



A unified framework for dataset shift diagnostics

Felipe Maia Polo ^{a,*}, Rafael Izbicki ^b, Evanildo Gomes Lacerda Jr ^c,
Juan Pablo Ibieta-Jimenez ^d, Renato Vicente ^e

^a Department of Statistics, University of Michigan, United States of America

^b Department of Statistics, Federal University of São Carlos, Brazil

^c Trustly, Brazil

^d Experian LatAm DataLab, Brazil

^e Department of Applied Mathematics, Institute of Mathematics and Statistics, University of São Paulo, Brazil



ARTICLE INFO

Keywords:

Dataset shift detection
Hypothesis testing
Transfer learning

ABSTRACT

Supervised learning techniques typically assume training data originates from the target population. Yet, in reality, *dataset shift* frequently arises, which, if not adequately taken into account, may decrease the performance of the supervised predictors. In this work, we propose a novel and flexible framework called *DetectShift*¹ that quantifies and tests for multiple dataset shifts, encompassing shifts in the distributions of (\mathbf{X}, Y) , \mathbf{X} , Y , $\mathbf{X}|Y$, and $Y|\mathbf{X}$. *DetectShift* equips practitioners with insights into data shifts, facilitating the adaptation or retraining of predictors using both source and target data. This proves extremely valuable when labeled samples in the target domain are limited. The framework utilizes test statistics with the same nature to quantify the magnitude of the various shifts, making results more interpretable. It is versatile, suitable for both regression and classification tasks, and accommodates diverse data forms—tabular, text, or image. Experimental results demonstrate the effectiveness of *DetectShift* in detecting dataset shifts even in higher dimensions.

1. Introduction

In machine learning applications, it is conventionally assumed that training data originates from the distribution of interest. We say a dataset shift has happened when that assumption does not hold. Formally, we have dataset shift when the joint distribution of features (\mathbf{X}) and labels (Y) associated with the training sample, also known as the source distribution, $P_{\mathbf{X}, Y}^{(1)}$, and the distribution of interest, also known as the target distribution, $P_{\mathbf{X}, Y}^{(2)}$, are different. Real-world applications, spanning fields like finance [42], health [10], technology [24], and physics [11], often grapple with the challenges posed by dataset shifts. Unfortunately, dataset shift may substantially decrease the predictive power of machine learning models if it is not adequately addressed [43].

When labeled data from the target distribution is scarce, addressing dataset shifts typically necessitates assumptions about the relationship between the source and target distributions. Different assumptions are translated into different types of shift [30], and every kind of shift demands specific adaptation methods [37,43,29]. For instance, if $P_{\mathbf{X}}^{(1)} \neq P_{\mathbf{X}}^{(2)}$ but $P_{Y|\mathbf{X}}^{(1)} = P_{Y|\mathbf{X}}^{(2)}$, dataset shift adaptation can be performed by using importance weighting on the training data [43,16,28]. Similarly, if $P_{Y}^{(1)} \neq P_{Y}^{(2)}$ but $P_{\mathbf{X}|Y}^{(1)} = P_{\mathbf{X}|Y}^{(2)}$,

* Corresponding author.

E-mail address: felipemaiapolo@gmail.com (F. Maia Polo).

¹ Our implementation for *DetectShift* can be found in <https://github.com/felipemaiapolo/detectshift>.

<https://doi.org/10.1016/j.ins.2023.119612>

Received 2 February 2023; Received in revised form 25 August 2023; Accepted 27 August 2023

Available online 1 September 2023

0020-0255/© 2023 Elsevier Inc. All rights reserved.

dataset shift adaptation can be performed by importance weighting [25] or re-calibrating posterior probabilities via Bayes' theorem [37,46]. Therefore, to successfully adapt prediction algorithms in a dataset shift setting, practitioners need to know not only *if* dataset shift occurs but also *which type* of shift happens for the data at hand.

This study introduces *DetectShift*, a novel unified framework capable of accurately quantifying and testing for diverse types of dataset shifts. This framework provides valuable insights to practitioners on optimally handling changes in the data, particularly in scenarios where labeled samples from the target domain are scarce. By using our framework, practitioners can detect specific types of shifts. Consequently, they can use that information to use both source and target data to retrain or adapt their predictors, which is especially appealing in situations where naive model retraining using only data points from $P_{X,Y}^{(2)}$ is not feasible.² In Section 4, we show in practice one example of how the insights provided by *DetectShift* can help practitioners in a prediction task in which target-labeled samples are not used to retrain the classifier.

The remainder of this paper is organized as follows. Firstly, in Section 2, we provide an overview of recent works relevant to our research, highlighting our contributions to the existing literature. Moving on to Section 3, we present our framework called *DetectShift*, explaining its conceptual basis and rationale. To demonstrate the effectiveness of our approach across various scenarios, we apply *DetectShift* to both artificial and real data in Section 4. Finally, in the concluding sections, we thoroughly discuss our findings, contributions, limitations, and potential avenues for further exploration.

1.1. Contributions

The main contributions of *DetectShift* are that:

- It facilitates quantification and formal testing of shifts in the distributions of (X, Y) , X , Y , $X|Y$, and $Y|X$. Identifying different types of shifts helps practitioners gain insights about changes in their data. Those insights allow them to leverage source and target data to retrain or adapt their predictors², which is especially beneficial when target labels are scarce.
- It utilizes test statistics with the same nature to quantify each type of dataset shift, enabling practitioners to compare the magnitude of different shifts and gain meaningful insights. Furthermore, empirical evidence demonstrates that these test statistics result in powerful shift detectors even in higher dimensions.
- It can be applied to regression and classification tasks. Additionally, it can be employed in conjunction with virtually any type of data, including tabular, text, and image data.

2. Related work

Dataset shift detection methods can have different objectives. For instance, specific approaches monitor the test error of a predictor over time, triggering an alarm when there is a significant degradation in its performance [12,32]. Other approaches are specifically designed to provide a deeper understanding of how data distribution has changed by detecting specific types of shifts, e.g., changes in the distribution of features. Our research aligns closely with the latter category, which we will delve into next.

Source and target distributions can be different in many ways. For instance, if $P_Y^{(1)} \neq P_Y^{(2)}$ but $P_{X|Y}^{(1)} = P_{X|Y}^{(2)}$, then the dataset shift can be characterized as label shift. If labels from the target domain are unavailable, label shift detection and quantification are not straightforward but can be accomplished under some conditions [25,46]. The method proposed by Lipton et al. [25], for example, exploits arbitrary black box classifiers and their confusion matrices to estimate $P_Y^{(2)}$. In a different direction, various methods are available when we are interested in detecting or understanding changes in the marginal distribution of possibly high-dimensional features X . Detecting shifts in P_X can be solved using two-sample tests [35,26], and some extensions are possible. For example, Jang et al. [17] proposes an online approach in which data is seen sequentially, and Wijaya et al. [49] proposes extracting concept-based features from images (e.g., rotation, shape) and then testing for shifts of those concepts. Moreover, Ginsberg et al. [13] proposes a method for *harmful* covariate shift detection while Luo et al. [27] proposes a martingale-based approach in which false alarms are controlled over time. For objectives centered on detecting shifts in conditional distributions like $P_{Y|X}$, several methods emerge. For instance, Schrouff et al. [38] proposes using a conditional independence test to detect shifts in conditional distributions in the context of algorithmic fairness. In contrast, Vovk [47] proposes a martingale-based approach to detect online changes in $P_{X|Y}$ considering classification problems.

While identifying specific types of shifts is essential, it is typically insufficient for making informed decisions. Ideally, practitioners should access a unified framework that evaluates multiple dataset shifts, guiding them on optimal next steps. That framework would also allow the practitioner to compare the strength of the several shift types, permitting a more interpretable analysis. Webb et al. [48] proposes a framework with those characteristics. The authors use the total variation (TV) distance between two probability distributions to measure shifts in P_Y , P_X , $P_{Y|X}$, and $P_{X|Y}$, to provide the practitioner with valuable insights. Their approach is limited, however. Estimating the TV distance between two probability distributions is challenging when variables are not discrete or their dimensionality is high (see Fig. 3 in our experiments section). Moreover, Webb et al. [48] does not suggest using formal hypothesis testing. Thus, false alarm control is not guaranteed. Building upon Webb et al. [48], our work augments their framework, offering adaptability to high-dimensional/continuous data and integrating rigorous hypothesis testing.

² See [37,25,43,29] for some examples on how to retrain or adapt predictors using not only target samples.

3. Methodology

We observe two datasets, $D^{(1)}$ and $D^{(2)}$, where

$$D^{(i)} = \left\{ \left(\mathbf{X}_1^{(i)}, Y_1^{(i)} \right), \dots, \left(\mathbf{X}_{n^{(i)}}^{(i)}, Y_{n^{(i)}}^{(i)} \right) \right\},$$

for $i = 1, 2$. We assume that all observations from the same dataset are i.i.d. and that the datasets are independent of each other. We denote by $P_{X,Y}^{(i)}$ the distribution associated with an observation from the i -th dataset, where $i = 1$ stands for source and $i = 2$ for the target. We aim to quantify and test which types of dataset shifts occur from source to target domain. The null hypotheses we want to test are

- **[Total Dataset Shift]** $H_{0,D} : P_{X,Y}^{(1)} = P_{X,Y}^{(2)}$
- **[Feature Shift]** $H_{0,F} : P_X^{(1)} = P_X^{(2)}$
- **[Response Shift]** $H_{0,R} : P_Y^{(1)} = P_Y^{(2)}$
- **[Conditional Shift - Type 1]** $H_{0,C1} : P_{X|Y}^{(1)} = P_{X|Y}^{(2)}$ ($P_Y^{(2)}$ - almost surely)
- **[Conditional Shift - Type 2]** $H_{0,C2} : P_{Y|X}^{(1)} = P_{Y|X}^{(2)}$ ($P_X^{(2)}$ - almost surely)

Section 3.1 describes how to obtain statistics that can quantify the amount of each type of dataset shift and Section 3.2 shows how to use the statistics to test each type of shift's occurrence.

