StaDRe and StaDRo: Reliability and Robustness Estimation of ML-based Forecasting using Statistical Distance Measures

Mohammed Naveed Akram[2]*, Akshatha Ambekar[3], Ioannis Sorokos[2]**, Koorosh Aslansefat[1], and Daniel Schneider[2]

[1] University of Hull, United Kingdom
[2] Fraunhofer IESE, Germany
[3] Technical University of Kaiserslautern, Germany

Abstract. Reliability estimation of Machine Learning (ML) models is becoming a crucial subject. This is particularly the case when such models are deployed in safety-critical applications, as the decisions based on model predictions can result in hazardous situations. As such, recent research has proposed methods to achieve safe, dependable and reliable ML systems. One such method is to detect and analyze distributional shift, and then measuring how such systems respond to these shifts. This was proposed in earlier work in SafeML. This work focuses on the use of SafeML for time series data, and on reliability and robustness estimation of ML-forecasting methods using statistical distance measures. To this end, distance measures based on the Empirical Cumulative Distribution Function (ECDF), proposed in SafeML, are explored to measure Statistical-Distance Dissimilarity (SDD) across time series. We then propose SDD-based Reliability Estimate (StaDRe) and SDD-based Robustness (StaDRo) measures. With the help of clustering technique, identification of similarity between statistical properties of data seen during training, and the forecasts is done. The proposed method is capable of providing a link between dataset SDD and Key Performance Indices (KPIs) of the ML models.

Keywords: SafeAI · Safe Machine Learning · Machine Learning Reliability · Artificial Intelligence Safety · Statistical Method

1 Introduction

The improved performance of hardware has triggered the increased use of ML-based components in many applications. This also includes many safety-critical application domains such as medical, industrial, transportation, aviation etc. This in turn requires new methods for increasing and assuring the reliability of ML-based components in safety-critical applications. One way to achieve this

---

* Corresponding author: email: naveed.akram@iese.fraunhofer.de

** Corresponding author: email: ioannis.sorokos@iese.fraunhofer.de

increase is by detecting dataset shift [32]. In previous work, SafeML has been introduced by the authors in [4], which uses statistical distance measures based on the Empirical Cumulative Distribution Function (ECDF) to detect distributional shift at runtime. Further, the approach estimates the performance degradation of a classifier based on those measures. The results showed that there is a correlation between performance and Statistical Distance Dissimilarity (SDD). The proposed method, however, was limited to only tabular data and classification tasks. In [3], SafeML has been improved to handle image classification tasks. Also, the authors have proposed to use a bootstrap-based p-value calculation to improve the reliability of statistical distance measures. In this work, we extend these measures for a time series data and use SDD to estimate reliability (StaDRe), and robustness (StaDRo) of a ML component.

The rest of the paper is divided as follows. Section 2 describes related literature, whereas in section 3 the problem statement is discussed. Section 4 details the proposed methodology, and section 5 describes the experiments conducted. Section 6 discusses the results, concluding the paper in 7.

2   Related Works

For an ML-based system to be used in safety-critical systems, certain qualities of ML solutions such as reliability, fairness, robustness, transparency and security are important. Several methods have been proposed to quantify and qualify these properties. For this work, the properties of ML-based system resulting from the dataset distribution and shifts are considered.

This section is divided into several subsections. Section 2.1 discusses the available methods for detecting dataset shift and distributional shift. Section 2.2 discusses the existing reliability metrics available in the literature. And section 2.3 discusses the existing robustness evaluation methods.

2.1   Dataset shift

There exist several names for the concept of dataset shift in the literature. Some of these are concept shift, contrast mining (in classification learning, fracture between data and fracture points and changing environments [23]). To overcome the inconsistencies, authors in [23] have defined dataset shift as the difference in the joint distribution in training and runtime datasets. Additionally, they identify three types of dataset shift: A covariate shift, prior probability shift, and a concept shift.

