

A Deep Dive into Perturbations as Evaluation Technique for Time Series XAI

Udo Schlegel^(✉)  and Daniel A. Keim 

University of Konstanz, Universitätsstraße 10, 78464 Konstanz, Germany
{u.schlegel,daniel.keim}@uni-konstanz.de

Abstract. Explainable Artificial Intelligence (XAI) has gained significant attention recently as the demand for transparency and interpretability of machine learning models has increased. In particular, XAI for time series data has become increasingly important in finance, healthcare, and climate science. However, evaluating the quality of explanations, such as attributions provided by XAI techniques, remains challenging. This paper provides an in-depth analysis of using perturbations to evaluate attributions extracted from time series models. A perturbation analysis involves systematically modifying the input data and evaluating the impact on the attributions generated by the XAI method. We apply this approach to several state-of-the-art XAI techniques and evaluate their performance on three time series classification datasets. Our results demonstrate that the perturbation analysis approach can effectively evaluate the quality of attributions and provide insights into the strengths and limitations of XAI techniques. Such an approach can guide the selection of XAI methods for time series data, e.g., focusing on return time rather than precision, and facilitate the development of more reliable and interpretable machine learning models for time series analysis.

Keywords: Explainable AI · XAI Evaluation · XAI for Time Series

1 Introduction

Artificial intelligence (AI) has become an integral part of our daily lives, from the personalized advertisement we receive on social media to conversational AI (chatbots) answering questions of users and customers using deep neural networks. However, as the complexity of deep neural network models increases, so does the difficulty in understanding how they get to their decisions [7]. A lack of interpretability can lead to severe consequences in critical domains such as finance, healthcare, and transportation, including financial losses, medical errors, and even loss of life by providing wrong decisions if complex models are deployed [16]. One promising approach to addressing such issues is through the usage of explainable artificial intelligence (XAI), which seeks to provide insights into the inner workings of complex models and the factors that drive

their decision-making [7]. One particular area of interest is time series data, which is characterized by the sequential nature of its observations and the interdependencies between them, as more sensors generate a massive amount of data and more tasks are tackled by complex models [25].

In recent years, a growing body of research has focused on developing XAI techniques tailored explicitly for time series data [25]. These techniques often rely on the concept of attributions, which aim to identify the contributions of individual features and time points to the overall prediction made by a model [25]. By providing insights into which parts of the input data are most relevant to the output, attributions can help users understand the reasoning behind the model’s decision-making process [18]. However, the evaluation of such attributions is not trivial [19]. To address the challenge of evaluating the quality of explanations for time series data, perturbation analysis has emerged as a promising evaluation technique [17, 22]. Perturbation analysis involves systematically modifying the input data and assessing the impact on the attributions generated by XAI methods [19]. By perturbing the input data, it is possible to evaluate the robustness of the explanations provided by XAI methods [25]. However, the effectiveness of perturbation analysis for evaluating the quality of attributions for time series data has not been extensively studied [19].

In this paper, we apply attribution techniques from various fields to a convolution neural network trained on time series classification data to evaluate and inspect the generated attributions in detail using perturbations, which involves systematically altering the input data and observing the effect on the model’s output. We investigate the performance of attribution techniques compared to each other based on the perturbation analysis result and explore the perturbation changes based on these attributions to gain insights into the model. Through such an analysis, we can identify spurious correlations and shortcuts in the complex models and thus enable developers to potentially improve models by debugging datasets. We show that our convolution neural network trained on time series classification learned certain shortcuts to achieve state-of-the-art performances. Based on these experiments and results, we provide guidelines for the application of attribution techniques for time series classification and release our evaluation framework to investigate other attribution techniques.

Thus, we contribute: (1) an in-depth analysis of attribution techniques on time series classification for deep learning models using a perturbation analysis, (2) insights into convolution neural networks trained on time series based on the generated attributions, (3) guidelines and a framework for applying attribution techniques for time series models with a perturbation analysis card for reporting. We first look into related work, and then we introduce the perturbation analysis methodology and the experiment setup we use for our deep dive. Here we also propose perturbation analysis cards as a guideline to report the results of an evaluation. Next, we present our results and discuss the impact of our conclusions for attribution techniques applied to time series. Lastly, in future work, we motivate new measures for the evaluation of attributions on time series data.