3.1. Test statistics

Our statistics are based on the Kullback-Leibler (KL) divergence [21,33], a well-known measure to describe discrepancies between probability measures. Formally, the KL divergence between two probability distributions P and Q is defined as

$$\text{KL}(P||Q) := \int \log \left(\frac{dP}{dQ} \right) dP,$$

where $\frac{dP}{dQ}$ is the Radon–Nikodym (R-N) derivative of P with respect to Q . If both P and Q are (Lebesgue) continuous distributions, the R-N derivative is simply a density ratio p/q . For the R-N derivative to be well defined, we need P to be absolutely continuous with respect to Q [33], meaning that the support of P is a subset of the support of Q . From now on, we assume that $P_{X,Y}^{(2)}$ is absolutely continuous with respect to $P_{X,Y}^{(1)}$. Then, we use the following quantities to measure each of the shifts described at the beginning of Section 3:

- **[Total Dataset Shift]** $\text{KL}_{X,Y} := \text{KL}(P_{X,Y}^{(2)} || P_{X,Y}^{(1)})$
- **[Feature Shift]** $\text{KL}_X := \text{KL}(P_X^{(2)} || P_X^{(1)})$
- **[Response Shift]** $\text{KL}_Y := \text{KL}(P_Y^{(2)} || P_Y^{(1)})$
- **[Conditional Shift - Type 1]** $\text{KL}_{X|Y} := \mathbb{E}_{P_Y^{(2)}} \left[\text{KL}(P_{X|Y}^{(2)} || P_{X|Y}^{(1)}) \right]$
- **[Conditional Shift - Type 2]** $\text{KL}_{Y|X} := \mathbb{E}_{P_X^{(2)}} \left[\text{KL}(P_{Y|X}^{(2)} || P_{Y|X}^{(1)}) \right]$

The following proposition states that we can rewrite the null hypotheses we want to test in terms of the quantities above.

Proposition 1. *The hypotheses $H_{0,D}$, $H_{0,F}$, $H_{0,R}$, $H_{0,C1}$, and $H_{0,C2}$ can be rewritten equivalently as follows*

$$H_{0,D} : \text{KL}_{X,Y} = 0, \quad H_{0,F} : \text{KL}_X = 0, \quad H_{0,R} : \text{KL}_Y = 0, \quad H_{0,C1} : \text{KL}_{X|Y} = 0, \quad H_{0,C2} : \text{KL}_{Y|X} = 0$$

The proof follows from the elementary KL divergence properties [33], e.g., $\text{KL}(P||Q) = 0$ if and only if $P = Q$. For the conditional shifts, see that $\text{KL}(P_{Y|X}^{(2)} || P_{Y|X}^{(1)})$ and $\text{KL}(P_{X|Y}^{(2)} || P_{X|Y}^{(1)})$ are non-negative random variables, consequently, their expected values are zero if and only if they are zero almost surely (with probability one).

To quantify and test the different types of shift, we use estimators of the parameters $\text{KL}_{X,Y}$, KL_X , KL_Y , $\text{KL}_{X|Y}$, and $\text{KL}_{Y|X}$ as test statistics. That is a reasonable choice because (i) our null hypotheses can be equivalently written in terms of such estimable parameters, and (ii) the magnitude of the statistics is directly related to the shift intensities. These suggest that tests based on these statistics will be powerful in detecting shifts. Moreover, all parameters are integrals computed using the target distribution. Thus, they give more weight to regions of the feature/label space to which most target data points belong, letting us focus on regions that matter.

To estimate $\text{KL}_{X,Y}$, KL_X , and KL_Y , we first use³ training data and the probabilistic classification method for density ratio estimation [45,7,8], also known as odds-trick, to estimate the Radon–Nikodym derivative between the two probability distributions. Then, we use test data to estimate the divergences. More precisely, we first create the augmented dataset

$$D = \{(\mathbf{X}_1, Y_1, Z_1), \dots, (\mathbf{X}_n, Y_n, Z_n)\},$$

where each (\mathbf{X}_k, Y_k) corresponds to a different observation taken from $D^{(1)} \cup D^{(2)}$ and $Z_k \in \{1, 2\}$ indicates from which dataset (\mathbf{X}_k, Y_k) comes from. We then randomly split D into two sets: D^{tr} (training set) and D^{te} (test set). We use D^{tr} to train a probabilistic classifier that predicts Z . The features used to predict Z are (i) (\mathbf{X}, Y) to estimate the amount of total dataset shift ($\text{KL}_{X,Y}$), (ii) \mathbf{X} to estimate the amount of feature shift (KL_X), and (iii) Y to estimate the amount of response shift (KL_Y). The estimated Radon–Nikodym derivative in the case of total dataset shift (the other ones are analogous) is given by

$$\widehat{\frac{dP_{X,Y}^{(2)}}{dP_{X,Y}^{(1)}}}(\mathbf{x}, y) := \frac{n_{tr}^{(1)}}{n_{tr}^{(2)}} \frac{\widehat{\mathbb{P}}(Z = 2 | \mathbf{X} = \mathbf{x}, Y = y)}{\widehat{\mathbb{P}}(Z = 1 | \mathbf{X} = \mathbf{x}, Y = y)},$$

where $\widehat{\mathbb{P}}$ denotes the trained probabilistic classifier and $n_{tr}^{(i)}$ is the number of samples from population i in D^{tr} . If target-labeled samples are scarce, training the classifier $\widehat{\mathbb{P}}$ from scratch can be challenging when \mathbf{X} is high-dimensional. However, if target unlabeled samples are abundant, one possible solution is first training a classifier only depending on \mathbf{X} and then using $h(\mathbf{x}) = \widehat{\mathbb{P}}(Z = 2 | \mathbf{X} = \mathbf{x})$ to reduce the dimensions of \mathbf{X} before training the classifier dependent on both \mathbf{X} and Y . This solution does not affect the reliability of our method since false alarm control is not affected.

Finally, we use empirical averages⁴ over the test dataset D^{te} to estimate the KL divergence (again in the case of total dataset shift; the other ones are analogous):

$$\widehat{\text{KL}}_{X,Y} := \frac{1}{|D_2^{te}|} \sum_{(\mathbf{X}_k, Y_k, Z_k) \in D_2^{te}} \log \left(\frac{\widehat{\frac{dP_{X,Y}^{(2)}}{dP_{X,Y}^{(1)}}}(\mathbf{X}_k, Y_k)}{\widehat{\frac{dP_{X,Y}^{(1)}}{dP_{X,Y}^{(2)}}}(\mathbf{X}_k, Y_k)} \right),$$

where $D_2^{te} := \{(\mathbf{X}_k, Y_k, Z_k) \in D^{te} : Z_k = 2\}$ denotes the test samples from the second population.

This approach however cannot be used to estimate $\text{KL}_{Y|X}$ or $\text{KL}_{X|Y}$. Instead, we rely on the KL divergence decomposition, given in the following proposition extracted from Polyanskiy and Wu [33] (Theorem 2.13). To keep this text as self-contained as possible, we include a proof in the appendix.

Proposition 2. Let KL_Y , KL_X , $\text{KL}_{Y|X}$, $\text{KL}_{X|Y}$, $\text{KL}_{X,Y}$ be defined as they were in Section 3.1. Then

$$\text{KL}_{X,Y} = \text{KL}_{Y|X} + \text{KL}_X = \text{KL}_{X|Y} + \text{KL}_Y$$

This result shows that the KL divergences of the conditional distributions can be estimated via $\widehat{\text{KL}}_{X|Y} := \widehat{\text{KL}}_{X,Y} - \widehat{\text{KL}}_Y$ and $\widehat{\text{KL}}_{Y|X} := \widehat{\text{KL}}_{X,Y} - \widehat{\text{KL}}_X$.

3.2. Hypothesis testing

Once we have statistics that can quantify the magnitude of different types of dataset shifts, we can use them to formally test the hypotheses described in Section 3.1. In this section, Y can be discrete or continuous, except when obtaining the p -values for the hypothesis $P_{X|Y}^{(1)} = P_{X|Y}^{(2)}$, in which we assume it is discrete. This is needed since Algorithm 1, in the appendix, relies on this assumption. If Y is continuous or has few repeated values, the conditional shift of type 1 can be tested by discretizing/binning the label for computing the statistic and applying the algorithm⁵ – we give more details and references at the end of this section.

Consider the datasets D^{te} and D_2^{te} as defined in Section 3.1, and let $T(D_2^{te})$ be a test statistic of interest computed using D_2^{te} . Namely, $T(D_2^{te})$ can represent any of the following random quantities, depending on which type of shift we are testing for: $\widehat{\text{KL}}_Y$, $\widehat{\text{KL}}_X$, $\widehat{\text{KL}}_{Y|X}$, $\widehat{\text{KL}}_{X|Y}$, $\widehat{\text{KL}}_{X,Y}$. We test each of the hypotheses of interest by computing a p -value of the form

$$p(D^{te}) = \frac{1 + \sum_{j=1}^B \mathbb{I} \left[T(D_2^{te}) \leq T(\tilde{D}_2^{te(j)}) \right]}{B + 1}, \quad (1)$$

where each $\tilde{D}_2^{te(j)}$ is a modified version of D_2^{te} , which can depend on the whole test set D^{te} . The modification that is done depends on the hypothesis we are testing:

³ When Y is discrete, KL_Y can also be estimated by using a plug-in estimator described in the appendix.

⁴ The same approach is used to estimate divergences by Sønderby et al. [41] in the context of generative models, for example.

⁵ Binning is not needed when training the classifiers though.

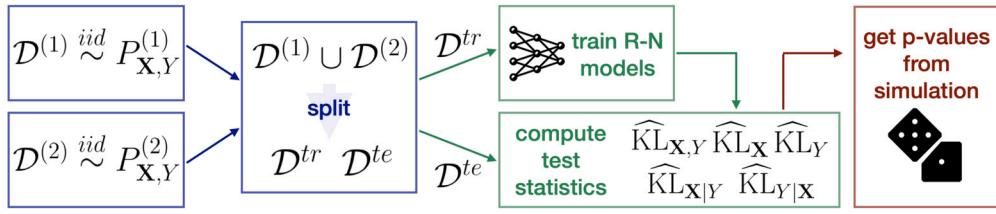


Fig. 1. *DetectShift* in a nutshell. First, data are sampled from source and target distributions, mixed together, and split into a training and a test set. Then, using the training set, the Radon-Nikodym (R-N) models are trained, and all the test statistics are computed using the test set. Finally, *p*-values to conduct the tests are obtained through simulation procedures.

- To test the hypotheses related to unconditional distributions ($H_{0,D}$, $H_{0,F}$ and $H_{0,R}$), $\tilde{D}_2^{te(j)}$ is obtained randomly permuting Z_k 's on D^te and then selecting the samples with $Z_k = 2$ to form the modified version of D_2^{te} . In this case, p is the *p*-value associated with a permutation test, a method commonly used to perform two-sample tests [23].
- To test $H_{0,C1}$, $\tilde{D}_2^{te(j)}$ is obtained randomly permuting the values of Z_k 's *within each level of Y* on D^te and then selecting the samples with $Z_k = 2$ to form the modified version of D_2^{te} . Thus, we require Y to be discrete to apply this test. In this case, p is the *p*-value associated with a *conditional independence* local permutation test [18].
- To test $H_{0,C2}$, we first estimate the conditional distribution of $Y|X$ using the whole training set D^tr of labeled samples. Let $Q_{Y|X}$ denote such an estimate, which can be obtained using any probabilistic classifier, such as logistic regression, neural networks, CatBoost classifier [34], or conditional density estimators [15] and GANs [3] if Y is continuous. We then obtain $\tilde{D}_2^{te(j)}$ by replacing each Y_k in D_2^{te} by a random draw from $Q_{Y|X=x_k}$. This test is known as the conditional randomization test (CRT) [5,4].