Several techniques for dataset shift detection are proposed in the literature for forecasting and classification problems. The work in [6] proposes measures that are applicable only to linear and multiple linear regression models, work on exponentially-weighted moving average (EWMA) charts [29] is limited by analysis power in presence of larger shifts, the proposed technique in [26] for image classification problem is restricted by lack of support for online dataset shift detection and the technique based on Particle Swarm Optimization (PSO)

[25] overshadows its usefulness due to an overhead of retraining every time a concept drift occurs. In addition to these limitations, the existing methods do not consider statistical distance measures for computing distributional shift in time series.

2.2   Reliability

A widely regarded work in dependable systems, the authors in [5], define reliability as continuity of correct service. In classical systems, quality measures such as Mean Time Between Failures (MTBF) is usually used for reliability measurement. For an ML system, authors in [9] define reliability by any qualitative attribute related to a "vital performance indicator" of that system. They included measures such as accuracy or inaccuracy, availability, downtime rate, responsiveness etc. In the context of ML-model-driven systems, we consider continuity of correct service, or ML performance evaluation, to not only encompass the service provision, but also correctness of the ML model response itself.

The existing approaches proposed in literature for reliability estimation can be classified into two classes: model-agnostic or model-specific approaches. While model-specific techniques [16,24] are based on specific model designs and learning algorithms, they utilize probabilistic interpretations for reliability, in terms of confidence measurement. In our work, we focus on model-agnostic techniques.

The model-agnostic approaches are generic and provide distributional or indicator-based reliability estimates that are less probabilistically interpretable. Some examples of model-agnostic reliability estimation techniques are sensitivity analysis [8], local cross-validation [13], confidence estimation based on neighbors' errors and variance in the environment [11], bootstrapping [15] and an ML-based approach [1]. Authors in [11] propose the *"CONFINE"* metric to measure reliability. In contrast to previous heuristic-based techniques, the CONFINE measure incorporates actual model prediction errors for reliability estimation, and estimations are based on the error of the test instance's nearest neighbors. However, this metric does not cover the effect of distributional shift on reliability. Hence, we propose extending the CONFINE metric with the rate of change of dissimilarity.

2.3   Robustness

In machine learning literature, different aspects of model's robustness such as robustness against adversarial attacks and robustness against dataset shift are evaluated. To begin, adversarial robustness is used to measure the robustness of an ML model against adversarial attacks. The authors of [28] describe model robustness as the average size of the minimum adversarial perturbation over many samples. The majority of work in this field focuses on various kinds of adversarial attacks and building ML models that are robust against these.

However, in this work, we focus on robustness against dataset shift. A well-generalized network should be robust to data subpopulations it has not yet

encountered [31]. A measure termed *relative accuracy* to compute model's performance is presented in work [31] with a focus on image classification problems. Furthermore, the work in [17] argues that accuracy can suffer when the train and test distributions are dissimilar, therefore implement a data augmentation technique at training to improve robustness of image classifiers, as well as use *flip probability* as a measure of perturbation robustness. The authors of [33] introduced measure called *effective robustness* which is the model accuracy difference in the presence of distributional shift, taking into consideration an accuracy baseline determined using available standard models (without robustness interventions). The authors of [4] examined the use of statistical distance measures to evaluate the robustness of ML models. The work in [19] propose verification of ML models in safety-critical applications, by verifying generalized model performance using Neuron Region Distance (NRD) as an aid for dissimilarity measure.

A ML model that is robust should be able to deliver acceptable performance under conditions like dataset shift. This notion of robustness is also closely related to definition of robustness presented in [20] and [5] for classical systems. Though the existing methods cover a plethora of robustness measures, currently there are not many that use statistical distance measures to account for robustness against dataset shift. In our work, we propose using a performance vs dissimilarity curve for robustness evaluation.

3   Problem Definition

When an ML model encounters an input differing from its training data, out-of-distribution (OOD) uncertainty emerges. This can be taken into account by estimating the SDD, with which reliability and robustness can be estimated. Hence, in this work, the following research questions are established:

RQ1: SDD-Accuracy Correlation. How does SDD relate to the performance of a model specifically for time series application?

RQ2: SDD based Reliability and Robustness. How can SDD be incorporated in reliability and robustness measures of a model?

