

Ensembling Shift Detectors: an Extensive Empirical Evaluation

Simona Maggio*¹[0000–0001–5653–937X]
Léo Dreyfus-Schmidt*²[0000–0001–8271–1217]

¹ Dataiku Lab, Paris, France simona.maggio@dataiku.com

² Dataiku Lab, Paris, France leo.dreyfus-schmidt@dataiku.com

* Equal contribution

Abstract. The term *dataset shift* refers to the situation where the data used to train a machine learning model is different from where the model operates. While several types of shifts naturally occur, existing shift detectors are usually designed to address only a specific type of shift. We propose a simple yet powerful technique to ensemble complementary shift detectors, while tuning the significance level of each detector’s statistical test to the dataset. This enables a more robust shift detection, capable of addressing all different types of shift, which is essential in real-life settings where the precise shift type is often unknown. This approach is validated by a large-scale statistically sound benchmark study over various synthetic shifts applied to real-world structured datasets.

Keywords: dataset shift · shift detector · ensemble · hypothesis testing

1 Introduction

It is crucial for ML practitioners to detect possibly harmful changes of the data consumed by ML models, thus identifying situations of *dataset shift*, where the joint distribution $\mathbb{P}(X, Y)$ of the input features X and output Y in the source domain S , used to train the model, is different from the distribution in the target domain T , where the model operates.

There are many reasons why target domain data is different in practice from carefully collected source data. For instance, by sample selection bias, when some samples are less likely to be included in the source datasets because of a specific selection process implicitly dependent on the output variable. This is often the case for data collected through web surveys, where self-selection and under-representation of people with limited internet access prevent reliable inference [5]. Another example of dataset shift affecting the model validity is when the target variable distribution is not stationary. This is the case for pneumonia diagnosis models, which are sensitive to seasonal outbreaks [14].

Dataset shifts are usually categorised as follows: *(i)* prior shift indicates that the only changing factor is the prior distribution $\mathbb{P}_S(Y) \neq \mathbb{P}_T(Y)$, *(ii)* covariate shift indicates that the only changing factor is the covariates distribution

$\mathbb{P}_S(X) \neq \mathbb{P}_T(X)$, while (iii) sample selection bias impacts both prior and covariate distributions. Another cause for dataset change is concept shift, where the dependence of the target variable on the features changes from source to target domain. In this work we focus on unsupervised approaches, while concept shift detection requires target labels and so is not covered in this study.

To protect deployed models from dataset shift, we need to compare source and target data. The most challenging aspect of shift detection stems from the absence of ground truth for the target variable, preventing the use of simple model metric monitoring techniques. Thus, in order to detect changes in the output distribution, it is essential to find robust proxies.

There are two dominant approaches for generic shift detection, one comparing feature distributions, either directly or through a domain-discriminating classifier trained on the features from the two domains [21], the second comparing the distributions of the model predictions as proxies for the true prior distributions, called Black-Box Shift Detector (BBSD) [14]. The feature-based approaches are more sensitive to shifts impacting the feature distribution, while the BBSD has been specifically designed to deal with prior shift. In both cases, hypothesis testing is employed to seek statistically significant differences. Although these two approaches have different designs and purposes, they are usually compared on generic shift detection tasks mainly on datasets for visual classification [21].

The key idea of this paper is that ensembling shift detectors and adapting the significance level of their statistical tests is a more robust approach, especially suited for real-life settings where the precise shift type is unknown. This result is supported by the following contributions:

- We propose various ensemble schemes to combine feature- and prediction-based shift detectors and validate detectors ensembling as a more robust solution.
- Motivated by the observation that the statistical tests needed for drift detection have a different behaviour across datasets, we propose the adaptation of the significance level required for detection to the specific dataset under study.
- We perform a statistically sound benchmark study of base and ensembled shift detectors on structured datasets (21 OpenML¹ and Kaggle² classification datasets, largely differing in number of features and class proportions). Indeed the vast majority of deployed models in production consume structured data, highlighting the need for studies on model and data monitoring approaches for this modality. The shift detectors are evaluated against various types of synthetic drifts, including selection bias drift, which is less explored in the existing shift detection benchmarks although very frequent in real-world scenarios.
- Last but not least, we also investigate both theoretically and empirically the degradation of the BBSD detector, a core component of the detectors ensembles, as the predictive power of the underlying classifier drops.

¹ www.openml.org

² www.kaggle.com

After presenting related work, Section 3 describes the shift detection approaches in more details, as well as promoting the benefit of adapting the statistical tests to the dataset. Section 4 presents the experimental setting with a focus on synthetic drift generation, while Section 5 discusses the main outcomes and observations from the experiments.

2 Related Work

Dataset shift [20,16] has been addressed extensively in the context of **domain adaptation**, by exploring techniques aiming to align the model representations of source and target data, especially without target labels [13,22].

Label shift correction is a sub-field of domain adaptation that recently contributed to novel shift detection approaches. In particular, [14] proposed a shift detection approach based on the Black Box Shift Learning technique to correct label shift. [21] presented a general framework for shift detection and evaluated the BBSD and other state-of-the-art shift detection techniques on image datasets. A similar approach to correct prior shift is the Regularized Learning under Label Shifts [3], also relying on classifiers predictions. An alternative maximum likelihood approach to estimate target label distributions highlights the benefits of calibration under prior shift [1,11].

Best practices for responsible deployment of ML models include **data monitoring**, by inspecting incoming data for any signs of deviation from the expected scenario [2,12]. While dataset shift detection is an important part of ML monitoring system, recent work [24,10] has focused on directly predicting models performance drop for specific task-relevant drifts.

Besides monitoring possible distributional changes, the analysis of individual samples against the source dataset introduces the **anomaly detection** problem, extensively treated in the machine learning literature [23,7,15].

Another related field is **out-of-distribution detection** [25], which together with shift detection task has been used recently to evaluate ML models **predictive uncertainty** [18]. The study of predictive uncertainty offers a unifying view of the robustness a ML model should be equipped with to address any type of unexpected abnormality in input data. Finally **hypothesis testing** [26] is a core part of all shift detection techniques.

3 Shift Detectors

3.1 Notations and Problem Setup

Let \mathcal{X} be an input space and \mathcal{Y} be a label space where $|\mathcal{Y}|$ is finite. Let $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ be random variables. We denote by $f : X \rightarrow Y$ a predictor of the classification task at hand and we will refer to it as the *primary model*.

3.2 Feature-Based Detection

The simplest feature-based detection is univariate hypothesis testing on individual features. For n features, this requires n univariate hypotheses tests to separately compare the distributions of features between the source and the target domain, and eventually aggregate the n p -values through Bonferroni correction.

However it can quickly become cumbersome when the feature space is large. For a more scalable solution, a more compact data representation can be obtained through standard dimensionality reduction techniques such as Principal Component Analysis (PCA) or Sparse Random Projection (SRP) [21].

In this work we do not consider detectors based on multivariate hypothesis testing as they were found to offer comparable performance to aggregated univariate tests [21].