Algorithm 1, in the appendix, details the steps to obtain the *p*-values for each case. For all the cases, we assume the procedure involving the training of probabilistic classifiers, described in Section 3.1, has already been executed. That is, we have a test statistic T for every test we want to perform. Also, for the case we are testing for conditional shift type 2, we assume the estimated conditional distribution $Q_{Y|X}$ has been computed. We fix a significance level $\alpha \in (0, 1)$, and after calculating the *p*-value p for a specific null hypothesis of interest, reject that hypothesis if $p \leq \alpha$. Fig. 1 summarizes how our framework, *DetectShift*, works.

The following proposition shows that such tests are valid (that is, they control Type I error rate). In other words, false alarms are controlled. The only exception is the test for $H_{0,C2}$, which approximately controls Type I error probabilities as long as $Q_{Y|X}$ is a good approximation of $P_{Y|X}$. For the following result, we consider that T and $Q_{Y|X}$ are given and fixed, as both are obtained from the training set. Also, we make sure that $(T(D_2^{te}), T(\tilde{D}_2^{te(1)}), \dots, T(\tilde{D}_2^{te(B)}))$ has no repeated values by adding small centered Gaussian noises⁶ to $T(D_2^{te})$ and $T(\tilde{D}_2^{te(j)})$, for every j .

Proposition 3. *Let $p(D^te)$ be a *p*-value obtained from Algorithm 1 (appendix) with fixed T and $Q_{Y|X}$. Then, for every $\alpha \in (0, 1)$,*

- Under $H_{0,D}$, $H_{0,F}$, $H_{0,R}$, and $H_{0,C1}$ (if Y is discrete),

$$\mathbb{P}(p(D^te) \leq \alpha) \leq \alpha$$

- Under $H_{0,C2}$

$$\mathbb{P}(p(D^te) \leq \alpha) \leq \alpha + \mathbb{E}_{\tilde{P}_X^{(2)}} [d_{TV}(\tilde{Q}_{Y|X}, \tilde{P}_{Y|X})],$$

where d_{TV} is the total variation distance, $\tilde{P}_X^{(2)} := \prod_{(X_k, Y_k, Z_k) \in D_2^{te}} P_{X_k}^{(2)}$, $\tilde{Q}_{Y|X} := \prod_{(X_k, Y_k, Z_k) \in D_2^{te}} Q_{Y|X=x_k}$, and $\tilde{P}_{Y|X} = \prod_{(X_k, Y_k, Z_k) \in D_2^{te}} P_{Y|X=x_k}^{(2)}$.

Here, \prod denotes products of probability distributions. Proposition 3 borrows well-studied results from the Statistics literature. The results for $H_{0,D}$, $H_{0,F}$, and $H_{0,R}$ are directly obtained by the fact that we use a permutation test [23], while the result for $H_{0,C1}$ is obtained because of properties of local permutation [18] and the result for $H_{0,C2}$ is obtained by adapting the results of Berrett et al. [4] (Theorem 4) to our context. If Y is continuous or has few repeated values, if we discretize/bin it to test the conditional shift of type 1, our approach leads to an approximate test for $H_{0,C1}$ in the sense it approximately controls the Type I error. See Kim et al. [18] (Theorems 2 and 3) for more details. To make this work self-contained, we include a proof for Proposition 3 in the appendix.

4. Experiments

This section presents numerical experiments with both artificial and real data. In all the experiments in which Y is discrete, we use the plug-in estimator (appendix) to estimate KL_Y .

⁶ This step is performed to break ties among data modifications, making it easy to show Type-I error is controlled.

4.1. Artificial data experiments

4.1.1. Detecting different types of shifts in isolation

In the first experiment, we set

$$P_Y^{(1)} = \text{Ber}(1/2), \quad P_Y^{(2)} = \text{Ber}(1/2 + \delta)$$

and

$$P_{X|Y}^{(1)} = \mathcal{N}(Y \cdot \mathbf{1}_d, \mathbf{I}_d), \quad P_{X|Y}^{(2)} = \mathcal{N}((Y + \gamma) \cdot \mathbf{1}_d, \mathbf{I}_d),$$

where $\text{Ber}(p)$ denotes the Bernoulli distribution with mean p , $\mathbf{1}_d$ denotes a vector of ones of size $d = 3$, \mathbf{I}_d the identity matrix of dimension $d = 3$, and $\mathcal{N}(\mu, \Sigma)$ denotes the normal distribution with mean vector μ and covariance matrix Σ . This way, δ controls the amount of response shift, while γ controls the amount of conditional shift of type 1. Indeed, it is possible to show that KL_Y positively depends only on $|\delta|$ while $\text{KL}_{X|Y}$ positively depends only on $|\gamma|$. In the second experiment, we set

$$P_X^{(1)} = \mathcal{N}(0, 1), \quad P_X^{(2)} = \mathcal{N}(\lambda, 1)$$

and

$$P_{Y|X}^{(1)} = \mathcal{N}(X, 1), \quad P_{Y|X}^{(2)} = \mathcal{N}(X + \theta, 1),$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 . In this way, λ controls the amount of feature shift, while θ controls the amount of conditional shift of type 2. Indeed, it is possible to show that KL_X positively depends only on $|\lambda|$ while $\text{KL}_{Y|X}$ positively depends only on $|\theta|$.

We vary (δ, γ) and (λ, θ) in a grid of points for experiments 1 and 2, respectively. For each point in the grid, we perform 100 Monte Carlo simulations to estimate the tests' powers, that is, the probabilities of rejecting the null hypotheses. For each pair (δ, γ) or (λ, θ) and Monte Carlo simulation, we: (i) draw training and test sets from both joint distributions, with size 2500 each; (ii) train a logistic regression⁷ model as a probabilistic classifier to estimate the Radon-Nikodym derivatives using the training sets; (iii) use the test set from the target population to estimate $\text{KL}_{X,Y}$, $\text{KL}_{X|Y}$ or $\text{KL}_{Y|X}$, and KL_Y or KL_X ; (iv) and use the test set to calculate the p -values using⁸ Algorithm 1 in the appendix setting $B = 100$; (v) reject the null hypothesis if the p -value is smaller than the level of significance $\alpha = 5\%$.

Fig. 2 shows the power estimates for each test as a function of (δ, γ) or (λ, θ) . Our procedure to test the presence of different types of dataset shift is well-behaved: the power is close to the nominal level $\alpha = 5\%$ when (δ, γ) or (λ, θ) is close to the origin, i.e., when no shift happens and grows to 1 when $\|(\delta, \gamma)\|$ or $\|(\lambda, \theta)\|$ gets larger. Moreover, our procedure could also detect types of shifts in isolation: the power of our tests increases for conditional shift (types 1 and 2) and response/feature shift detection when increasing $|\gamma|$ or $|\lambda|$ and $|\delta|$ or $|\theta|$ separately. As expected, the tests are not affected by the shifts that are not being tested at that moment.

In these experiments, we showed that the novel unified framework could reliably detect different dataset shifts in isolation.

4.1.2. Comparisons with existing approaches

Now, we compare our framework with existing methods for detecting shifts. We do this by comparing the power of the different hypotheses tests fixing $\alpha = 5\%$. For this experiment, we use the same data-generating process used in the first two experiments and set the sample sizes of training and test sets to 500. Moreover, we use 200 Monte Carlo simulations to estimate power and set $B = 100$ for Algorithm 1. When our objective is to detect response and feature shifts, we vary δ and λ but fix $\gamma = \theta = 0$; when we aim to see both types of conditional shifts, we vary γ and θ but set $\delta = \lambda = 0$.

The main alternative approach we compare our method with is the total variation (TV) approach proposed by Webb et al. [48], which also quantifies different types of dataset shift. To apply their method, we discretized the data and used our Algorithm 1 (appendix) to obtain p -values. When testing for response shift, we also include comparisons with a Z-test to compare two proportions [23], a χ^2 test [35], and a classification-based two-sample test [26]. When testing for feature shift, we also include comparisons with a Kolmogorov-Smirnov (KS) test [19,40], an MMD-based test [14, Corollary 16], and a classification-based two-sample test [26]. For the classification-based two-sample tests, we use Lopez-Paz and Oquab [26]'s formulation to obtain the p -values. Finally, when we test for conditional shifts 1 and 2, we include two instances of the local permutation test⁹ (LPT) [18] and the conditional randomization test (CRT) [4,5,3] using statistics based on the classification approach [26]. More details can be found in the appendix.

Fig. 3 shows that our method had similar power curves to the alternative approaches when testing for response and feature shift. However, when testing for both types of conditional shift, our method achieved a significantly higher power when compared with the alternative approaches.

Next, we investigate the role of the dimensionality of the feature space in the performance of the three methods used to test for feature shift, and that can be easily extended to multidimensional cases. More specifically, our goal in the example is to detect feature shifts using the settings from the second experiment of this section when $\lambda = .24$. We concatenate to X a standard Gaussian

⁷ We use Scikit-Learn's [31] default configuration with no hyperparameter tuning.

⁸ For the conditional randomization test, we train a linear regression model with Gaussian errors to estimate the conditional distribution of Y given X using the full training set.

⁹ Permuting data within each level of Y .