4   Proposed Method

In this section, we use the SDD measure to obtain the StaDRe and StaDRo. Finally, we describe the experimental design for our evaluation of the proposed approaches on a univariate time series dataset.

For the rest of the paper, we will use the following terms. Let us assume $X_{all} \in \mathbb{R}$ be the overall univariate time series dataset, $X \in \mathbb{R}$ be the training portion of it, and $Y \in \mathbb{R}$ be the validation part. At runtime, we use X^* for then available data. Let f be any ECDF-based distance measure such as Wasserstein, Kolmogorov-Smirnov etc. Let M be the ML model function, and $Y' \in \mathbb{R}$ be the

prediction for a given input x. Due to its unbounded nature, geometry considerations, and method which is essentially the area measured between two ECDFs, this work primarily uses Wasserstein distance as ECDF function for estimating reliability and robustness [7].

4.1 SDD vs Performance

First, let us consider the behavior of the ML model performance with its SDD against the data observed during training. To observe this behavior, the performance of the validation data Y is compared against its SDD from the training set ($f(X, Y)$). The performance can be measured either in root mean square error (RMSE), or mean absolute percentage error (MAPE). To obtain a sequence of observations, the validation set Y is divided into several subsets $Y^* \in Y$ of length l. The performance of the model M can be obtained, by obtaining the prediction over each subset $M(Y^*) \ \forall Y^* \in Y$. The corresponding SDD of subset Y^* is obtained using $f(X, Y^*)$. From these, a curve of performance vs SDD can be plotted. Section 5 discusses the detailed experiments.

4.2 Statistical Distance based Reliability estimate - StaDRe

In [11], reliability of ML model is given in terms of confidence, by "CONFidence estimation based on the Neighbors' Errors" (CONFINE) as in equation 1.

$$cs_{CONFINE}(X^*) = 1 - \frac{1}{m} \sum_{i=1}^{m} \epsilon_i^2 \,[11] \tag{1}$$

Where, X^* is a data instance for which reliability is needed, ϵ_i is the error of a nearest neighbor i, m are the total nearest neighbors. However, this does not take into account the SDD of the instance. Moreover, CONFINE uses Euclidean distance to obtain the nearest neighbors.

 We extend this approach by taking into account the SDD, and using the DTW-based clustering for computing nearest neighbors of X^*. Clustering is a technique commonly used to group together similar patterns in time series using distance measures like DTW, Fréchet, Kullback-Leibler divergence [2,14,22]. Among these, K-Means clustering [10] with a DTW distance metric [27,30] was found to be adequate for univariate time series. Due to space limitations, the explanation of the DTW-based clustering technique used is omitted, but is available on GitHub (see section 7). Instead of using Euclidean distance, we use the DTW-based clustering technique to identify the appropriate cluster c for the input instance X^* obtained from Y'. Then, the mean squared error (MSE) is obtained for the m members of the identified cluster c. Finally, the ECDF-based distance is computed between the cluster c and X^*, and the cluster c and the reference O. The StaDRe is then given by equation 2. The algorithm 1 describes the general procedure for StaDRe computation using statistical distance measures. For this work, the following assumptions are made.

1. The members of the cluster c to which X^* gets assigned to, are the nearest neighbors of X^*, where X^* is a data instance.
2. To compute d_{origin} in algorithm 1 initialised for Wasserstein distance, the origin is a time series containing 0 at every time step and whose shape is the same as the cluster's center.

$$StaDRe(X^*) = \frac{2 - \dfrac{1}{m}\sum_{i=1}^{m}\epsilon_i^2 - \dfrac{f(c,X^*)}{f(c,O)}}{2} \qquad (2)$$

Where, ϵ_i is the error of a given neighbor i.