A different solution to detect feature distributional changes in high dimensional datasets relies on domain-discriminating models [4]. The *domain classifier* is a model trained to discriminate source and target domains, using the very same features employed by the primary model. By design a domain classifier is meant to identify covariate shift and sample selection bias, but not prior shift which assumes unchanged feature distribution.

High accuracy of the classifier is a symptom of dataset shift, therefore a Binomial test with a null hypothesis of 0.5 accuracy for indistinguishable domains and balanced datasets is used to make sure the observed difference is statistically significant. One limitation of the domain classifier is the need of retraining for any new incoming target dataset.

3.3 Prediction-Based Detection

BBSD exploits the primary model to measure the distributions of predictions on source and target features, which are used as proxies for the true source and target label distributions. There are two variants of BBSD, one using the probability outputs (BBSDs), and the other using the hard-thresholded predictions (BBSDh).

For BBSDh, a χ^2 test is used to compare source and target predicted class distributions, while for BBSDs the Kolmogorov-Smirnov (KS) test is employed to compare per-class probability distributions. The primary model predictions are k -dimensional with $k = |\mathcal{Y}|$ the number of classes in the primary task, thus the BBSDs shift detection requires k univariate hypothesis tests. As in the previous section, Bonferroni correction [21] can be used to aggregate the k KS tests.

Complementarily to the domain classifier, the BBSD has been specifically designed to address prior shift situations, but the distribution of predictions can also serve as proxy for the distribution of covariates [14].

3.4 Limitations of Shift Detectors

BBSD degrades as primary performances drop. We present two limitations of the above shift detection method. Firstly, the statistical power of BBSD is limited by the predictive power of the primary model. We theoretically prove in the

Appendix (cf. Section A.1) that there exist prior shifts such that the smaller the predictive power of the model is, the higher the p -values of the BBSD test are, reducing the sensitivity of the test. This phenomenon also occurs on more general type of shifts as validated experimentally (cf. Section 4).

Detection power reduced by Bonferroni correction. Another limitation stems from the use of Bonferroni correction. Indeed, when the p -values of k hypotheses are aggregated by Bonferroni correction, a significance level of α/k is required for each hypothesis to ensure that the Type-I error rate stays below α . However if the hypotheses are not independent, this correction is too conservative and the family-wise Type-I error is controlled at a level strictly lower than α , resulting in a loss of power [17,19]. This is the case for the BBSDs test as the underlying tests are done on the distribution of predicted probabilities of each class. For binary classification, the null hypothesis for one class and its complementary class are the same so that the Type-I error rate is actually bounded by $\alpha/2$ (instead of the expected α), reducing the test’s power. This is also the case for the univariate hypothesis testing on individual features, where the Bonferroni correction is conservative because of features correlation. We propose to address the latter limitation by adapting the significance level to the dataset.

3.5 Dataset Adaptive Significance Level

To improve the BBSDs test’s power, keeping the Type-I error below a desired rate (i.e. 5%), we set the significance level to the 5% quantile of the empirical p -value distribution under the null hypothesis. We achieve this by performing 100 runs of shift detection comparing different random splits of the source dataset only.

A dataset-specific significance level is not only beneficial for shift detectors requiring aggregation of multiple univariate tests, but also when dealing with small datasets (i.e. from 10 to 100 samples). Indeed, performing statistical tests on small datasets yields p -values that can only assume few values. This quantization leads to non-uniform distributions of p -values under the null hypothesis so that a dataset-specific significance level is also beneficial.

The drawback of this adaptation represents the initial setup cost of running detection experiments on the source data, but the selected significance level can be used for all subsequent shift detection tasks for the given dataset.

Finally, for very imbalanced datasets (*creditcard*, *pc2* and *mc1* where the minority class represents less than 0.7%), the distribution of the BBSDs p -values under H_0 is extremely skewed towards 1. The study of the impact of imbalanced datasets on prediction-based detectors is left for future research.

As a natural extension of base shift detectors, the detectors ensembles also benefit from the adaptation of the significance level.

3.6 Detectors Ensembles

We evaluate two ensembling strategies combining feature and prediction-based shift detectors, in order to exploit their complementary detection abilities:

- domain classifier trained on both features and primary model predictions;
- pair of statistical tests from complementary shift detectors with Bonferroni correction.

Ensembling strategies are naturally robust to various shift types, and represent a better fit to real-life settings where the precise shift type is often unknown. We confirm experimentally that ensembling retains the advantages of feature- and prediction-based shift detectors (cf. Section 4), providing a single practical tool for reliable drift monitoring.

4 Experimental Setup

The purpose of our experiments is to compare the different detection approaches on various types of simulated drift over a large collection of structured datasets. Specifically we aim at highlighting differences in the approaches related to both the detection power and efficiency, the latter expressed in terms of minimum number of samples required to detect a shift. All the details of the experiments setup and reproducibility of this work are available in the Appendix (cf. Section A.2).

4.1 Datasets and Shift Simulation

In our experiments we use 21 datasets for classification tasks from OpenML and Kaggle as listed in Table 7 of the Appendix.

In order to simulate dataset drift we synthetically apply different types of shift to a split of the original dataset. The 10 simulated shifts belong to the following categories and are detailed in Table 1:

- Prior shift: generated by changing the fraction of samples belonging to a class.
- Covariate shift due to Gaussian noise: generated by adding Gaussian noise to some numeric features of a fraction of samples.
- Covariate shift due to Adversarial noise: generated by applying Adversarial noise to numeric features. It’s a more subtle kind of noise, slightly changing the features but inducing the primary model to switch its predicted class.
- Selection bias: generated by selecting samples with a probability dependent on the sample features and implicitly dependent on the class, possibly over-sampling by interpolation of existing observations.

The severity of those shifts is controlled by the parameters shown in Table 1. Our experiments explore the effect of changing those parameters values generating a total of 19 drifts, then averaging the outcome from the same shift type to build the results for the 10 shifts types in Table 1.

The combination of different shift types is also relevant as it can occur in real-life settings, for this reason the study of the detectors response to composite shifts will be considered in future work.

Table 1: Simulated shift types.

Category	Shift type	Parameters	Description
Prior shift	Knock-out	$s = 25\%$ $s = 40\%$	Remove $s\%$ of majority class.
	Only-One		Select only-one minority class.
Covariate shift	Small Gaussian	$s = 50\%$ $f = 50\%$	Small amount of Gaussian Noise applied on $s\%$ of samples
		$s = 50\%$ $f = 100\%$	and $f\%$ of features.
		$s = 100\%$ $f = 50\%$	
		$s = 100\%$ $f = 100\%$	
	Medium Gaussian	$s = 50\%$ $f = 50\%$	Medium amount of Gaussian Noise applied on $s\%$ of samples
		$s = 50\%$ $f = 100\%$	and $f\%$ of features.
Selection bias	Joint Subsampling		Keeps an observation with probability decreasing as points are away from the samples mean.
		$f = 100\%$	Subsample with low probability samples with low feature values separately for $f\%$ features.
	Under-sampling	$s = 50\%$	Keep $s\%$ of samples, selecting samples close to the minority class (NearMiss3 heuristics).
	Over-sampling	$s = 50\%$	Replace $s\%$ of samples with samples interpolated from the remaining part.