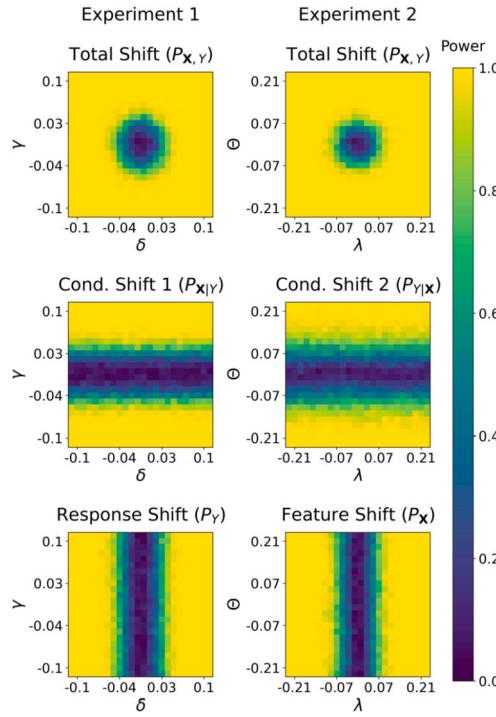


Fig. 2. Power surfaces. In these experiments, δ controls response shift, γ controls conditional shift 1, λ controls feature shift, and θ controls conditional shift 2. While the tests can control Type I errors, their power gets close to 1 when the shifts are bigger. Moreover, our procedure can detect the types of shifts in isolation. *Therefore, the novel unified framework can reliably detect different types of dataset shifts.*

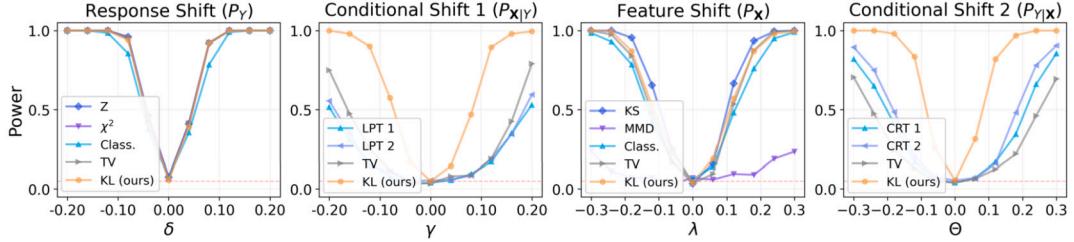


Fig. 3. Comparing our framework with existing approaches plotting their power curves when $\alpha = 5\%$. Our method had power curves similar to the alternative approaches when testing for response and feature shifts. However, our method achieved significantly higher power when testing both conditional shifts. *In summary, (i) our approach has the advantage of being unified, i.e., a single framework is used to test all hypotheses using test statistics with the same nature, and (ii) when compared to specialized tests, the novel unified framework does not lose in power when detecting marginal shifts while it wins when detecting conditional shifts.*

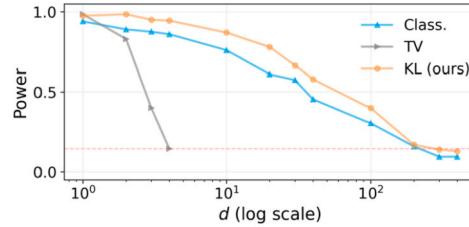


Fig. 4. Role of the feature space's dimensionality in the tests' power. The performance of our method and the classification approach when $d = 400$ is equivalent to the performance of the TV approach when $d = 4$. *Our method scales well to high-dimensional data and consistently outperforms other approaches.*

random vector (independent of the original X), ending up with an updated version of X with size d . Then, we compare the various tests in terms of their power to test $H_{0,F}$ when $\alpha = 5\%$. We use 200 Monte Carlo simulations to estimate power and set $B = 100$ for Algorithm 1 (appendix). This experiment allows us to isolate the dimensionality because the divergence between the distributions remains the same when adding this noise.

Fig. 4 indicates that the performance of our method and the classification approach does not suffer as much from increasing d compared to the TV approach. We compare the TV approach with $d \in \{1, 2, 3, 4\}$ with our approach $d \in \{1, 2, 3, 4, 10, 20\}$,

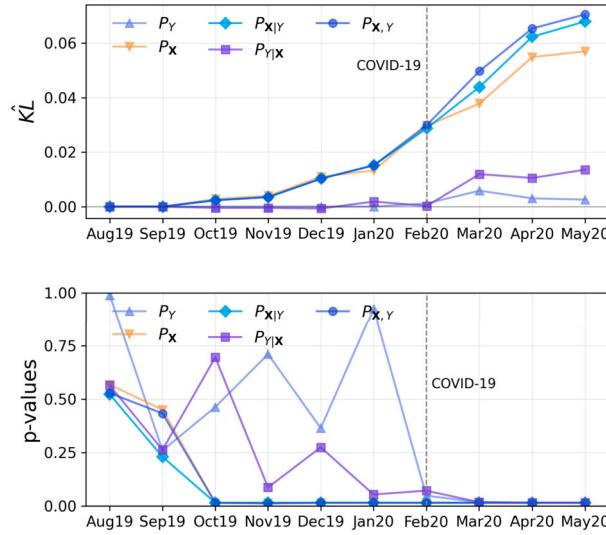


Fig. 5. Detecting different types of dataset shift using credit data from the Latin American Experian DataLab, based in Brazil. The vertical dashed line marks the beginning of the COVID-19 crisis in Brazil. We highlight the decoupling between the total shift and covariate or conditional shift type 1 after February/2020. This behavior is due to a bigger shift in the marginal and conditional distribution of Y and is possibly associated with the economic consequences of the pandemic. *Because all tests use test statistics with the same nature, we can easily compare the magnitude of the different types of shifts. The testing procedure becomes interpretable.*

$30, 40, 100, 200, 300, 400$. We stop at $d = 4$ for the TV approach because the quantity of bins increases geometrically with the number of dimensions, and if $d \geq 5$, we would expect to find less than two data points per bin. Interestingly, the performance of the TV approach when $d = 4$ is equivalent to the performance of the other methods when $d = 400$. Moreover, our method consistently outperforms the other two approaches.

In summary, (i) our approach has the advantage of being unified, i.e., a single framework is used to test all hypotheses using test statistics with the same nature while maintaining good power; (ii) our method scales well to high-dimensional data and consistently outperforms other approaches.

4.2. Real data experiments

4.2.1. Insights from credit data

In this experiment, we use our method to extract insights into how probability distributions can differ in a financial application. The dataset used in this experiment is a credit scoring dataset and was kindly provided by the Latin American Experian DataLab, based in Brazil. It contains financial data of one million Brazilians collected every month going from August/2019 to May/2020.

The features in this dataset are related to past financial data, e.g., amount of loans and credit card bills not paid on time, and the *label variable* informs whether a consumer will *delay a debt payment* for 30 days in the *next three months*, i.e., we have a binary classification problem. In this experiment, we kept 20k random data points each month, with 20% of them going to the test set. Also, we kept the top 5 most essential features of the credit risk prediction model. These specific features are related to payment punctuality for credit card bills, the number of active consumer contracts, and the monetary values involved. We used CatBoost [34] both to estimate the Radon-Nikodym derivative and the conditional distribution of $Y|X$.

The results in Fig. 5 indicate increasing covariate, conditional shift type 1, and total dataset shift from the beginning. This is expected because the features contain information about how people use their credit (e.g., loan amount, credit card use). The way people use their credit is a function of changes in the economy that can occur rapidly, such as fluctuating inflation/interest/exchange rates, extra expenses due to holidays, etc. From February 2020 onward, i.e., labels relative to months post-March 2020, it is possible to notice a decoupling between shift curves in the coming months after the first official COVID-19 case detected in Brazil and the beginning of the economic crisis. The decoupling means that a more significant share of the total shift is due to a shift in the marginal and conditional distribution of Y . We speculate that this decoupling is due to measures taken by banks and credit bureaus to help consumers during the pandemic. Some measures include, but are not limited to, longer payment intervals and lower interest rates.

To conclude, because all test statistics have the same nature, we can easily compare the magnitude of the different types of shifts. The testing procedure becomes interpretable.

4.2.2. Using dataset shift diagnostics to improve predictions

In this experiment, we evaluate our method as a guide for dataset shift adaptation using the MNIST and USPS datasets [22,50]. Both datasets contain images, i.e., pixel intensities, and labels for the same ten digits (i.e. 0 to 9). Our interest is (i) to use our framework to quantify and formally test the presence of all types of dataset shift using the MNIST distribution as the source and a mixture between MNIST and USPS (with increasing proportions of USPS participation) as target distributions and then (ii) adapt our predictors using the insights given by our diagnostics to achieve better out-of-sample performance. We aim to show how detecting

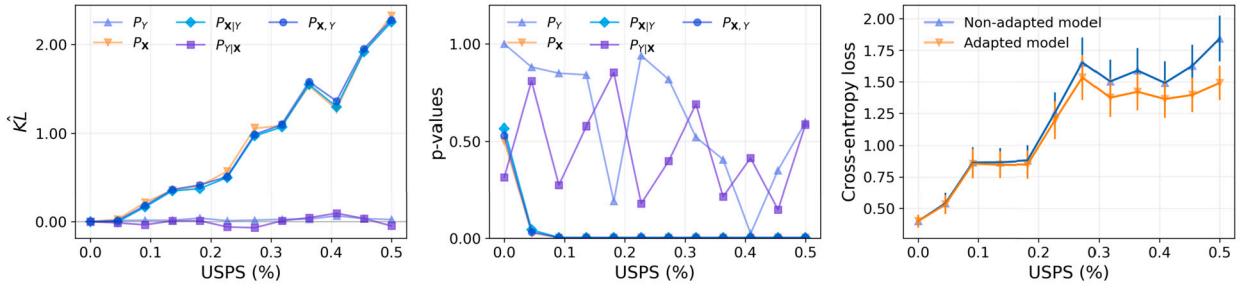


Fig. 6. Detecting different types of dataset shift using MNIST and USPS data and then adapting a predictor. We define the MNIST distribution ($\text{USPS} \% = 0$) as the source distribution and the mixed distributions ($\text{USPS} \% > 0$) as multiple target distributions. Response and conditional shifts of type 2 (shift in $P_{Y|X}$) are not evident, while others are. The second plot indicates that adapting for feature shift should be enough (as $P_{Y|X}$ seems static). In the third plot, we compare the performance of two logistic regression models trained using pure MNIST samples with one adapted for feature shift. The adapted model gives better predictions. *Thus, by properly leveraging the insights given by our framework, one can improve the predictive power of the supervised learning model.*

specific types of shift help practitioners correct their models. In this experiment, we use 256 features (pixel intensities of 16×16 images) and partition the data in 12 smaller disjoint datasets of size 3.1k in a way that the proportion of USPS samples increases linearly from 0% to 50%. We split each dataset with 10% of the samples to test and use CatBoost [34] to estimate R-N derivatives and conditional distribution. We use the first dataset, with only MNIST samples, as our baseline and compare it to the mixed datasets.

The plots in Fig. 6 indicate that the total dataset shift (shift in $P_{X,Y}$), conditional shift 1 (shift in $P_{X|Y}$), and feature shift (shift in P_X) are promptly detected while response shift (shift in P_Y) and conditional shift 2 (shift in $P_{Y|X}$) are not. That observation is consistent with that (i) the distribution of Y is similar in both MNIST and USPS populations, which also implies that (ii) two similar pixel configurations should induce similar posterior distributions of labels regardless of the origin distribution. Finally, the plots indicate that adapting for feature shifts should be enough to achieve better predictions on the target domain. Indeed, that is the case here – in the third plot, we compare the performance of two logistic regression models⁶. Both are trained using pure MNIST samples, but one is correct for feature shift using importance weighting [44]. *The weights are obtained via the classifier used to estimate KL_X , with no need to fit an extra model for the weights.*

In this experiment, we show that, by properly leveraging the insights given by our framework, one can improve the predictive power of the supervised learning model.