Algorithm 1: Algorithm for StaDRe

Result: Reliability
X = GetTrainingSet();
X^* = GetDataInstance();
clusters = GetClusters(X);
for c *in clusters* **do**
    comparison = CompareClusters(c, X^*);
    if *comparison==True* **then**
        d = getStatisticalDistance(c, X^*);
        d_origin = getStatisticalDistance(c,O);
        break;
    end
end
MeanSquaredError = GetClusterMeanSquareError(c);
Reliability = $\dfrac{(2 - MeanSquaredError - \dfrac{d}{d\_origin})}{2}$;
return Reliability;

4.3   Statistical Distance based Robustness - StaDRo

In this section, we introduce a method to obtain an evaluation of robustness against distributional shift using statistical distance measures. For this, the performance vs SDD curve obtained in 4.1 is used. For a required minimum performance P_{min}, using the performance vs SDD curve, the corresponding SDD distance for the P_{min} is obtained. For a given data instance X^* for which robustness is required, the ECDF distance measure is obtained against the instance X^* and the training set X. StaDRo for an instance X^* at a given P_{min} is given by the equation 3.

$$StaDRo(X^*, P_{min}) = \begin{cases} True & \dfrac{f(X,X^*)}{d_{Pmin}} \leq 1 \\[2em] False & \dfrac{f(X,X^*)}{d_{Pmin}} > 1 \end{cases} \qquad (3)$$

Where, $StaDRo(X^*, P_{min})$ gives the robustness of the model for the instance X^* at a given minimum required performance P_{min}. The algorithm 2 shows a general procedure for StadRo estimation using statistical distance measures.

Algorithm 2: Algorithm for StaDRo

Result: Robust
X = GetTrainingSubSet();
Y = GetValidationSet();
for *Sequence in Y* **do**
  P = GetPerformance(Sequence);
  d = getStatisticalDistance(X, Sequence);
  Performance.append(P);
  SDD.append(d);
end
Curved = FitCurve(plot(SDD, Performance));
X^* = GetDataInstance();
d_instance = getStatisticalDistance(X, X^*);
MINIMUM_REQUIRED_PERFORMANCE_LIMIT = CONSTANT;
d_from_curve = getDistanceForPerformance(Curved,
  MINIMUM_REQUIRED_PERFORMANCE_LIMIT);
Ratio = d_instance / d_from_curve;
if *Ratio <= 1* **then**
  Robust = True;
else
  Robust =False;
end
return Robust;

---

5 Experiments

We performed experiments on stock price prediction application (an example of regression applications) and used stock price data of several companies. The closing price of these stocks is considered. We extracted this using the National Stock Exchange (NSE) library and Yahoo Finance. We excluded the dataset that did not exhibit any kind of shift between data at training and validation. The data was divided into 80% training data and 20% was used for the validation. We predicted the stock closing price for the next day, at a given point the model is used. All the data was normalized between -1 and 1. The two forecasting models we implemented are Long Short-Term Memory (LSTM) [18] and Gated Recurrent Unit (GRU) [12]. The choice of hyperparameters and other information is given in Table 1. All the models are trained using the Adam optimizer [21] with a learning rate of 0.01 over 100 epochs for LSTM and 50 epochs for GRU.

To obtain StaDRo as described in Algorithm 2, the experiments are divided into three phases. In the first phase, we measure how accuracy *acc* of the model M changes with shift in the subset of the validation set Y. In the second phase,

$f(X, Y)$ is computed for the same subsets of Y as in the first phase. For this work, we instantiated Wasserstein distance as an ECDF measure f. In the third phase, we fit a polynomial curve of degree 2 to derive the relationship between dataset SDD and the forecasting model's performance. We then obtain a statistical distance for a desired minimum performance from the fitted curve. For a desired sequence X^*, we compute the SDD between X^* and X, then calculate the ratio and estimate robustness.

For reliability estimation using the general procedure described in Algorithm 1, we applied clustering on training data of several companies and then selected some validation points randomly at different time steps to compute StadRe for these points. When computing MSE, we chose the number of nearest neighbors (m) as equal to the number of samples of X that get assigned to the smallest cluster c'. We applied random sampling for selecting m neighbors when the assigned cluster c is not c'.

TABLE 1: Details of several stocks and hyperparameters of LSTM and GRU model.

