathcal{X} \rightarrow \mathbb{R}^D$ defined by

$$\Phi(x) = \sqrt{\frac{2}{D}} [\cos(\omega_i^T x), \sin(\omega_i^T x)]_{i=1}^{D/2} \quad (10)$$

is such that $k(x, y) = \mathbb{E}[\tilde{\Phi}(x)^T \tilde{\Phi}(y)]$, where the expectation is taken with respect to $(\omega_i)_{i=1}^{D/2}$.

Another vectorisation used in practice is

$$\Phi(x) = \sqrt{\frac{2}{D}} [\cos(\omega_i^T x + b_i)]_{i=1}^D, \quad (11)$$

where b_i are i.i.d samples from the uniform distribution on $[0, 2\pi]$. See [23] for the detailed computation. Even, if this vectorisation is popular in practice, we would like to point out that the second version yields worst results both in term of variance and upper-bound for the Gaussian kernel [41].

D Additional figures

Figure 3 shows the three types of contamination tested in Section 4.2.

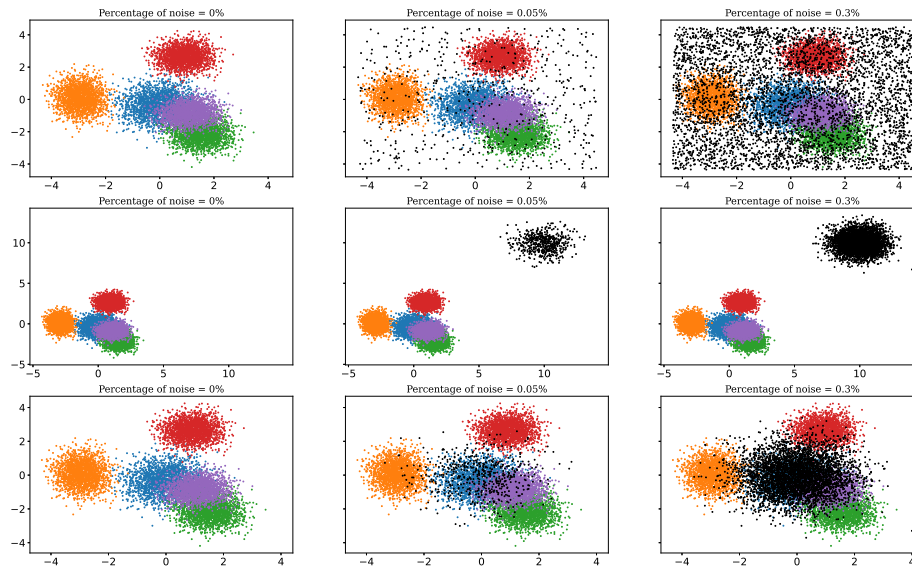


Fig. 3. The first row represents the background uniform noise for different values of ϵ . The second row represents the case where a new class appears far from the other distributions. Finally the last row is the scenario where the new class appears close to the other distributions.