Results and source code of the experiments is online available at: <https://github.com/visual-xai-for-time-series/time-series-xai-perturbation-analysis>.

2 Related Work

Explainable AI (XAI) accelerated through several surveys [1,7] and techniques, e.g., LIME [15] and SHAP [12] in the last few years. Especially, attributions are prevalent in the image domain as heatmap explanations are easy to understand for users [10]. Some theoretical works dig into the backgrounds of why models learn certain shortcuts to solve tasks [6] and thus enable further explanations for decisions. However, evaluating explanations is still a slowly growing area with limited work toward benchmarking different techniques against each other [8]. Further, shortcuts or spurious correlations are not trivial to detect in explanations and need an in-depth analysis to be able to identify these [29].

Some works started to collect possible evaluation techniques [14] and categorized these into five measurements: mental model, explanation usefulness and satisfaction, user trust and reliance, human-AI task performance, and computational measures. The first few measures focus on evaluating with or in cooperation with humans and are thus heavily influenced by human factors. The computational measures exclude human factors and focus on purely automatic evaluation of explanations. In this work, we inspect the computational measures and, more precisely, the explainer fidelity of the attribution technique on the model to show how the attributions fit the model.

XAI for time series classification (TSC), on the one hand, incorporates previously proposed explanation techniques from other fields and introduces the time dependence into some of the techniques [25]. Theissler et al. [25] categorize possible explanations for TSC into time point, subsequence, and instance explanations. All these operate on a different level of the time series and are thus unique in their explanation and evaluation. In this work, we tackle time point explanations and, to be more precise, attributions to highlight and explore shortcuts and spurious correlations. As Schlegel et al. [17] and others [13,22,25] demonstrated, attributions techniques such as LIME [15], SHAP [12], LRP [4], GradCAM [21], Integrated Gradients [24], and more [20], produce working attributions on time series to extract explanations from a model. However, in most cases, only purely computational measures are applied to the attributions, which are not further inspected, e.g., Mercier et al. [13] to gain deeper insights.

Schlegel et al. [17] started by using a perturbation analysis on attribution techniques applied to TSC using various perturbation functions to highlight that techniques for images and text are also working on time series. Based on such preliminary experiments, they enhanced their approach with additional perturbation functions to showcase deeper insights into the fidelity evaluation [19]. Mercier et al. [13] enhanced these perturbations with further measures from the image domain, such as (in)fidelity and sensitivity [27]. Simic et al. [22] extended the proposed methods by Schlegel et al. [19] with out-of-distribution detecting functions and gave guidelines for the selection of attribution techniques and the size of the window for the perturbation. Turbe et al. [26] enhance previous approaches with another metric to improve the comparison of the attribution techniques and the ability to demonstrate their fidelity towards the model. However, all of these approaches do not look into the attributions and the produced

values to investigate further into the techniques behind the attributions and the models. Thus, an in-depth analysis is needed to investigate the attributions generated for time series classification models.

3 Perturbation Analysis

We use the perturbation analysis approach by Schlegel et al. [19] to generate attributions, verify, and compare them using the proposed perturbation function strategies [17, 22]. We extend the comparison by calculating the Euclidean and cosine distance between the original and the perturbed time series instance and the Euclidean and cosine distance between the original attributions of the dataset and the attributions of the perturbed instances of the dataset. Collecting these results can help us have a more in-depth analysis of the attribution techniques and reveal relations between attributions and models. However, we first need to establish the general perturbation analysis.

Let $D = (X, Y)$ be a time series classification dataset with X as the time series samples and Y as the time series labels. $X = \{ts_1, ts_2, \dots, ts_n\}$ contains n time series samples with m time points for each sample represented as $ts = \{tp_1, tp_2, \dots, tp_m\}$, where tp_1 is the value for the i th time point of ts . $Y = \{l_1, l_2, \dots, l_n\}$ contains n labels one label for each time series sample. Let $M(ts, \theta) = y'$ be a time series classification model which predicts a label y' based on a time series input ts and has the parameters θ . Let $A(X, M, \theta)$ be an XAI technique for generating attributions for the time series data. The original attributions for X generated by A can be represented as $A(X, M, \theta) = \{a_1, a_2, \dots, a_m\}$, where a_i is the attribution score for the i th time point of X , M the time series classification model for which the attributions are calculated, and θ the parameters of the attribution technique.