4.2 Experiments

For each run of drift experiment, we test all the mentioned shift detection approaches on subsets of the source and target datasets with number of samples in [10, 100, 500, 1000, 2000]. The different detection approaches are reported using the following labels:

- *BBSDs*: soft version of BBSD with Random Forest primary model and KS test.
- *BBSDh*: hard version of BBSD with Random Forest primary model and χ^2 test.
- *Test_X*: KS test on input features.
- *Test_PCA*: KS test on PCA-projected features.
- *Test_SRP*: KS test on SRP-projected features.
- *DC*: Random Forest domain classifier with Binomial test.

The ensembling strategies are reported using the following labels:

- *BBSDs + X*: *BBSDs* and *Test_X* with Bonferroni correction.
- *BBSDs + DC*: *BBSDs* and *DC* with Bonferroni correction.
- *DC**: *DC* trained on both features and primary predictions.

For any detector, its adaptive variant using a significance level tailored to the dataset (cf. Subsection 3.5) is referred to with the suffix (*adapt*), i.e. *BBSDs (adapt)*.

Overall, we have collected p -values from shift detection experiments for 19 different drifts, run 5 times, for 5 different sizes, for each of the 21 datasets. Furthermore the results over the different runs are averaged, as well as the results from the same types of drift with different drift parameters to yield 210 (21 datasets \times 10 shift types) detection results per dataset size per detector.

4.3 Metrics

The p -values of the detectors are averaged over multiple runs and when their value is less than the significance level, we consider that a drift has been detected. The significance level used in the non-adaptive version of the statistical test is 0.05.

Considering shift detectors as binary classifiers where one sample is one comparison of source and target datasets, we can measure its accuracy and true positive rate (TPR). When the accuracy is reported the evaluated shift scenarios include a negative (no shift) situation per dataset.

To measure the data efficiency of a shift detector, we define an **efficiency score** as the minimum size level required to detect the shift. As the detectors are evaluated at different datasets sizes [10, 100, 500, 1000, 2000], they are assigned a score from 5 to 1. A detector failing to find the shift at any size is assigned a score of 0.

4.4 Statistical Comparison

In order to fairly compare the detection approaches, we use a Friedman test to statistically assess whether the detectors have different performances. When the conclusion is positive, a Nemenyi post-hoc test is used to determine pairwise equivalence of the methods [9].

Based on the efficiency score defined in Section 4.3, the average ranks for each detector are computed across all datasets. The Friedman test then checks whether the detectors average ranks are significantly different from the mean rank expected under the null-hypothesis. When the former test rejects the null-hypothesis, the Nemenyi test is used to seek for pairwise differences between ranks.

In order to have independent observations from different datasets, we perform these statistical tests separately for each of the 10 shift types. We use the library *autorank*³ to perform this statistical comparison.

5 Results

We first give the experiment results comparing ensembled and base shift detectors over all datasets and by shift types. We then highlight the complementarity of feature- and prediction-based shift detectors motivating the ensembling strategy. Finally, we present ablation studies comparing base detectors and ensembles to their adaptive counterparts.

³ <https://github.com/sherbold/autorank>

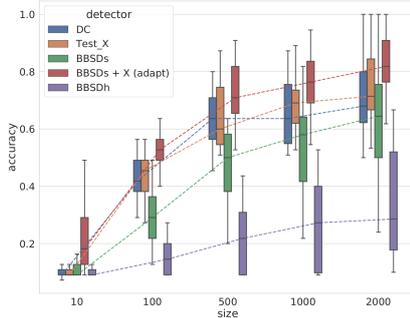


Fig. 1: Shift detectors accuracy by dataset size.

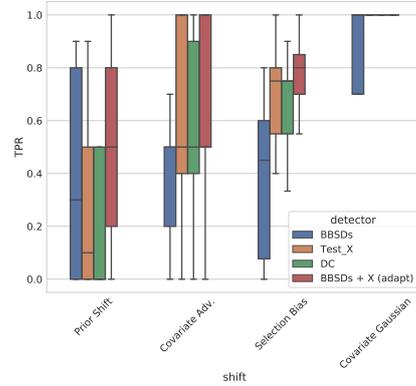


Fig. 2: Shift detectors true positive rate (TPR) by type of shift.

5.1 Ensembling Shift Detectors

When observing the average detection accuracy across datasets and shift types for different dataset sizes (Figure 1) some simple detection approaches stand out (*Test_X* and *DC*), although slightly less accurate than the adaptive ensemble technique (*BBSDs+X (adapt)*). However the true positive rate for different shift types for a fixed size of 1000 (Figure 2) reveals a more complex landscape, where the performance of base shift detectors is not uniform. For an easier reading of Figures 1 and 2, some detectors have been omitted and the various shifts have been grouped by type (cf. Table 1), but all the detailed TPR and accuracy results for individual shift detectors and shift types can be found in Tables 8 and 9 of the Appendix.

In Tables 2 and 3 we report the mean accuracy by dataset size and mean TPR by type of shift for the base detectors and the detectors ensembles with adaptation of the significance level to the dataset. The ensemble *BBSDs+X (adapt)* comes out as the most accurate shift detector overall.

In addition to be sensitive at lower sizes *BBSDs+X (adapt)* is also able to capture alone most types of shifts. This approach represents a more robust solution able to detect the shift situations missed by either of the base shift detectors. It is interesting to notice that simpler ensembles *BBSDs+X (adapt)* and *BBSDs+DC (adapt)* perform very similarly to *DC**, while not requiring an additional training. Ensembling feature- and prediction-based shift detectors is a promising strategy for an effective and robust drift monitoring.

5.2 Comparison of Base Shift Detectors

In order to have a global comparison of the performance of base detectors, we aggregate the results from different dataset sizes with the efficiency score and

Table 2: Accuracy across all datasets by size (mean \pm std). Best values and all values within 95% confidence interval are in bold.

Size	BBSDs	Test_X	DC	Ensembles (adapt)		
				DC*	BBSDs+DC	BBSDs+X
10	0.11 \pm 0.03	0.10 \pm 0.02	0.10 \pm 0.02	0.20 \pm 0.24	0.14 \pm 0.03	0.22 \pm 0.11
100	0.30 \pm 0.10	0.45 \pm 0.10	0.44 \pm 0.09	0.46 \pm 0.17	0.50 \pm 0.10	0.54 \pm 0.10
500	0.46 \pm 0.14	0.65 \pm 0.12	0.64 \pm 0.10	0.63 \pm 0.14	0.67 \pm 0.12	0.73 \pm 0.11
1000	0.54 \pm 0.16	0.69 \pm 0.11	0.66 \pm 0.11	0.63 \pm 0.13	0.72 \pm 0.12	0.75 \pm 0.12
2000	0.63 \pm 0.19	0.75 \pm 0.13	0.71 \pm 0.14	0.67 \pm 0.15	0.79 \pm 0.12	0.83 \pm 0.12

Table 3: TPR across all datasets by shift type at dataset size of 1000 (mean \pm std). Best values and all values within 95% confidence interval are in bold.