Characteristic	Dataset				
	Reliance	*Google*	*Airtel*	*JP Morgan*	*MRF*
Start date (dd.mm.yyyy)	04.01.2010	03.01.2005	04.01.2010	02.01.2003	03.01.2005
No. of Data points	2895	4300	3008	4803	4214
Window size	20	50	20	20	20
No. of clusters	7	5	6	5	5
LSTM *(# layers; hidden size)*	2; 32	2; 40	2; 32	2; 64	2; 32
GRU *(# layers; hidden size)*	2; 32				

6   Results and Discussion

This section presents the results of the experiments performed. Figure 1 shows how model performance (as MAPE, RMSE) deteriorates as SDD increases. Each performance measure is fit to a curve using second-degree curve fitting. As clearly seen in Figures 1(a) and 1(b), there is an increase in performance error with increasing SDD. Figure 2 displays the reliability measure against model performance. As with Figure 1, the plots are fitted with a second-degree curve. The figures clearly shows increasing performance with increasing reliability, StaDRe, thereby validating the measure.

Table 2 shows the results of the StaDRo computation on some example stocks. The table consists of a sequence of validation sets across 2 stocks, corresponding performance metrics (RMSE, MAPE) and StaDRo robustness results for each metric. It can be seen from the table that for a validation subset which StaDRo identifies as not robust against SDD (listed as 'FALSE'), a higher error is observed. This indicates that StaDRo is a useful measure for robustness against statistical dissimilarity in datasets.
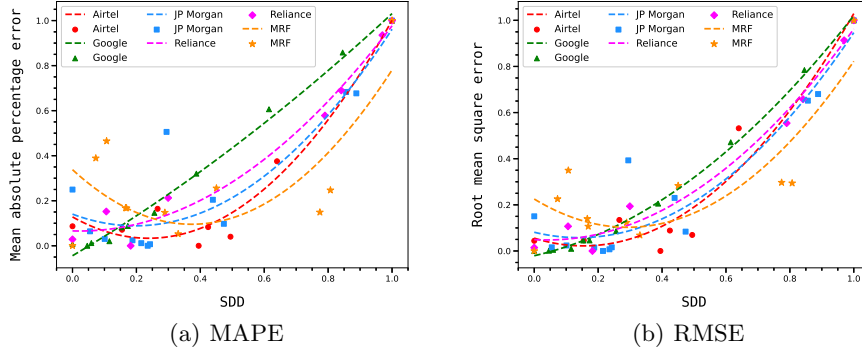
FIG. 1: The curve displaying performance vs SDD for several stocks normalized to 0-1 range.
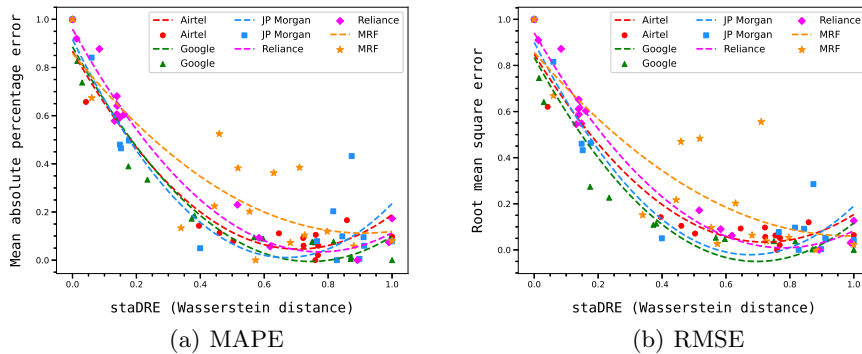

FIG. 2: The curve displaying performance vs StaDRe for several stocks normalized to 0-1 range.

7   Conclusion and Future Works

In this paper, we describe the use of SafeML for SDD-based assessment of ML models for time series applications. The performance behavior against SDD is investigated. It was found that the performance is correlated to the SDD. We propose SDD-based metrics for reliability as StaDRe and robustness as StaDRo. The results on example applications shows the effectiveness of these measures on univariate time series.