4.2.3. Detecting shifts with deep models

We use our framework to detect shifts in image and text datasets using deep learning models as classifiers to estimate the Radon-Nikodym derivative and the conditional distribution of $Y|X$. We use large pre-trained models as feature extractors, freezing all the layers except the output one, given by a logistic regression model. The pre-trained models are VGG-16 [39] for images and XLM-ROBERTa [6] for texts. We use the Tiny ImageNet [9,1] and CIFAR-10 [20] as image¹⁰ datasets and “Amazon Fine Food Reviews,” available on Kaggle, as our text dataset. Then, the first two datasets we use are composed of RGB images from $K = 10$ different classes.¹¹ In contrast, the third dataset is composed of product reviews in the form of short texts and a rating, varying from 0 to 4, given by consumers, thus having $K = 5$ classes.¹² The sample sizes are 30k data points with 10% going to test.

We derive the source and target datasets from the original datasets as follows. First, we fix $\delta \in (0, .5)$ and then create a list `LIST` of K numbers (one for each class) where the first element of the list is δ , the last is $1 - \delta$, and the intermediate ones are given by linear interpolation of δ and $1 - \delta$. Then, for each $k \in \{0, \dots, K - 1\}$, we randomly select `LIST`[k] of the samples of class k to be in the source dataset, while the rest goes to the target dataset. In this way, we explicitly introduce label shift (shift in P_Y but not in $P_{X|Y}$), which is expected to induce feature shift and conditional shift type 2 as well. After we have data from both populations (source and target), we detect the shifts as usual. We repeated the same procedure for all $\delta \in \{.5, .45, .4, .35, .3\}$ and for five different random seeds.

The results are shown in Fig. 7, where the lines represent averages across repetitions and error bars give the standard deviations. We were able to detect¹³ all types of shift except conditional shift 1 using our approach. This result was expected because, given the class, the distribution of the features must not be affected by how we introduced the shift. We also run results considering the classification-based two-sample test proposed by [26] to test shifts related to joint/marginal distributions and CRT 1 and LPT 1, explained in detail in the appendix, to test for conditional shifts. For the benchmarks, we cannot detect a conditional shift of type 2. A lower power of the benchmarks compared to our method can explain that fact (see artificial data experiments). Furthermore, we cannot compare shift magnitudes using benchmark approaches since their statistics do not have the same nature. We report KL

¹⁰ Tiny ImageNet contains 64×64 images from 200 classes while CIFAR-10 contains 32×32 images from 10 classes.

¹¹ We group, in increasing order, the classes from Tiny ImageNet in 10 meta-classes.

¹² Regarding the Amazon dataset, we subsampled the data to guarantee all the classes have roughly the same number of examples.

¹³ Except when δ is close to .5 (small or no shift).

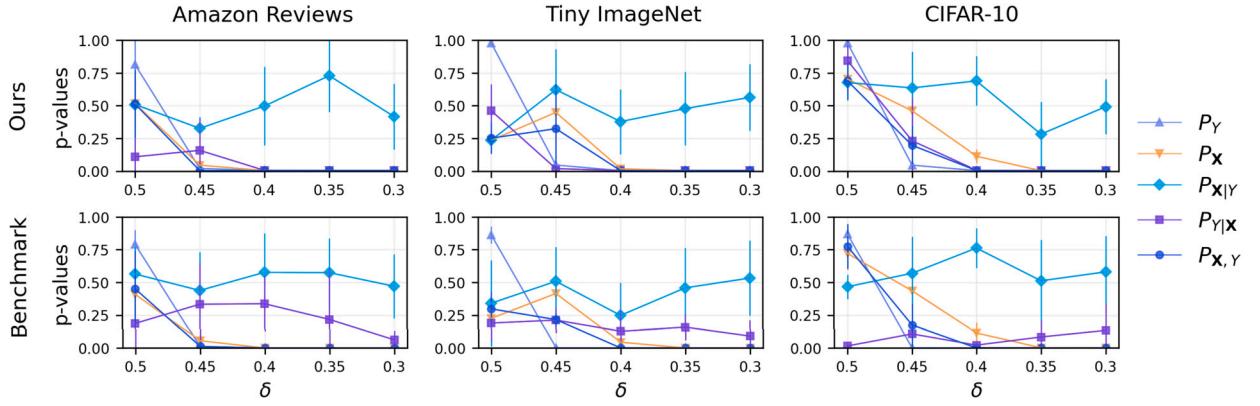


Fig. 7. Detecting different types of dataset shift using complex high-dimensional data. For many values of δ , we could detect all types of shift except conditional shift 1, which is expected because we introduced response shift in isolation ($\delta \neq .5$ is equivalent to label shift). The benchmark approaches cannot detect a conditional shift of type 2, which can be explained by lower power. *Then, our framework can accurately detect the correct shifts even in complex datasets and using deep models.*

estimates of *DetectShift* in the appendix. We also left in the appendix an extra (and similar) experiment in which our framework obtained similar p-values compared to the benchmarks.

In summary, we show that our framework can accurately detect the correct shifts even in complex datasets and using deep models.

5. Discussion

The proposed *DetectShift* framework quantifies and formally tests different dataset shifts, including response, feature, and conditional shifts. Our approach sheds light not only on *if* a prediction model should be retrained/adapted but also on *how*, enabling the practitioners to tackle shifts in their data objectively and efficiently. Our method is versatile and applicable across various data types and supervised tasks. Unlike our framework, existing dataset shift detection methods are only designed to detect specific types of shift or cannot formally test their presence, sometimes even requiring both labels and features to be discrete or low-dimensional.

Our experiments have provided compelling evidence of the effectiveness of our framework. Through artificial data experiments, we have demonstrated the versatility of *DetectShift* in isolating different types of dataset shifts and leveraging supervised learning to construct powerful tests. Furthermore, our framework has shown remarkable performance in higher dimensions, surpassing possible benchmarks. In real data experiments, we have illustrated how *DetectShift* can extract valuable insights from data, assisting practitioners in adapting their models to changing distributions. Moreover, our approach seamlessly integrates with deep learning models, making it applicable to real-world scenarios. These findings highlight the significant contributions of our framework, showcasing its efficacy in detecting dataset shifts and underscoring its potential for practical applications in various domains.

6. Final remarks and possible extensions

In conclusion, we reflect on key facets and potential limitations inherent to our framework.

Why and how to obtain better Radon-Nikodym derivatives estimates? While Type I error control is not affected by Radon-Nikodym derivatives estimation (as the theoretical results show), the power of the tests is. This happens because the classifiers' performance in predicting probabilities directly influences how well we approximate the KL divergence and, consequently, can detect shifts in the data. Thus, we can use the cross-entropy (CE) loss on a validation set to pick the best classifier.

Modular framework and alternative statistics choices. Our hypothesis tests are agnostic to the choice test statistics; therefore, our framework is modular. We consider the KL statistics desirable, among other things, because we have the additivity property (Proposition 2), making the analyses more interpretable. A natural alternative to the KL statistics, when we do not want to favor one of the distributions, is the symmetrized KL, for example.

Dealing with streaming data. While the primary design of our framework does not cater to streaming data, it possesses the flexibility for adaptation in such settings. One way of doing this is to group the data into batches (e.g., every hour, day, month) and then apply the proposed approach to compare two or more data batches. Multiple testing methods [36] can be used along with our framework if the practitioner desires.

Labeled data in the target domain. Even though there are many situations in which at least some labeled data are available for the target domain, there are cases in which that is not true. We recognize this scenario as a limitation, warranting further exploration in subsequent research.

CRediT authorship contribution statement

Felipe Maia Polo: Conceptualization, Methodology, Writing, Experiments, Software. **Rafael Izbicki:** Conceptualization, Methodology, Writing. **Evanildo Gomes Lacerda Jr:** Credit scoring experiments. **Juan Pablo Ibieta-Jimenez:** Credit scoring experiments. **Renato Vicente:** Conceptualization, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Resources

Every dataset is available besides the credit scoring dataset, which is confidential.

Our package for dataset shift detection can be found in <https://github.com/felipemaiapolo/detectshift>

The source code and data used in this paper can be found in https://github.com/felipemaiapolo/dataset_shift_diagnostics

Acknowledgements

We sincerely thank the “Information Sciences” editor and anonymous reviewers for their invaluable contributions to this paper. We are grateful for their time and dedication in carefully evaluating our manuscript and providing valuable suggestions for improvement.

We are thankful for the credit dataset provided by the Latin American Experian Datalab, the Serasa Experian DataLab. We want to thank the developers of Grammarly and ChatGPT, which were helpful tools when polishing our text.

FMP is grateful for the financial support of CNPq (32857/2019-7) and Fundunesp/Advanced Institute for Artificial Intelligence (AI2) (3061/2019-CCP) during his master’s degree at the University of São Paulo (USP). Part of this work was written when FMP was at USP.

RI is grateful for the financial support of CNPq (309607/2020-5 and 422705/2021-7) and FAPESP (2019/11321-9).

Appendix A. Methodology

A.1. Alternative estimator for KL_Y when Y is discrete

Assume \mathcal{Y} is the range of Y in the target domain, where \mathcal{Y} is finite. Define $p_y^{(i)} = P_Y^{(i)}(\{y\})$, for $i = 1, 2$. Then, we can write $KL_Y = \sum_{y \in \mathcal{Y}} p_y^{(2)} \log \frac{p_y^{(2)}}{p_y^{(1)}}$. Having observed two datasets (in practice, we use the test datasets), $D^{(1)}$ and $D^{(2)}$, we define $\hat{p}_y^{(i)}$ to be the relative frequency of the label y in dataset i . Then, a plug-in estimator for KL_Y is given by $\widehat{KL}_Y = \sum_{y \in \mathcal{Y}} \hat{p}_y^{(2)} \log \frac{\hat{p}_y^{(2)}}{\hat{p}_y^{(1)}}$. This estimator is consistent.

A.2. Algorithm to obtain p-values

Algorithm 1: DetectShift: obtaining p-values.