Shift type	BBSDs	Test_X	DC	Ensembles (adapt)		
				DC*	BBSDs+DC	BBSDs+X
Knock-Out	0.23 \pm 0.30	0.11 \pm 0.24	0.00 \pm 0.00	0.01 \pm 0.05	0.25 \pm 0.30	0.34 \pm 0.40
Only-One	0.60 \pm 0.46	0.55 \pm 0.49	0.46 \pm 0.47	0.51 \pm 0.48	0.70 \pm 0.44	0.73 \pm 0.41
Small Gaussian	0.74 \pm 0.37	0.96 \pm 0.14	1.00 \pm 0.00	0.95 \pm 0.12	1.00 \pm 0.00	0.99 \pm 0.04
Medium Gaussian	0.87 \pm 0.32	0.98 \pm 0.06	1.00 \pm 0.00	1.00 \pm 0.00	0.99 \pm 0.04	0.99 \pm 0.04
Adv. ZOO	0.21 \pm 0.40	0.43 \pm 0.51	0.42 \pm 0.50	0.47 \pm 0.50	0.43 \pm 0.51	0.45 \pm 0.49
Adv. Boundary	0.72 \pm 0.42	0.77 \pm 0.35	0.70 \pm 0.41	0.72 \pm 0.40	0.74 \pm 0.40	0.98 \pm 0.06
Subsampling Joint	0.04 \pm 0.08	0.20 \pm 0.30	0.10 \pm 0.19	0.30 \pm 0.35	0.16 \pm 0.26	0.34 \pm 0.30
Subsampling	0.42 \pm 0.41	0.71 \pm 0.45	0.70 \pm 0.44	0.17 \pm 0.29	0.72 \pm 0.43	0.72 \pm 0.45
Under-sampling	0.44 \pm 0.44	1.00 \pm 0.00				
Over-sampling	0.57 \pm 0.40	1.00 \pm 0.00				

show the global average ranks by type of shift in Table 4. A first observation is that overall the feature-based detectors have lower rank, indicating that they are more effective than other approaches, requiring less data to detect a drift situation.

Indeed for all the types of shift the Friedman test at $p = 0.05$ reported that the approaches are not statistically equivalent.

We thus look at the results of the post-hoc Nemenyi test, represented in Figure 3 for the case of under-sampling shift. The Critical Distance (CD) between the average ranks reveals two groups of different detectors, with feature-based detectors outperforming the BBSd approaches. The *Test_PCA*, *Test_X*, *DC* detectors have statistically higher efficiency than *BBSdh* and *BBSDs*. The Nemenyi test at $p = 0.05$ is not powerful enough to draw conclusions about the other detectors.

Overall the Nemenyi tests (Figure 4) highlight the following differences:

- *Prior shift*: prediction-based detectors perform better than other approaches when the label distribution is affected, as we could expect by design.
- *Gaussian noise*: domain classifier and direct input features testings are more effective than other approaches.
- *Adversarial Boundary*: BBSDs and direct feature testing perform slightly better than domain classifier and significantly better than *BBSdh*. This perturbation is too subtle to be spotted by the domain classifier, although

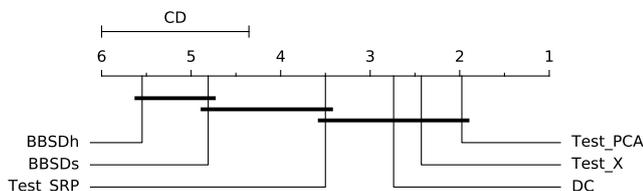


Fig. 3: Pairwise comparison of all detectors with the Nemenyi test for *Under-sampling* shift.

Table 4: Average ranks of 6 shift detectors based on efficiency score.

Shift type	BBSDs	BBSDh	Test_X	Test_PCA	Test_SRP	DC
Knock-Out	2.71	3.02	3.74	3.67	3.62	4.24
Only-One	2.64	3.07	3.45	3.90	3.90	4.02
Small Gaussian	2.93	5.43	1.95	4.93	4.26	1.50
Medium Gaussian	2.36	5.52	2.26	4.79	4.14	1.93
Adversarial ZOO	3.81	4.21	2.60	3.93	3.40	3.05
Adversarial Boundary	2.95	4.55	2.31	4.64	3.55	3.00
Joint Subsampling	3.98	3.98	3.26	2.52	3.29	3.98
Subsampling	4.12	4.86	2.69	3.10	3.00	3.24
Under-sampling	4.81	5.55	2.43	1.98	3.50	2.74
Over-sampling	4.90	5.69	2.26	2.14	3.52	2.48
All	3.54	4.63	2.63	3.58	3.65	2.96

its important impact on class distributions is easily detected by both BBSDs and Test_X.

- *Subsampling, Under-sampling, Over-sampling*: features-based detectors perform significantly better than prediction-based detectors on perturbations simulating selection bias.

The previous findings highlight the complementarity of feature- and prediction-based drift detectors on different drift scenarios, motivating the proposed ensembling strategy (cf. Section 3.6).

The individual results from each drift scenario per dataset are available in the paper repository ⁴ for full inspection.

5.3 Impact of Dataset-Adaptive Significance Level

The dataset-adaptive significance level in BBSDs is consistently higher than the Bonferroni-corrected significance level of α/k . For instance, for all binary

⁴ https://github.com/dataiku-research/drift_detectors_benchmark/disaggregated_results.md