For future work, we are planning to experiment with the proposed measures on multivariate time series for both classification and regression tasks. Adversarial attacks can also be used to consider StaDRo for robustness against adversarial examples as well. The idea of this paper can be extended for videos (image-time-

TABLE 2: Results displaying StaDRo estimates on some stocks

Data Instance	Google (P_{min} (RMSE) = 500.0)				JP Morgan (P_{min} (RMSE) = 8.5)			
	RMSE	WD	Rate of change of SDD	Robust	RMSE	WD	Rate of change of SDD	Robust
0:70	24.85	634.73	0.39	TRUE	2.23	61.25	0.63	TRUE
70 : 140	31.12	740.50	0.45	TRUE	2.32	62.63	0.64	TRUE
140 : 210	26.35	717.43	0.44	TRUE	2.53	53.44	0.55	TRUE
210 : 280	36.73	842.07	0.52	TRUE	2.41	59.38	0.61	TRUE
280 : 350	82.03	910.24	0.56	TRUE	2.43	63.12	0.65	TRUE
350 : 420	82.46	945.59	0.58	TRUE	3.28	78.90	0.81	TRUE
420 : 490	133.10	1094.77	0.67	TRUE	7.16	66.65	0.68	TRUE
490 : 560	284.06	1330.72	0.82	TRUE	4.12	46.56	0.48	TRUE
560 : 630	620.01	1739.06	1.07	FALSE	2.43	50.30	0.52	TRUE
630 : 700	1014.97	2153.03	1.32	FALSE	5.11	76.55	0.78	TRUE
700 : 770	1287.46	2430.59	1.49	FALSE	10.40	105.03	1.08	FALSE
770 : 840	-	-	-	-	10.76	107.21	1.1	FALSE
840 : 910	-	-	-	-	14.77	114.84	1.18	FALSE

series-based inputs) and tested for autonomous vehicles in CARLA[4] simulation. Moreover, the proposed idea can be used as a foundation for time-series safety-related explainability and interpretability.

All results, code, additional results, additional experiments, and exhaustive explanations will be made available in GitHub[5] repository.

Acknowledgements

References

1. Adomavicius, G., Wang, Y.: Improving reliability estimation for individual numeric predictions: a machine learning approach. INFORMS Journal on Computing (2021)
2. Anh, D.T., Thanh, L.H.: An efficient implementation of k-means clustering for time series data with dtw distance. International Journal of Business Intelligence and Data Mining **10**(3), 213–232 (2015)
3. Aslansefat, K., Kabir, S., Abdullatif, A., Vasudevan, V., Papadopoulos, Y.: Toward improving confidence in autonomous vehicle software: A study on traffic sign recognition systems. Computer **54**(8), 66–76 (2021)

---

[4] www.carla.org

[5] https://github.com/n-akram/TimeSeriesSafeML

4. Aslansefat, K., Sorokos, I., Whiting, D., Tavakoli Kolagari, R., Papadopoulos, Y.: Safeml: Safety monitoring of machine learning classifiers through statistical difference measures. In: International Symposium on Model-Based Safety and Assessment. pp. 197–211. Springer (2020)

5. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. IEEE transactions on dependable and secure computing **1**(1), 11–33 (2004)

6. Becker, A., Becker, J.: Dataset shift assessment measures in monitoring predictive models. Procedia Computer Science **192**, 3391–3402 (2021). https://doi.org/https://doi.org/10.1016/j.procs.2021.09.112, `https://www.sciencedirect.com/science/article/pii/S1877050921018512`, knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 25th International Conference KES2021

7. Bellemare, M.G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., Munos, R.: The cramer distance as a solution to biased wasserstein gradients. arXiv preprint arXiv:1705.10743 (2017)

8. Bosnić, Z., Kononenko, I.: Estimation of individual prediction reliability using the local sensitivity analysis. Applied intelligence **29**(3), 187–203 (2008)

9. Bosnić, Z., Kononenko, I.: An overview of advances in reliability estimation of individual predictions in machine learning. Intelligent Data Analysis **13**(2), 385–401 (2009)

10. Bradley, P.S., Fayyad, U.M.: Refining initial points for K-Means clustering. In: Proc. 15th International Conf. on Machine Learning. pp. 91–99. Morgan Kaufmann, San Francisco, CA (1998)

11. Briesemeister, S., Rahnenführer, J., Kohlbacher, O.: No longer confidential: estimating the confidence of individual regression predictions. PloS one **7**(11), e48723 (2012)

12. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: Encoder–decoder approaches. In: Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation. pp. 103–111. Association for Computational Linguistics, Doha, Qatar (Oct 2014). https://doi.org/10.3115/v1/W14-4012, `https://aclanthology.org/W14-4012`

13. DEMUT, I.R.: Reliability of predictions in regression models. Doktorandske dny'10 (2010)

14. Driemel, A., Krivošija, A., Sohler, C.: Clustering time series under the fréchet distance. In: Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms. pp. 766–785. SIAM (2016)

15. Efron, B.: Bootstrap methods: another look at the jackknife. In: Breakthroughs in statistics, pp. 569–593. Springer (1992)

16. Gammerman, A., Vovk, V., Vapnik, V.: Learning by transduction. CoRR **abs/1301.7375** (2013), `http://arxiv.org/abs/1301.7375`

17. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: A simple data processing method to improve robustness and uncertainty. arXiv preprint arXiv:1912.02781 (2019)

18. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Comput. **9**(8), 1735–1780 (nov 1997). https://doi.org/10.1162/neco.1997.9.8.1735, `https://doi.org/10.1162/neco.1997.9.8.1735`

19. Hond, D., Asgari, H., Jeffery, D., Newman, M.: An integrated process for verifying deep learning classifiers using dataset dissimilarity measures. International Journal of Artificial Intelligence and Machine Learning (IJAIML) **11**(2), 1–21 (2021)

20. IEEE: Standard glossary of software engineering terminology. IEEE Std 610.12-1990 pp. 1–84 (1990). https://doi.org/10.1109/IEEESTD.1990.101064
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
22. Lee, T., Xiao, Y., Meng, X., Duling, D.: Clustering time series based on forecast distributions using kullback-leibler divergence. In: International Institute of Forecasters (IIF). Web (2014)
23. 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)
24. Nouretdinov, I., Melluish, T., Vovk, V.: Ridge regression confidence machine. In: Proceedings of the Eighteenth International Conference on Machine Learning. pp. 385–392 (2001), ed. by C. E. Brodley and A. P. Danyluk. San Francisco, CA: Morgan Kaufmann.
25. Oliveira, G.H., Cavalcante, R.C., Cabral, G.G., Minku, L.L., Oliveira, A.L.: Time series forecasting in the presence of concept drift: A pso-based approach. In: 2017 IEEE 29th International Conference on Tools with Artificial Intelligence (ICTAI). pp. 239–246. IEEE (2017)
26. Rabanser, S., Günnemann, S., Lipton, Z.: Failing loudly: An empirical study of methods for detecting dataset shift. Advances in Neural Information Processing Systems **32** (2019)
27. Ratanamahatana, C.A., Keogh, E.: Everything you know about dynamic time warping is wrong. In: Third workshop on mining temporal and sequential data. vol. 32. Citeseer (2004)
28. Rauber, J., Brendel, W., Bethge, M.: Foolbox: A python toolbox to benchmark the robustness of machine learning models. arXiv preprint arXiv:1707.04131 (2017)
29. Raza, H., Prasad, G., Li, Y.: Ewma model based shift-detection methods for detecting covariate shifts in non-stationary environments. Pattern Recognition **48**(3), 659–669 (2015)
30. Sakoe, H., Chiba, S.: Dynamic programming algorithm optimization for spoken word recognition. IEEE Transactions on Acoustics, Speech, and Signal Processing **26**(1), 43–49 (1978). https://doi.org/10.1109/TASSP.1978.1163055
31. Santurkar, S., Tsipras, D., Madry, A.: Breeds: Benchmarks for subpopulation shift. arXiv preprint arXiv:2008.04859 (2020)
32. Storkey, A.: When training and test sets are different: characterizing learning transfer. Dataset shift in machine learning **30**, 3–28 (2009)
33. Taori, R., Dave, A., Shankar, V., Carlini, N., Recht, B., Schmidt, L.: Measuring robustness to natural distribution shifts in image classification. Advances in Neural Information Processing Systems **33**, 18583–18599 (2020)