To perform perturbation analysis, we introduce a perturbation function g that modifies X in a controlled manner. Specifically, we define a perturbed time series dataset X' as:

$$X' = g(X, A, \xi) \quad (1)$$

Our perturbation function g modifies the dataset X based on the attributions A and a threshold ξ . The value for the modification can be changed and depends on the selected function g , e.g., exchange to zero. The threshold ξ can be set to a value by hand or some other function, e.g., using the 90-percentile of the attributions so that the attributions, e.g., a_i the i th element, above the threshold, will be modified to the previously set value, e.g., zero. Figure 1 demonstrates the approach with zero perturbations on attributions with high values.

The original X and the perturbed dataset X' get predicted with the model M to get $M(X) = Y'$ and $M(X') = Y''$. Based on Schlegel et al. [19], we incorporate a quality metric qm , e.g., accuracy, to compare the performance of the model M with the original X and the perturbed dataset X' . For the time series classification, we assume that the qm decreases after the original data changes, and thus the labels are not fitting anymore [17]. We further assume




Fig. 1. Starting from a time series ts , we use a selected attribution technique A to get attributions. Based on the attributions, we use a selected perturbation function g to set highly relevant time points, e.g., to zero. Further information in Schlegel et al. [19].

a suiting attribution technique decreases the performance more heavily as the most relevant parts of the input data get perturbed [8]. Thus, we assume:

$$qm(Y', Y) \leq qm(Y'', Y) \quad (2)$$

However, in some cases, the scores are very similar [13, 17], and a deeper investigation into the attributions is necessary to find similarities or dissimilarities in the relevances of the techniques. Thus, we do not only compare the quality metrics but also focus on the distances between the original X and the perturbed X' datasets. We apply the Euclidean and cosine distances to the datasets as these are common distance functions for time series [2] to collect the changes of the perturbation function g . We define the Euclidean distance as:

$$Euc(X, X') = \sqrt{\sum_{i=1}^n (ts_i - ts'_i)^2} \quad (3)$$

where $X = ts_1, ts_2, \dots, ts_n$ and $X' = ts'_1, ts'_2, \dots, ts'_n$ are the two time series being compared. And we define the cosine distance as:

$$Cos(X, X') = 1 - \frac{\sum_{i=1}^n ts_i \times ts'_i}{\sqrt{\sum_{i=1}^n ts_i^2} \times \sqrt{\sum_{i=1}^n ts_i'^2}} \quad (4)$$

where $X = ts_1, ts_2, \dots, ts_n$ and $X' = ts'_1, ts'_2, \dots, ts'_n$ are the two time series being compared. These changes enable us to compare the attributions not only on a performance level but on a raw level directly on the data.

4 Experiments with Perturbation Analysis

For our analysis, we look into the time series that changed and those that did not change during the perturbation analysis. We especially want to understand the attribution distributions to investigate the attribution techniques responsible for fitting explanations, with high fidelity [14], on the models. Fitting explanations in our assumptions are techniques that change the prediction of more samples in a perturbation analysis [13, 17, 19]. However, a general measure and metric for evaluating explanations are essential, but another factor is the attributions, as these can also hide information or present spurious correlations [29]. E.g., the question of how attributions are distributed over the techniques arises.

To answer such questions and others, we use the changes from Y (old prediction) to Y' (new prediction) to look into the samples that changed their prediction and those that do not change. We especially want to know when a change in the prediction happened, e.g., after how many removals based on the attributions and the perturbation strategy. Thus, we look at the prediction changes from one class to the other. E.g., in a binary classification with the assumption from above, the predictions change from one to the other class to demonstrate that the attributions highlight relevant time points for the model. Thus, we slowly perturb more and more values from the time series until there is a change in prediction. We use the percentile values (99, 98, ..., 25) as a threshold for the perturbation and record when the change happens. Further, we collect the skewness of the attributions of the changed and unchanged predictions. With such an exploration of the distributions of the attributions, we enable to inspect patterns inside of the attributions generated by different techniques. Also, the distributions of the skewness enable to have another factor for the comparison of the attribution techniques. Lastly, we do not only collect the skewness but also the Euclidean and the cosine distances of the original sample to the perturbed instance with regard to the changed and unchanged predictions. All these different collected statistics and properties can help us to identify various findings, insights, and correlations in the attribution techniques as we collect as much data from our perturbation analysis as possible.