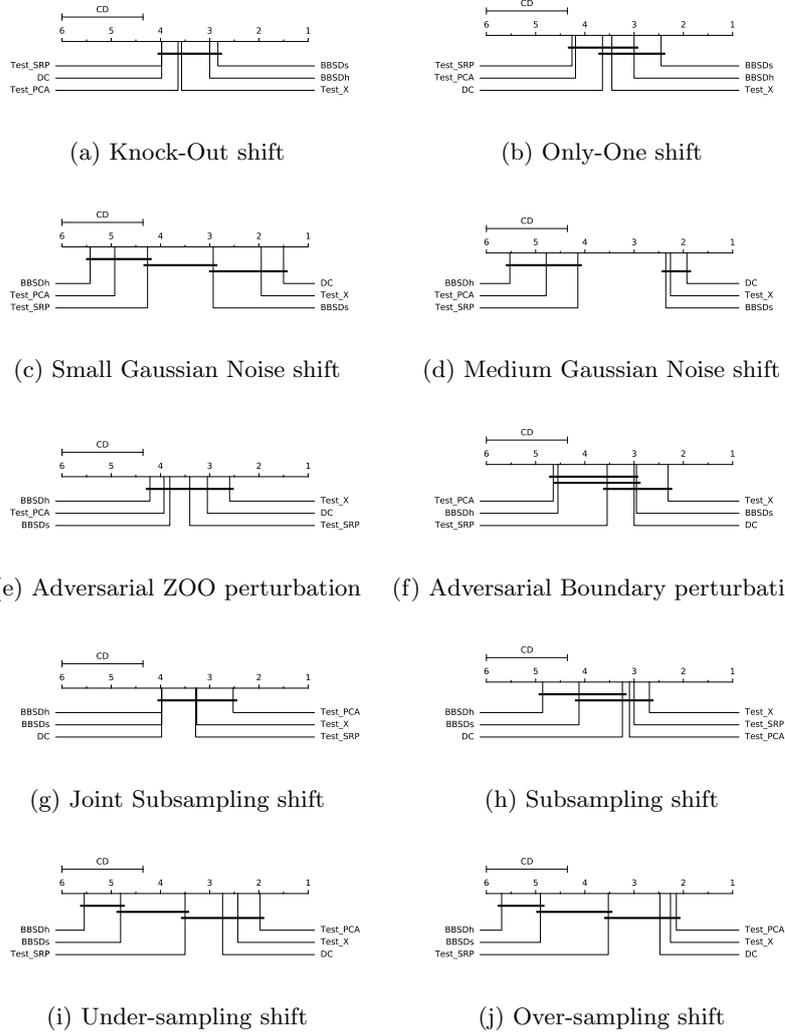


Fig. 4: Nemenyi post-hoc tests for the different types of drift.

classification datasets in Table 7 of the Appendix, the standard Bonferroni-corrected significance level is $\alpha/2 = 0.025$, while on average the computed adaptive level is 0.089: this gap shows the margin to improve the test's power, keeping the Type-I error below 5%. This is illustrated in Figure 5 on the *MagicTelescope* dataset, showing the BBSD detector as a binary classifier on all (positive) shift detection tests at size 2000 described in the experiments and additional 100 (negative) shift detection tests on randomly sampled validation and test sets of the same size, on which no shift is applied.

In order to analyze the improvement achieved by adapting the significance level of the statistical tests as described in Section 3.5, we compare the regular and adaptive versions of the top-3 base shift detectors through Friedman and Nemenyi post-hoc test. Table 5 and Figure 7 (in the Appendix) highlight the improved sensitivity of the adaptive detectors. The enhancement is more important for the prediction-based technique because of the conservative Bonferroni correction. The sensitivity is also improved for *Test_X*, where the correction is conservative because of features correlation. This adaptation of the significance level reduces the gap between the two approaches, bringing closer prediction-based and feature-based shift detectors.

As an ablation study we investigate the improvements of the detectors ensembles due to the adaptation of the significance level and confirm the benefit of this technique as shown in Table 6.

Table 5: Average ranks of 3 shift detectors based on efficiency score, comparing fixed and adaptive significance levels. The best values between regular and adaptive versions of the same shift detector are in bold, while the best values overall are underlined.

Shift type	BBSDs	BBSDs (adapt)	DC	DC (adapt)	Test_X	Test_X (adapt)
Knock-out	3.17	<u>2.12</u>	4.43	4.43	4.02	2.83
Only-One	3.10	<u>2.29</u>	4.50	4.36	3.95	2.81
Small Gaussian	5.29	3.67	3.10	2.64	3.98	<u>2.33</u>
Medium Gaussian	4.38	2.93	3.76	3.40	4.31	<u>2.21</u>
Adv. ZOO	4.40	3.64	3.76	3.64	3.10	<u>2.45</u>
Adv. Boundary	4.02	2.83	4.38	4.24	3.64	<u>1.88</u>
Joint Subsampling	4.02	3.50	4.02	3.86	3.29	<u>2.31</u>
Subsampling	4.60	3.98	3.81	3.57	3.19	<u>1.86</u>
Under-sampling	5.38	4.50	3.17	3.17	2.81	<u>1.98</u>
Over-sampling	5.67	4.60	3.02	2.69	2.81	<u>2.21</u>

5.4 Impact of Model Quality on BBSDs

To confirm the limitations of BBSDs on the primary model performance (cf. Section 3), we apply a random perturbation to the primary model with probability p , yielding a primary model with quality $1 - p$ in $[0.5, 1.0]$ and evaluate the power of the BBSDs test. As illustrated in Figure 6, showing the average TPR across 21 datasets with size of 1000 samples, the power decreases with the primary model quality, along with the variance of the results across the datasets. More difficult shifts (such as low intensity or a small amount of drifted samples) require a better primary model for the BBSDs test power to hold (shift-specific details in Table 10 of the Appendix). Regardless of the model quality, BBSDs is also less adequate in detecting some types of selection bias, as already observed in Figure 2.

Table 6: Average ranks of 3 ensemble shift detectors based on efficiency score, comparing fixed and adaptive significance levels. The best values between regular and adaptive versions of the same shift detector are in bold, while the best values overall are underlined.

Shift type	BBSDs+X	BBSDs+X (adapt)	BBSDs+DC	BBSDs+DC (adapt)	DC*	DC* (adapt)
Knock-out	3.55	<u>1.95</u>	3.43	2.93	4.57	4.57
Only-One	3.55	<u>2.62</u>	3.40	2.83	4.36	4.24
Small Gaussian	4.29	<u>2.26</u>	3.40	2.62	4.36	4.07
Medium Gaussian	4.19	<u>2.21</u>	3.98	2.24	4.50	3.88
Adv. ZOO	3.24	<u>2.57</u>	3.81	3.52	3.85	3.85
Adv. Boundary	3.21	<u>1.88</u>	3.86	3.17	4.62	4.15
Joint Subsampling	3.24	<u>2.43</u>	3.79	3.55	4.00	4.00
Subsampling	2.67	<u>2.19</u>	3.31	2.79	5.02	5.02
Under-sampling	3.43	<u>2.57</u>	4.00	3.57	3.71	3.71
Over-sampling	3.60	<u>2.88</u>	3.88	3.17	3.65	3.65

6 Conclusion

In this paper we propose a shift detectors ensembling technique capable of addressing all different types of studied drift scenarios. The key components of the proposed approach are the combination of complementary base detectors, designed to address different types of shift, and the adaptation of the significance level of the detectors statistical tests to the specific dataset under study.

The improved robustness of our approach is validated by a large-scale benchmark study comparing it to state-of-the-art shift detectors on 21 structured real-world datasets. For this purpose we simulate drifts with 10 different types of perturbations, including shift types not studied in previous works for tabular datasets, such as adversarial noise and sample selection bias. Our benchmark study highlights the complementarity of base drift detectors on different drift scenarios, motivating our ensembling approach. This benchmark also includes ablation studies showing that the dataset-adaptive significance level provides both base and ensembled shift detectors with higher detection power, while preserving the desired false positive rate.

Throughout our experiments, we observe that adaptive shift detectors ensembling represents the strongest strategy, robust to the various shift types, making this approach a natural choice for monitoring models in production in real-life settings, where the possible drift scenario is unknown.