Input: (i) Hypothesis to be tested and respective test statistic T , (ii) Test set D^{te} , (iii) number of iterations $B \in \mathbb{N}$, (iv) conditional distribution $Q_{Y|X}$ (in case of testing for conditional shift type 2), (v) Gaussian noise variance ν ;
Output: p-value $p = p(D^{te})$;

```

1 Initialize  $C = 0$  and obtain  $D_2^{te} = \{(\mathbf{X}_k, Y_k, Z_k) \in D^{te} : Z_k = 2\}$ ;
2 Get  $T_0 = T(D_2^{te}) + N(0, \nu)$ ;
3 for  $j$  in  $\{1, \dots, B\}$  do
4   if Testing for response shift, feature shift, or dataset shift then
5     Draw a random permutation  $\pi = (\pi_1 \dots \pi_{|D^{te}|})$  of natural numbers from 1 to  $|D^{te}|$ 
6     Set  $\widetilde{D}_2^{te(j)} = \{(\mathbf{X}_k, Y_k, Z_{\pi_k}) : (\mathbf{X}_k, Y_k, Z_k) \in D^{te} \text{ and } Z_{\pi_k} = 2\}$ 
7   else if Testing for conditional shift (type 1) then
8     Let  $\mathcal{Y}$  be a finite set which  $Y$  takes values
9     for  $y$  in  $\mathcal{Y}$  do
10      Get  $D^{(y)} = \{(\mathbf{X}_k, Y_k, Z_k) \in D^{te} : Y_k = y\}$ 
11      Draw a random permutation  $\pi^{(y)} = (\pi_1^{(y)} \dots \pi_{|D^{(y)}|}^{(y)})$  of natural numbers from 1 to  $|D^{(y)}|$ 
12      Set  $\widetilde{D}_2^{te(j)} = \bigcup_{y \in \mathcal{Y}} \{(\mathbf{X}_k, Y_k, Z_{\pi_k^{(y)}}) : (\mathbf{X}_k, Y_k, Z_k) \in D^{te} \text{ and } Z_{\pi_k^{(y)}} = 2\}$ 
13   else if Testing for conditional shift (type 2) then
14     Sample  $\tilde{Y}_k | \mathbf{X}_k \sim Q_{Y|X=\mathbf{X}_k}$ , for each  $\mathbf{X}_k$  from elements of  $D_2^{te}$ 
15     Set  $\widetilde{D}_2^{te(j)} = \{(\mathbf{X}_k, \tilde{Y}_k, Z_k) : (\mathbf{X}_k, Y_k, Z_k) \in D_2^{te}\}$ 
16   Get  $T_j = T(\widetilde{D}_2^{te(j)}) + N(0, \nu)$ ;
17   Update  $C = C + \mathbb{I}[T_0 \leq T_j]$ ;
18    $p = \frac{C+1}{B+1}$ 
19 return  $p$ .

```

In practice, $\nu > 0$ can be a very small number, e.g., 10^{-10} .

A.3. Proofs

The results presented in this section are not original; however, we decided to include proofs to make this work self-contained.

Proposition 2. Let KL_Y , KL_X , $\text{KL}_{Y|X}$, $\text{KL}_{X|Y}$, $\text{KL}_{X,Y}$ be defined as they were in Section 3.1. Then

$$\text{KL}_{X,Y} = \text{KL}_{Y|X} + \text{KL}_X = \text{KL}_{X|Y} + \text{KL}_Y$$

Proof. In this proof, we assume the distributions are absolutely continuous with respect to the Lebesgue measure. For a more general proof, see Polanskiy and Wu [33] (Theorem 2.13).

We start showing $\text{KL}_{X,Y} = \text{KL}_{Y|X} + \text{KL}_X$. Let $p_{X,Y}^{(1)}$ and $p_{X,Y}^{(2)}$ be the joint densities of $P_{X,Y}^{(1)}$ and $P_{X,Y}^{(2)}$ with respect to the Lebesgue measure. Then

$$\begin{aligned} \text{KL}_{X,Y} &= \int p_{X,Y}^{(2)}(\mathbf{x}, y) \log \frac{p_{X,Y}^{(2)}(\mathbf{x}, y)}{p_{X,Y}^{(1)}(\mathbf{x}, y)} d(\mathbf{x}, y) \\ &= \int p_{Y|X}^{(2)}(y|\mathbf{x}) p_X^{(2)}(\mathbf{x}) \log \frac{p_{Y|X}^{(2)}(y|\mathbf{x}) p_X^{(2)}(\mathbf{x})}{p_{Y|X}^{(1)}(y|\mathbf{x}) p_X^{(1)}(\mathbf{x})} d(\mathbf{x}, y) \\ &= \int p_X^{(2)}(\mathbf{x}) \int p_{Y|X}^{(2)}(y|\mathbf{x}) \log \frac{p_{Y|X}^{(2)}(y|\mathbf{x})}{p_{Y|X}^{(1)}(y|\mathbf{x})} dy d\mathbf{x} + \int p_X^{(2)}(\mathbf{x}) \log \frac{p_X^{(2)}(\mathbf{x})}{p_X^{(1)}(\mathbf{x})} d\mathbf{x} \\ &= \text{KL}_{Y|X} + \text{KL}_X \end{aligned}$$

We can show that $\text{KL}_{X,Y} = \text{KL}_{X|Y} + \text{KL}_Y$ analogously. \square

Proposition 3. Let $p(D^{te})$ be a p -value obtained from Algorithm 1 (appendix) with fixed T and $Q_{Y|X}$. Then, for every $\alpha \in (0, 1)$,

- Under $H_{0,D}$, $H_{0,F}$, $H_{0,R}$, and $H_{0,C1}$ (if Y is discrete),

$$\mathbb{P}(p(D^{te}) \leq \alpha) \leq \alpha$$

- Under $H_{0,C2}$

$$\mathbb{P}(p(D^{te}) \leq \alpha) \leq \alpha + \mathbb{E}_{\bar{P}_X^{(2)}} [d_{\text{TV}}(\bar{Q}_{Y|X}, \bar{P}_{Y|X})],$$

where d_{TV} is the total variation distance, $\bar{P}_X^{(2)} := \prod_{(\mathbf{x}_k, y_k, z_k) \in D_2^{te}} P_{\mathbf{x}_k}^{(2)}$, $\bar{Q}_{Y|X} := \prod_{(\mathbf{x}_k, y_k, z_k) \in D_2^{te}} Q_{Y|\mathbf{x}=\mathbf{x}_k}$, and $\bar{P}_{Y|X} = \prod_{(\mathbf{x}_k, y_k, z_k) \in D_2^{te}} P_{Y|\mathbf{x}=\mathbf{x}_k}^{(2)}$.

Proof. We start deriving the result for $H_{0,F}$ using the theory behind permutation tests for two-sample/independence tests (see Lehmann et al. [23, Chapter 15] for more details). The results for $H_{0,D}$ and $H_{0,R}$ obtained analogously. Define $\tilde{D}_2^{te(0)} := D_2^{te}$ and $T_j := T(\tilde{D}_2^{te(j)}) + N_j$, for all j , where N_j 's are i.i.d. $N(0, \nu)$. For the rest of the proof, T_j 's will always be noisy versions of the respective $T(\tilde{D}_2^{te(j)})$'s.

Recall that

$$p(D^{te}) = \frac{\sum_{j=0}^B \mathbb{I}[T_0 \leq T_j]}{B+1}$$

See that $(B+1)p(D^{te})$ gives the rank of T_0 among all T_j 's, that is, if $(B+1)p(D^{te}) = k$ it means that there are k values of j such that $T_0 \leq T_j$. If $P_X^{(1)} = P_X^{(2)}$ and the permutations are drawn uniformly from the set of all permutations of $\{1, \dots, |D^{te}|\}$, then T_0, \dots, T_B are exchangeable. Because we add N_j 's in T_j 's, we can guarantee that all T_j 's are different with probability 1. Consequently, $U_{B+1} := (B+1)p(D^{te})$ is uniformly distributed in $\{1, \dots, B+1\}$. Then,

$$\begin{aligned} \mathbb{P}(p(D^{te}) \leq \alpha) &= \mathbb{P}((B+1)p(D^{te}) \leq (B+1)\alpha) \\ &= \mathbb{P}(U_{B+1} \leq (B+1)\alpha) \\ &= \mathbb{P}(U_{B+1} \leq \lfloor (B+1)\alpha \rfloor) \\ &= \frac{\lfloor (B+1)\alpha \rfloor}{B+1} \\ &\leq \frac{(B+1)\alpha}{B+1} \end{aligned}$$

$$= \alpha$$

where $\lfloor \cdot \rfloor$ is the floor function.

The result for $H_{0,C1}$ is obtained similarly. Assume that $P_{X|Y=y}^{(1)} = P_{X|Y=y}^{(2)}$ for every $y \in \text{supp}(P_Y^{(2)})$. We can ignore values of y not in $\text{supp}(P_Y^{(2)})$ because $p(D^{te})$ does not depend on them. Given that we permute samples within each level of Y , we have that T_0, \dots, T_B are exchangeable, implying that (by our derivations above)

$$\mathbb{P}(p(D^{te}) \leq \alpha) \leq \alpha.$$

The result for $H_{0,C2}$ is based on Theorem 4 of Berrett et al. [4]. For each $j \in \{1, \dots, B\}$, denote $\widetilde{D}_2^{te(j)} = \{(\mathbf{X}_k, \widetilde{Y}_k, Z_k) : (\mathbf{X}_k, Y_k, Z_k) \in D_2^{te}\}$ where $\widetilde{Y}_k | \mathbf{X}_k \sim Q_{Y|X=X_k}$, for each \mathbf{X}_k from elements of D_2^{te} . Define $\widetilde{D}_2^{te(0)} := D_2^{te}$ and let $\widetilde{D}_2^{te(B+1)}$ be one extra dataset built in the same way as all others $\widetilde{D}_2^{te(j)}$ considering $j \in \{1, \dots, B\}$. See that all datasets $\widetilde{D}_2^{te(j)}$ considering $j \in \{0, \dots, B+1\}$ share the same values for the covariates; let \mathbb{X} denote a random matrix containing such values for all samples. Because the p-values only depend on the data coming from¹⁴ D_2^{te} (and not on the full test set D^{te}), we have that

$$\mathbb{P}(p(D^{te}) \leq \alpha | \mathbb{X}) \leq \mathbb{P}(\tilde{p}(D^{te}) \leq \alpha | \mathbb{X}) + d_{\text{TV}}(\bar{Q}_{Y|\mathbf{X}}, \bar{P}_{Y|\mathbf{X}})$$

where $\tilde{p}(D^{te})$ is the p-value calculated using $\widetilde{D}_2^{te(B+1)}$ instead of $\widetilde{D}_2^{te(0)}$. This step is justified by the definition of the total variation distance and by a conditional independence argument [4] (given \mathbb{X} , all data replications are conditionally independent).