Summary – Overall, we have the following dimensions we want to experiment on: a) attribution techniques, b) perturbation strategy. We collect and analyze the following properties: a) mean of the raw data samples of the changed and unchanged predictions; b) skewness of attributions based on the changed and unchanged predictions after the perturbation; c) new class distributions of the changed and unchanged predictions after the perturbation; d) amount of relevant attributions needed to perturb an instance to another class prediction. Figure 2 presents the collected properties using a perturbation analysis card with various statistics aggregated and visualized for easier analysis. We created these perturbation cards for all the experiments.

Hypotheses – After we established our experiment setup, we generated hypotheses around the results of the experiment on the basis of other work. Based on the preliminary experiments by Schlegel et al. [17], we generated the hypothesis that SHAP or SHAP derivatives will lead to the best results for the TSC task. Based on the results of Simic et al. [22], we will further look into the other attributions and double-check the results of Schlegel et al. [17] and the SHAP results even if SHAP results are less consistent [22]. Based on Simic et al. [22], we further look into the different perturbation strategies as we hypothesize that using one strategy is not enough to find a suitable attribution technique. Based on Geirhos et al. [6], we want also to check if there are patterns in the data the attributions show as relevant to find shortcuts the model learned to classify the data. E.g., using certain maximum or minimum values to separate one class from the other in a binary classification problem.

Perturbation Analysis Card – The perturbation analysis card is our proposed approach to reporting the results of our perturbation analysis strategies and techniques. Figure 2 shows such a perturbation analysis card with meta information (S), bar charts about the switch from one class to another (C), bar charts for the distribution of distances (D), statistics about the attributions (A), and a line plot visualization about the raw time series data (R).

Starting on top, Fig. 2(S), a short introduction presents a description of the dataset, the attribution technique, and the perturbation strategy. Right under the description, a stacked vertical bar chart shows a short glimpse of how good or bad the overall perturbation was. A good perturbation with an attribution technique presents just a lot of blue in this bar chart, while a bad perturbation shows a lot of orange in the visualization. Next to it, the exact numbers of the changed and unchanged samples are shown so that comparable numbers enhance the fast glance with other cards.

Figure 2(C) gives a detailed view of the perturbation and the changes there. The bar chart on the left visualizes the classes of the changed and unchanged predictions. For the changed prediction, the visualization also further presents the classes before and after the perturbation. Such visualization can help to identify spurious correlations as a model could, for instance, learn one feature of one class for the prediction. The bar chart on the right at (C) shows the number of perturbed values needed to change the prediction. The fewer changes needed, the better the attribution can identify relevant values.

In Fig. 2(D), the histogram of the distances between the perturbed and the original instances are shown. On top of (D), the Euclidean distances, and on the bottom of (D), the cosine distance can help to find clusters of needed changes for the perturbation of the samples by revealing a trend towards a certain distance range. Also, the distances can be used to compare the individual attribution techniques against each other. A smaller distance range, together with a lower number of perturbed values, presents a more focused technique.

Figure 2(A) visualizes more statistical information about the attributions. The plot on top of (A) shows the skewness of the attributions of the samples of the dataset. On the bottom, the means of the attributions are visualized. Through these, a general trend of the changed and unchanged samples and their attributions can be seen. Especially, outliers are interesting as a starting point for deeper analysis with other methods and visualizations.

Lastly, in Fig. 2(R), the time series time point means of the changed and unchanged samples can be inspected. So, for every time point in the time series, the mean of it over the subset (changed or unchanged) of the whole dataset is calculated and visualized. Thus, in the case of, e.g., FordA, with a standardization of the dataset, the samples slowly converge to zero. The visualization enables to spot large differences between the changed and unchanged samples.