A Appendix

A.1 Proof from Subsection 3.5

Proof (There exist prior shifts such that the smaller the predictive power of the model is, the higher the p -values of the BBSD test are). The true and predicted

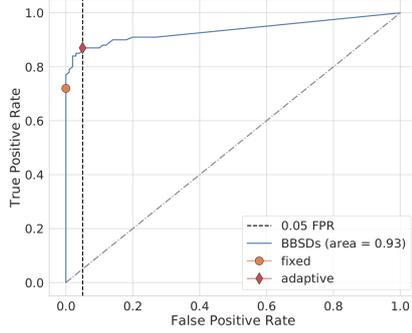


Fig. 5: ROC curve of BBSDs as a binary classifier on all positive and negative drift experiments for the *MagicTelescope* dataset.

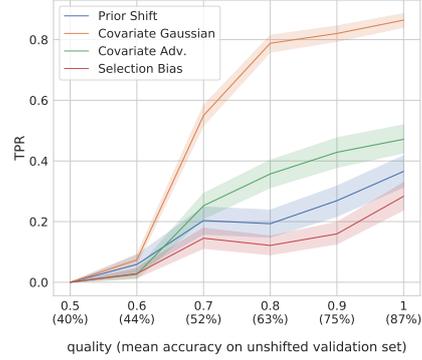


Fig. 6: Average TPR with standard deviation from the BBSDs test with perturbed primary model. Higher corruption probability degrades the model performances.

class distributions in the source (resp. target) domain are denoted by \mathbf{p}_S and $\hat{\mathbf{p}}_S$ (resp. \mathbf{p}_T and $\hat{\mathbf{p}}_T$).

Let us take a model $f : X \rightarrow Y$ solving a classification task with an ill-conditioned but still invertible confusion matrix \mathbf{C} , with eigenvalues λ_i and eigenvectors \mathbf{v}_i . Its minimum eigenvalue in absolute value is λ_{min} and the corresponding eigenvector is \mathbf{v}_{min} .

First, we make clear the connection between the model’s predictive performance and the norm of the minimum eigenvalue. We perturb the primary model f prediction’s by predicting random labels with probability p . The new model is denoted by \tilde{f} . As the perturbation p tends to 1, the eigenvalue $\lambda_{min}(\tilde{f})$ tends to 0. We then consider the family of prior shift $\mathbf{p}_T = \mathbf{p}_S + \alpha \cdot \mathbf{v}_{min}(\tilde{f})$ indexed by α such that \mathbf{p}_T is a probability measure on Y .

By Lemma 1 of [14] under the prior shift assumption and by the law of total probability we have that the following equation holds for both the source and the target domain:

$$\mathbb{P}(\hat{y}) = \sum_{y \in \mathcal{Y}} \mathbb{P}(\hat{y}|y)\mathbb{P}(y)$$

Considering the column-normalized confusion matrix \mathbf{C} as an estimator for $\mathbb{P}(\hat{y}|y)$, we can write in matrix form:

$$\hat{\mathbf{p}}_S = \mathbf{C} \cdot \mathbf{p}_S \quad \hat{\mathbf{p}}_T = \mathbf{C} \cdot \mathbf{p}_T$$

Let us assume $\mathbf{p}_S \neq \mathbf{p}_T$ with $\mathbf{p}_T = \mathbf{p}_S + \alpha \cdot \mathbf{v}_{min}$ with $\alpha \in [\alpha_{min}, \alpha_{max}]$, such that $p_{T_i} \in [0, 1] \forall i$. Note that α_{min} is always non positive and α_{max} is always non negative.

For the target distribution we have that:

$$\begin{aligned}\hat{\mathbf{p}}_T &= \mathbf{C} \cdot \mathbf{p}_T = \mathbf{C} \cdot \mathbf{p}_S + \mathbf{C} \cdot \alpha \cdot \mathbf{v}_{min} \\ &= \hat{\mathbf{p}}_S + \alpha \cdot \lambda_{min} \cdot \mathbf{v}_{min} \approx \hat{\mathbf{p}}_S\end{aligned}$$

Thus $\hat{\mathbf{p}}_S \approx \hat{\mathbf{p}}_T$ while $\mathbf{p}_S \neq \mathbf{p}_T$.

In particular:

$$\|\hat{\mathbf{p}}_S - \hat{\mathbf{p}}_T\|_2 = |\alpha \cdot \lambda_{min}| \approx 0$$

$$\|\mathbf{p}_S - \mathbf{p}_T\|_2 = |\alpha|$$

with $|\alpha| \in [0, \max\{|\alpha_{min}|, \alpha_{max}\}]$.

This shows that as the perturbation p tends to 1, the ℓ_2 norm $\|\hat{\mathbf{p}}_S - \hat{\mathbf{p}}_T\|_2$ tends to 0 while $\|\mathbf{p}_S - \mathbf{p}_T\|_2 = \alpha$. As the χ^2 statistic is controlled by ℓ_2 -norm, this means BBSD power is limited by the strength of the perturbation.

A.2 Experiments Details

Datasets Details on the 21 OpenML⁵ and Kaggle⁶ structured classification datasets are available in Table 7.

Preprocessing and Algorithms Hyperparameters For each dataset (see Table 7), a fixed training set of 1000 samples is randomly selected. The remainder of each dataset is then randomly split into a validation set and test set of 2000 samples each. The former is used as source dataset, while the latter is perturbed by applying a particular type of drift and constitutes the target dataset. Apart from simple imputation of missing values and one-hot encoding of categorical features, no preprocessing is performed. The training set is used to fit PCA and SRP parameters, as well as to train the primary model, a Random Forest, with default *sklearn* hyper-parameters. Random Forest is an industry standard model, extensively employed in enterprise machine learning for its performance and interpretability. The number of dimensions of PCA- and SRP-processed features is set to the same value corresponding to a variance retention rate in PCA of 0.8. The validation-test split is repeated over 5 random seeds so that each drift detection experiment is repeated 5 times.

Reproducibility We leveraged various open source libraries to generate those shifts in our experiments: the prior and Gaussian shifts rely on the library provided in the Failing Loudly github repository⁷, the adversarial shifts are generated using the Adversarial Robustness Toolbox⁸, while the under-sampling

⁵ www.openml.org

⁶ www.kaggle.com

⁷ <https://github.com/steveab/failing-loudly>

⁸ <https://github.com/IBM/adversarial-robustness-toolbox>

Table 7: Classification datasets from OpenML and Kaggle (denoted by a *).

Dataset	classes	features	size	% minority class
adult	2	14	48842	23.9
Amazon employee access	2	9	32769	5.8
Artificial characters	10	7	10218	5.9
Click prediction small	2	9	39948	16.8
codrna	2	8	488565	33.3
creditcard	2	30	284807	0.17
electricity	2	8	45312	42.5
higgs	2	29	98050	47.1
homesite-quote-conversion*	2	298	260753	18.8
JapaneseVowels	9	14	9961	7.9
jm1	2	21	10885	19.3
letter	26	16	20000	3.7
MagicTelescope	2	10	19020	35.2
mc1	2	38	9466	0.72
mozilla4	2	5	15545	32.9
optdigits	10	64	5620	9.9
otto-group-product-classification*	9	94	61878	3.1
pc2	2	36	5589	0.41
phoneme	2	5	5404	29.3
santander-customer-satisfaction*	2	370	76020	4.0
waveform_5000	3	40	5000	33.1

and over-sampling shifts rely on techniques from imbalanced-learn⁹. For the sake of completeness and reproducibility, all the code used to generate the shifts and to perform drift experiments is accessible in a public repository¹⁰.