By construction, T_1, \dots, T_{B+1} are exchangeable given \mathbb{X} , we have that $\mathbb{P}(\tilde{p}(D^{te}) \leq \alpha | \mathbb{X}) \leq \alpha$. Then, taking expectations on both sides of the inequality above, we get

$$\mathbb{P}(p(D^{te}) \leq \alpha) \leq \alpha + \mathbb{E}_{\bar{P}_{\mathbf{X}}^{(2)}} [d_{\text{TV}}(\bar{Q}_{Y|\mathbf{X}}, \bar{P}_{Y|\mathbf{X}})] \quad \square$$

Realize that the last result is true even if $H_{0,C2}$ does not hold. However, if $H_{0,C2}$ is not true, we do not expect $\mathbb{E}_{\bar{P}_{\mathbf{X}}^{(2)}} [d_{\text{TV}}(\bar{Q}_{Y|\mathbf{X}}, \bar{P}_{Y|\mathbf{X}})]$ to be small.

Appendix B. Experiments

B.1. Comparisons with existing approaches (more details)

Some details about the experiments were omitted from the main text: (i) how we choose $Q_{Y|\mathbf{X}}$ when testing for conditional shift 2; (ii) how the local permutation tests (LPT) and conditional randomization tests (CRT) alternatives work and what statistic they use.

B.1.1. How do we choose $Q_{Y|\mathbf{X}}$ in this set of experiments?

Given that our main objective in this experiment is comparing the power of different tests that use the same approximated distribution $Q_{Y|\mathbf{X}} = P_{Y|\mathbf{X}}^{(0)} = \mathcal{N}(X, 1)$, where $P_{Y|\mathbf{X}}^{(0)}$ is the conditional distribution of $Y|X$ under $H_{0,C2}$.

B.1.2. How do the LPT and CRT alternative tests work, and what statistics do they use?

We start explaining the two CRT alternative tests, which work in the same way but have different test statistics.

1. We split our dataset $\mathcal{D} = \{(X_i, Y_i, Z_i)\}_{i=1}^n$ in a training set $\mathcal{D}^{tr} = \{(X_i, Y_i, Z_i)\}_{i=1}^{n^{tr}}$ and a test set $\mathcal{D}^{te} = \mathcal{D} - \mathcal{D}^{tr}$;
2. We build an artificial training set $\widetilde{\mathcal{D}}^{tr} = \{(X_i, \widetilde{Y}_i, Z_i)\}_{i=1}^{n^{tr}}$, where $\{\widetilde{Y}_i\}$ are sampled from $Q_{Y|\mathbf{X}}$;
3. We train a probabilistic classifier \hat{h} to distinguish samples from \mathcal{D}^{tr} and $\widetilde{\mathcal{D}}^{tr}$, where the original set receives label 1 and the artificial data receives label 0;
4. For $B \in \mathbb{N}$, we test each of the hypotheses of interest by computing a p-value of the form

$$p(D^{te}) = \frac{1 + \sum_{j=1}^B \mathbb{I}\left[T(D^{te}) \leq T\left(\widetilde{D}^{te(j)}\right)\right]}{B + 1},$$

where T is a test statistic depending on \hat{h} and each $\widetilde{D}^{te(j)}$ is obtained sampling different labels from $Q_{Y|\mathbf{X}}$. When

$$T(D^{te}) = \frac{1}{|\mathcal{D}^{te}|} \sum_{(X, Y, Z) \in \mathcal{D}^{te}} \mathbb{I}[\hat{h}(X, Y, Z) > 1/2]$$

we have CRT 1. When

$$T(D^{te}) = \frac{1}{|\mathcal{D}^{te}|} \sum_{(X, Y, Z) \in \mathcal{D}^{te}} \hat{h}(X, Y, Z)$$

¹⁴ If that was not true, the following inequality would hold if we condition not only on \mathbb{X} but also on a matrix composed of Z 's (as done in Berrett et al. [4]).

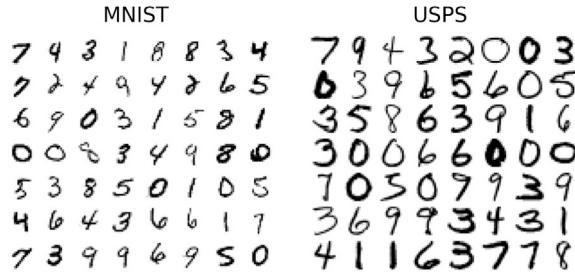


Fig. 8. Image samples from MNIST and USPS.

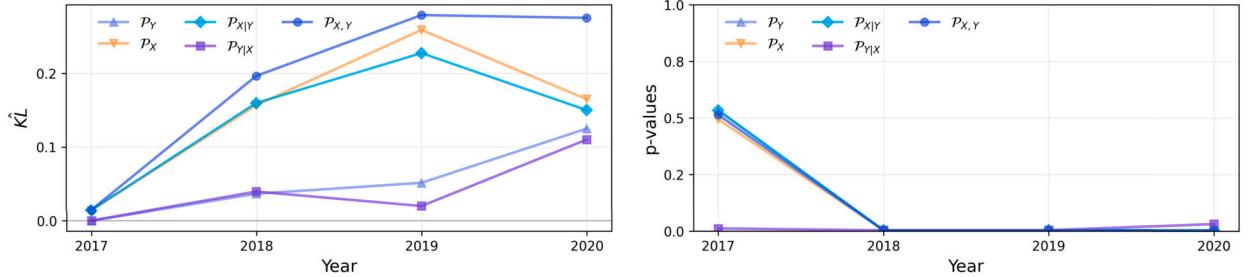


Fig. 9. Detecting different types of dataset shift in a regression experiment. Using data from ENEM, the “Brazilian SAT,” we do a similar analysis compared to the credit one, comparing the probability distributions of 2018, 2019, and 2020 with the one in 2017. In this experiment, Y is given by the student’s math score on the logarithmic scale, while X comprises students’ personal and socioeconomic features. It is possible to see that we detected all kinds of shifts in every moment after 2017. This result indicates that a model trained in 2017 might not generalize well to other years, and practitioners may consider retraining their models.

we have CRT 2.

The LPT alternative tests work in the same way. The only difference is that the “artificial” variables $\{\tilde{X}_i\}$ are obtained via local permutation within Y levels.

In this experiment, \hat{h} represents a CatBoost classifier. When \hat{h} is a logistic regressor (like in the KL tests), at least one test is too conservative.

B.2. Digits experiment

MNIST samples tend to have more white pixels than USPS (Fig. 8); thus, the distributions of X are different in both datasets.

B.3. A regression experiment

We present a regression experiment using data from 2017, 2018, 2019, and 2020 of ENEM,¹⁵ the “Brazilian SAT.” In each of the years, Y is given by the students’ math score in a logarithmic scale while X is composed of six of their personal and socioeconomic features: gender, race, school type (private or public), mother’s education, family income and the presence of a computer at home. We randomly subsample the data in each one of the years to 30k data points with 10% of them going to the test portion and then use the CatBoost algorithm to estimate the Radon-Nikodym derivative and the conditional distribution of $Y|X$. When estimating the distribution of $Y|X$, we first fit a regressor to predict Y given X , and then, using a holdout set, we fit a Gaussian model on the residuals. When testing for a shift in the $X|Y$ distribution, we discretize Y in 10 bins, evenly splitting the data. Even though we use the binned version of Y to get the p -value, we report $\hat{KL}_{X|Y}$ in the first panel of Fig. 9. In this experiment, we do a similar analysis compared to the credit one, comparing the probability distributions of 2018, 2019, and 2020 with the one in 2017. From Fig. 9, it is possible to see that we detected all kinds of shifts after 2017. This result indicates that a model trained in 2017 might not generalize well to other years, and practitioners may consider retraining their models from scratch using more recent data.

B.4. Detecting shifts with deep models

B.4.1. Extra plot

See Fig. 10.

¹⁵ Data extracted from <https://www.gov.br/inep/pt-br/acesso-a-informacao/dados-abertos/microdados/enem>.

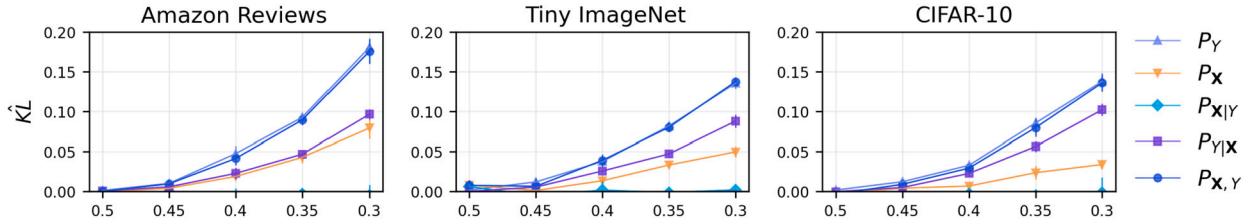


Fig. 10. Detecting different types of dataset shift using complex high-dimensional data. For many values of δ , we could detect all types of shift except conditional shift 1, which is expected because we introduced response shift in isolation ($\delta \neq .5$ is equivalent to label shift).

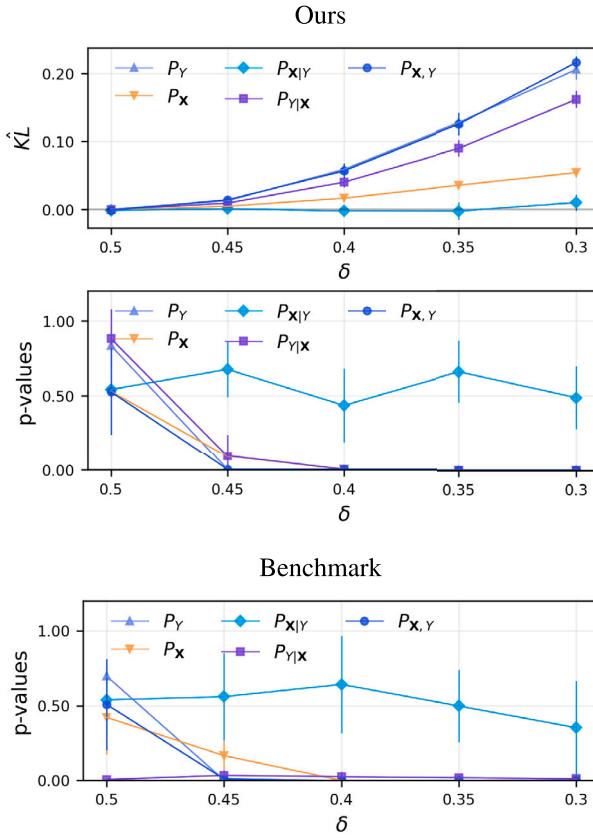


Fig. 11. Detecting different types of dataset shift using Stack Overflow text data. For many values of δ , we could detect all types of shift except concept shift 1, which is expected because we introduced label shift in isolation ($\delta \neq .5$ is equivalent to label shift).