A.3 Detailed Results from Subsections 5.2, 5.3 and 5.4

Values are reported with *mean* \pm *std*. Best values and all values within 95% confidence interval are in bold.

Table 8: Accuracy across all datasets by size.

Size	BBSDh	BBSDs	DC	Test_PCA	Test_SRP	Test_X
10	0.10 \pm 0.02	0.11 \pm 0.03	0.10 \pm 0.02	0.13 \pm 0.03	0.11 \pm 0.03	0.10 \pm 0.02
100	0.15 \pm 0.06	0.30 \pm 0.10	0.45 \pm 0.09	0.36 \pm 0.06	0.29 \pm 0.10	0.45 \pm 0.10
500	0.22 \pm 0.11	0.46 \pm 0.14	0.64 \pm 0.10	0.47 \pm 0.11	0.47 \pm 0.14	0.65 \pm 0.12
1000	0.27 \pm 0.15	0.54 \pm 0.16	0.66 \pm 0.11	0.51 \pm 0.13	0.54 \pm 0.16	0.69 \pm 0.11
2000	0.33 \pm 0.19	0.63 \pm 0.19	0.71 \pm 0.14	0.55 \pm 0.20	0.60 \pm 0.19	0.75 \pm 0.13

⁹ <https://github.com/scikit-learn-contrib/imbalanced-learn>

¹⁰ https://github.com/dataiku-research/drift_detectors_benchmark

Table 9: TPR across all datasets by shift type at dataset size of 1000.

Shift type	BBSDh	BBSDs	DC	Test_PCA	Test_SRP	Test_X
Knock-Out	0.15 \pm 0.27	0.23 \pm 0.30	0.00 \pm 0.00	0.06 \pm 0.18	0.09 \pm 0.21	0.11 \pm 0.24
Only-One	0.50 \pm 0.52	0.60 \pm 0.46	0.46 \pm 0.47	0.36 \pm 0.43	0.47 \pm 0.48	0.55 \pm 0.49
Small Gaussian	0.04 \pm 0.14	0.74 \pm 0.37	1.00 \pm 0.00	0.24 \pm 0.33	0.41 \pm 0.45	0.96 \pm 0.14
Medium Gaussian	0.28 \pm 0.33	0.87 \pm 0.32	1.00 \pm 0.00	0.49 \pm 0.45	0.66 \pm 0.45	0.98 \pm 0.06
Adv. ZOO	0.14 \pm 0.36	0.21 \pm 0.40	0.42 \pm 0.50	0.22 \pm 0.40	0.32 \pm 0.47	0.43 \pm 0.51
Adv. Boundary	0.47 \pm 0.45	0.72 \pm 0.42	0.70 \pm 0.41	0.33 \pm 0.38	0.53 \pm 0.45	0.77 \pm 0.35
Subsampling Joint	0.00 \pm 0.00	0.04 \pm 0.08	0.10 \pm 0.19	0.38 \pm 0.39	0.20 \pm 0.28	0.20 \pm 0.30
Subsampling	0.29 \pm 0.35	0.42 \pm 0.41	0.70 \pm 0.45	0.66 \pm 0.41	0.70 \pm 0.41	0.71 \pm 0.45
Under-sampling	0.04 \pm 0.13	0.45 \pm 0.45	1.00 \pm 0.00	1.00 \pm 0.00	0.74 \pm 0.45	1.00 \pm 0.00
Over-sampling	0.02 \pm 0.06	0.57 \pm 0.40	1.00 \pm 0.00	1.00 \pm 0.00	0.81 \pm 0.35	1.00 \pm 0.00

Table 10: BBSDs TPR across all datasets by quality and shift at 1000 samples.

Quality (mean accuracy)	0.5 (40%)	0.6 (44%)	0.7 (52%)	0.8 (63%)	0.9 (75%)	1.0 (87%)
Shift type						
Knock-Out	0.00 \pm 0.00	0.03 \pm 0.07	0.13 \pm 0.09	0.05 \pm 0.08	0.15 \pm 0.23	0.28 \pm 0.31
Only-One	0.00 \pm 0.00	0.08 \pm 0.10	0.39 \pm 0.29	0.51 \pm 0.49	0.57 \pm 0.50	0.59 \pm 0.49
Small Gaussian	0.00 \pm 0.00	0.05 \pm 0.07	0.42 \pm 0.27	0.70 \pm 0.34	0.76 \pm 0.35	0.82 \pm 0.34
Medium Gaussian	0.00 \pm 0.00	0.10 \pm 0.12	0.68 \pm 0.29	0.88 \pm 0.30	0.88 \pm 0.31	0.95 \pm 0.18
Adv. ZOO	0.00 \pm 0.00	0.03 \pm 0.07	0.17 \pm 0.15	0.15 \pm 0.31	0.19 \pm 0.35	0.21 \pm 0.39
Adv. Boundary	0.00 \pm 0.00	0.02 \pm 0.04	0.33 \pm 0.19	0.57 \pm 0.36	0.67 \pm 0.36	0.73 \pm 0.37
Subsampling Joint	0.00 \pm 0.00	0.03 \pm 0.07	0.13 \pm 0.10	0.03 \pm 0.07	0.00 \pm 0.00	0.04 \pm 0.08
Subsampling	0.00 \pm 0.00	0.00 \pm 0.00	0.13 \pm 0.10	0.03 \pm 0.07	0.04 \pm 0.10	0.10 \pm 0.15
Under-sampling	0.00 \pm 0.00	0.04 \pm 0.08	0.14 \pm 0.11	0.18 \pm 0.32	0.20 \pm 0.34	0.45 \pm 0.45
Over-sampling	0.00 \pm 0.00	0.04 \pm 0.08	0.16 \pm 0.10	0.23 \pm 0.35	0.36 \pm 0.38	0.52 \pm 0.42

References

- Alexandari, A., Kundaje, A., Shrikumar, A.: Maximum Likelihood with Bias-Corrected Calibration is Hard-To-Beat at Label Shift Adaptation. arXiv e-prints arXiv:1901.06852 (Jan 2019)
- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., Mané, D.: Concrete Problems in AI Safety. arXiv e-prints arXiv:1606.06565 (Jun 2016)
- Azizzadenesheli, K., Liu, A., Yang, F., Anandkumar, A.: Regularized learning for domain adaptation under label shifts. International Conference on Learning Representations (2019)
- Ben-David, S., Blitzer, J., Crammer, K., Pereira, F.: Analysis of Representations for Domain Adaptation. Proceedings of the 20th Annual Conference on Neural Information Processing Systems. (2006)
- Bethlehem, J.: Selection Bias in Web Surveys. International Statistical Review 78 (2), 161–188 (2010)
- Brendel, W., Rauber, J., Bethge, M.: Decision-based adversarial attacks: Reliable attacks against black-box machine learning models. 6th International Conference on Learning Representations, ICLR 2018 pp. 1–12 (2018)
- Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Computing Surveys (CSUR) (2009)

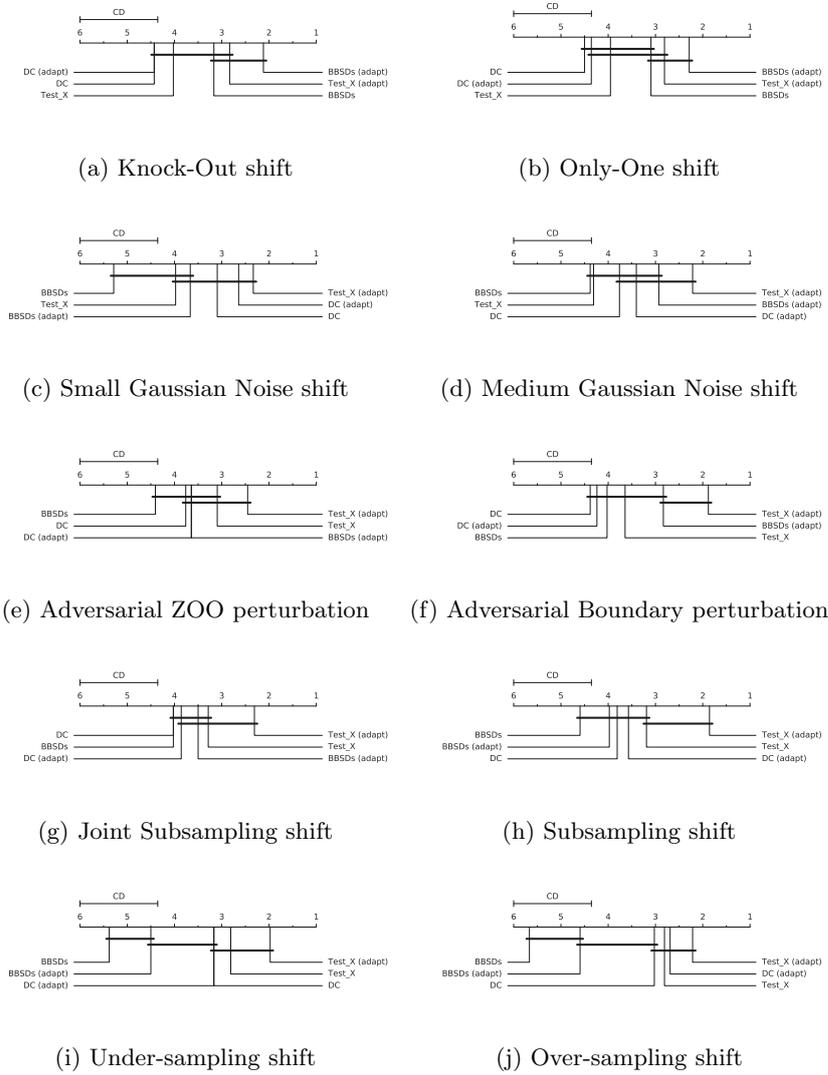


Fig. 7: Nemenyi post-hoc test comparing regular and adaptive significance level.

8. Chen, P.Y., Zhang, H., Sharma, Y., Yi, J., Hsieh, C.J.: ZOO: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. *AISec 2017 - Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security* pp. 15–26 (2017)
9. Demšar, J.: Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research* 7, 1-30 (2006)
10. Elshahar, H., Gallé, M.: To Annotate or Not? Predicting Performance Drop under Domain Shift. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on*

- Natural Language Processing (EMNLP-IJCNLP). pp. 2163–2173. Association for Computational Linguistics, Hong Kong, China (Nov 2019)
11. Garg, S., Wu, Y., Balakrishnan, S., Lipton, Z.C.: A Unified View of Label Shift Estimation. arXiv e-prints arXiv:2003.07554 (Mar 2020)
 12. Klaise, J., Van Looveren, A., Cox, C., Vacanti, G., Coca, A.: Monitoring and explainability of models in production. arXiv e-prints arXiv:2007.06299 (Jul 2020)
 13. Kouw, W.M., Loog, M.: A review of domain adaptation without target labels. arXiv e-prints arXiv:1901.05335 (Jan 2019)
 14. Lipton, Z.C., Wang, Y.X., Smola, A.: Detecting and Correcting for Label Shift with Black Box Predictors. arXiv e-prints arXiv:1802.03916 (Feb 2018)
 15. Markou, M., Singh, S.: Novelty detection: A review: Part 1: Statistical approaches. *Signal Processing* (2003)
 16. Moreno-Torres, J.G., Raeder, T., Alaiz-Rodríguez, R., Chawla, N.V., Herrera, F.: A unifying view on dataset shift in classification. *Pattern Recognition* 45 (1), 521–530 (2012)
 17. Nakagawa, S.: A farewell to Bonferroni: the problems of low statistical power and publication bias. *Behavioral Ecology* 15 (6), 1044–1045 (2004)
 18. Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J.V., Lakshminarayanan, B., Snoek, J.: Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift. arXiv e-prints arXiv:1906.02530 (Jun 2019)
 19. Perneger, T.V.: What’s wrong with Bonferroni adjustments. *BMJ (Clinical research ed.)* 316 (7139), 1236–1238 (1998)
 20. Quinero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: *Dataset Shift in Machine Learning* (2009)
 21. Rabanser, S., Günnemann, S., Lipton, Z.C.: Failing Loudly: An Empirical Study of Methods for Detecting Dataset Shift. arXiv e-prints arXiv:1810.11953 (Oct 2018)
 22. Redko, I., Morvant, E., Habrard, A., Sebban, M., Bennani, Y.: A survey on domain adaptation theory: learning bounds and theoretical guarantees. arXiv e-prints arXiv:2004.11829 (Apr 2020)
 23. Ruff, L., Kauffmann, J.R., Vandermeulen, R.A., Montavon, G., Samek, W., Kloft, M., Dietterich, T.G., Müller, K.R.: A Unifying Review of Deep and Shallow Anomaly Detection. arXiv e-prints arXiv:2009.11732 (Sep 2020)
 24. Schelter, S., Rukat, T., Biessmann, F.: Learning to Validate the Predictions of Black Box Classifiers on Unseen Data. In: *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. p. 1289–1299. SIGMOD ’20, Association for Computing Machinery, New York, NY, USA (2020)
 25. Shafaei, A., Schmidt, M., Little, J.J.: A less biased evaluation of out-of-distribution sample detectors. *30th British Machine Vision Conference 2019* (2018)
 26. Steinebach, J.: EL Lehmann, JP Romano: Testing statistical hypotheses. *Metrika* 64.2 pp. 255–256 (2006)