B.4.2. Extra results

We include one extra experiment using the “60k Stack Overflow Questions with Quality Rating” dataset¹⁶ [2] (Fig. 11). This dataset has three classes of rating from Stack Overflow questions. We repeat the same procedure to obtain results for the other deep-learning experiments.

B.5. Experiments running time

All the experiments were run in a MacBook Air (M1, 2020) 16 GB, except for the credit analysis experiment in an 80 CPUs Intel Xeon Gold 6148 cluster. We consider one iteration as all the steps needed to compute all the p -values used for a specific experiment. In the artificial data experiments, on average, each iteration performed by our framework took less than 1s, while in real data experiments, each iteration took less than 200s.

¹⁶ <https://www.kaggle.com/datasets/imoore/60k-stack-overflow-questions-with-quality-rate>.

References

- [1] Zoheb Abai, Nishad Rajmalwar, Densenet models for tiny ImageNet classification, arXiv preprint, arXiv:1904.10429, 2019.
- [2] Issa Annamoradnejad, Jafar Habibi, Mohammadamin Fazli, Multi-view approach to suggest moderation actions in community question answering sites, *Inf. Sci.* (ISSN 0020-0255) 600 (2022) 144–154, <https://doi.org/10.1016/j.ins.2022.03.085>, <https://www.sciencedirect.com/science/article/pii/S0020025522003127>.
- [3] Alexis Bellot, Mihaela van der Schaar, Conditional independence testing using generative adversarial networks, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [4] Thomas B. Berrett, Yi Wang, Rina Foygel Barber, Richard J. Samworth, The conditional permutation test for independence while controlling for confounders, *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 82 (1) (2020) 175–197.
- [5] Emmanuel Candès, Yingying Fan, Lucas Janson, Jinchi Lv, Panning for gold: ‘model-x’ knockoffs for high dimensional controlled variable selection, *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 80 (3) (2018) 551–577.
- [6] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, Veselin Stoyanov, Unsupervised cross-lingual representation learning at scale, in: *ACL*, 2020.
- [7] Kyle Cranmer, Juan Pavez, Gilles Louppe, Approximating likelihood ratios with calibrated discriminative classifiers, arXiv preprint, arXiv:1506.02169, 2015.
- [8] Niccolò Dalmasso, David Zhao, Rafael Izbicki, Ann B. Lee, Likelihood-free frequentist inference: bridging classical statistics and machine learning in simulation and uncertainty quantification, arXiv preprint, arXiv:2107.03920, 2021.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, Li Fei-Fei, ImageNet: a large-scale hierarchical image database, in: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2009, pp. 248–255.
- [10] Samuel G. Finlayson, Adarsh Subbaswamy, Karandeep Singh, John Bowers, Annabel Kupke, Jonathan Zittrain, Isaac S. Kohane, Suchi Saria, The clinician and dataset shift in artificial intelligence, *N. Engl. J. Med.* 385 (3) (2021) 283.
- [11] Peter E. Freeman, Rafael Izbicki, Ann B. Lee, A unified framework for constructing, tuning and assessing photometric redshift density estimates in a selection bias setting, *Mon. Not. R. Astron. Soc.* 468 (4) (2017) 4556–4565.
- [12] Joao Gama, Pedro Medas, Gladys Castillo, Pedro Rodrigues, Learning with drift detection, in: *Brazilian Symposium on Artificial Intelligence*, Springer, 2004, pp. 286–295.
- [13] Tom Ginsberg, Zhongyuan Liang, Rahul G. Krishnan, A learning based hypothesis test for harmful covariate shift, in: *The Eleventh International Conference on Learning Representations*, 2023, <https://openreview.net/forum?id=rdfgqiwz7lZ>.
- [14] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, Alexander Smola, A kernel two-sample test, *J. Mach. Learn. Res.* 13 (1) (2012) 723–773.
- [15] Rafael Izbicki, Ann B. Lee, Converting high-dimensional regression to high-dimensional conditional density estimation, *Electron. J. Stat.* 11 (2) (2017) 2800–2831.
- [16] Rafael Izbicki, Ann B. Lee, Peter E. Freeman, Photo-z estimation: an example of nonparametric conditional density estimation under selection bias, *Ann. Appl. Stat.* 11 (2) (2017) 698–724.
- [17] Sooyong Jang, Sangdon Park, Insup Lee, Osbert Bastani, Sequential covariate shift detection using classifier two-sample tests, in: *International Conference on Machine Learning*, PMLR, 2022, pp. 9845–9880.
- [18] Ilmun Kim, Matey Neykov, Sivaraman Balakrishnan, Larry Wasserman, Local permutation tests for conditional independence, arXiv preprint, arXiv:2112.11666, 2021.
- [19] A.N. Kolmogorov, Sulla determinazione empirica di una legge di distribuzione, *G. Ist. Ital. Attuari* 4 (1933) 83–91.
- [20] Alex Krizhevsky, Geoffrey Hinton, et al., Learning multiple layers of features from tiny images, 2009.
- [21] Solomon Kullback, Richard A. Leibler, On information and sufficiency, *Ann. Math. Stat.* 22 (1) (1951) 79–86.
- [22] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* 86 (11) (1998) 2278–2324.
- [23] Erich Leo Lehmann, Joseph P. Romano, George Casella, *Testing Statistical Hypotheses*, vol. 3, Springer, 2005.
- [24] Yan Li, Hiroyuki Kambara, Yasuharu Koike, Masashi Sugiyama, Application of covariate shift adaptation techniques in brain–computer interfaces, *IEEE Trans. Biomed. Eng.* 57 (6) (2010) 1318–1324.
- [25] Zachary Lipton, Yu-Xiang Wang, Alexander Smola, Detecting and correcting for label shift with black box predictors, in: *International Conference on Machine Learning*, PMLR, 2018, pp. 3122–3130.
- [26] David Lopez-Paz, Maxime Oquab, Revisiting classifier two-sample tests, arXiv preprint, arXiv:1610.06545, 2016.
- [27] Rachel Luo, Rohan Sinha, Ali Hindy, Shengjia Zhao, Silvio Savarese, Edward Schmerling, Marco Pavone, Online distribution shift detection via recency prediction, arXiv preprint, arXiv:2211.09916, 2022.
- [28] Felipe Maia Polo, Renato Vicente, Effective sample size, dimensionality, and generalization in covariate shift adaptation, *Neural Comput. Appl.* (2022) 1–13.
- [29] Subha Maity, Diptavo Dutta, Jonathan Terhorst, Yuekai Sun, Moulinath Banerjee, A linear adjustment based approach to posterior drift in transfer learning, arXiv preprint, arXiv:2111.10841, 2021.
- [30] Jose G. Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V. Chawla, Francisco Herrera, A unifying view on dataset shift in classification, *Pattern Recognit.* 45 (1) (2012) 521–530.
- [31] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al., Scikit-learn: machine learning in python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [32] Aleksandr Podkopaev, Aaditya Ramdas, Tracking the risk of a deployed model and detecting harmful distribution shifts, arXiv preprint, arXiv:2110.06177, 2021.
- [33] Yury Polyanskiy, Yihong Wu, *Information Theory: From Coding to Learning*, 2022.
- [34] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, Andrey Gulin, CatBoost: unbiased boosting with categorical features, arXiv preprint, arXiv:1706.09516, 2017.
- [35] Stephan Rabanser, Stephan Günnemann, Zachary Lipton, Failing loudly: an empirical study of methods for detecting dataset shift, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [36] Treva K. Rice, Nicholas J. Schork, D.C. Rao, Methods for handling multiple testing, *Adv. Genet.* 60 (2008) 293–308.
- [37] Marco Saerens, Patrice Latinne, Christine Decaestecker, Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure, *Neural Comput.* 14 (1) (2002) 21–41.
- [38] Jessica Schrouff, Natalie Harris, Sanmi Koyejo, Ibrahim M. Alabdulmohsin, Eva Schnider, Krista Opsahl-Ong, Alexander Brown, Subhrajit Roy, Diana Mincu, Christina Chen, et al., Diagnosing failures of fairness transfer across distribution shift in real-world medical settings, *Adv. Neural Inf. Process. Syst.* 35 (2022) 19304–19318.
- [39] Karen Simonyan, Andrew Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv preprint, arXiv:1409.1556, 2014.
- [40] Nikolai V. Smirnov, Estimate of deviation between empirical distribution functions in two independent samples, *Bull. Moscow Univ.* 2 (2) (1939) 3–16.
- [41] Casper Kaae Sønderby, Jose Caballero, Lucas Theis, Wenzhe Shi, Ferenc Huszár, Amortised map inference for image super-resolution, arXiv preprint, arXiv: 1610.04490, 2016.
- [42] Skyler Speakman, Srihari Sridharan, Isaac Markus, Three population covariate shift for mobile phone-based credit scoring, in: *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*, 2018, pp. 1–7.
- [43] Masashi Sugiyama, Motoaki Kawanabe, *Machine Learning in Non-stationary Environments: Introduction to Covariate Shift Adaptation*, MIT Press, 2012.
- [44] Masashi Sugiyama, Matthias Krauledat, Klaus-Robert Müller, Covariate shift adaptation by importance weighted cross validation, *J. Mach. Learn. Res.* 8 (5) (2007).

- [45] Masashi Sugiyama, Taiji Suzuki, Takafumi Kanamori, *Density Ratio Estimation in Machine Learning*, Cambridge University Press, 2012.
- [46] Afonso Fernandes Vaz, Rafael Izbicki, Rafael Bassi Stern, Quantification under prior probability shift: the ratio estimator and its extensions, *J. Mach. Learn. Res.* 20 (1) (2019) 2921–2953.
- [47] Vladimir Vovk, Testing for concept shift online, *arXiv preprint, arXiv:2012.14246*, 2020.
- [48] Geoffrey I. Webb, Loong Kuan Lee, Bart Goethals, François Petitjean, Analyzing concept drift and shift from sample data, *Data Min. Knowl. Discov.* 32 (5) (2018) 1179–1199.
- [49] Maleakhi A. Wijaya, Dmitry Kazhdan, Botty Dimanov, Mateja Jamnik, Failing conceptually: concept-based explanations of dataset shift, *arXiv preprint, arXiv: 2104.08952*, 2021.
- [50] Yiming Xu, Diego Klabjan, Concept drift and covariate shift detection ensemble with lagged labels, in: 2021 IEEE International Conference on Big Data (Big Data), IEEE, 2021, pp. 1504–1513.