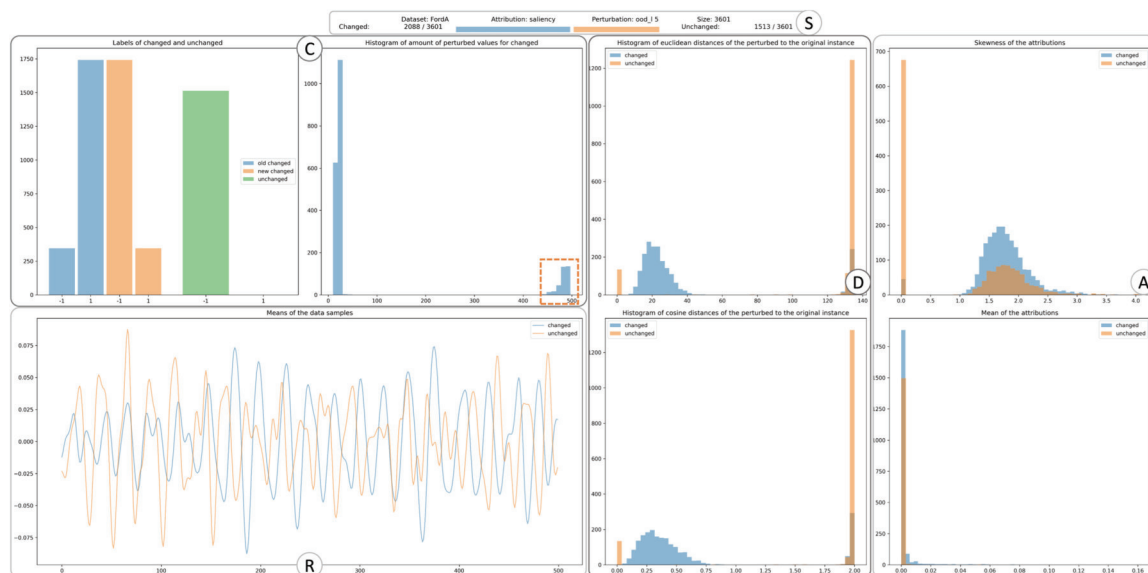


Fig. 2. Perturbation analysis card for the FordA dataset: the top description (S) contains general statistics for the card, starting with the dataset, the attribution technique, and the perturbation strategy. Beneath are the statistics for the amount of changed and unchanged sample predictions encoded as a progress bar and with numbers. The plots in (C) give a more detailed insight into the class changes after the perturbation. The left plot presents the amount of changed and unchanged samples for each class and also visualizes the class change for changed predictions. The right plot shows the number of perturbed values when a change in prediction happens. In (D), the distances of the original to the perturbed instance are shown. The top presents the Euclidean distance between the pairs, and the bottom shows the cosine distance. (A) presents the skew (top) and mean (bottom) of the attributions for the changed and unchanged sample predictions. In (R), the mean of every value at a specific time point is visualized for the changed and unchanged samples. (Color figure online)

5 Results and Discussion

Our current experiment setup evolves around an in-depth analysis of the attributions of seven attribution techniques (Saliency, IntegratedGradients, DeepLift, Occlusion, GradientShap, DeepLiftShap, KernelShap) based on the implementations in Captum¹. We incorporate 16 perturbation strategies, two based on Simic et al. [22], six based on Schlegel et al. [19], and eight extensions we describe later. We implemented nine single time point perturbations (zero, mean, inverse, dataset mean, dataset max, dataset min, OOD high, OOD low, random between min and max) and seven subsequence perturbations (zero, subsequence mean, dataset mean, inverse, OOD high, OOD low, random between min and max). The subsequence length is fixed to ten percent of the length of the data.

We focus on the UCR benchmark datasets [5] and take three of the most extensive datasets (FordA, FordB, ElectricDevices) to investigate data characteristics. However, our approach can be applied to any time series classification

¹ Captum is a Pytorch-based XAI module for Python: <https://captum.ai/>.

dataset. The FordA and FordB are sensor data with a length of 500 and provide an anomaly detection binary classification task. FordA has 3601 samples in the training set and 1320 in the test set. FordB has 3636 samples in the training set and 810 in the test set. The ElectricDevices dataset is shorter, with only 96 time points. However, the dataset has 8926 training samples and 7711 test samples.

We investigate two architectures of convolutional neural networks. The first architecture tackles the FordA and FordB datasets. The model consists of three 1D convolutional layers with a kernel size of three and increases the channels from one to 10 to 50 to 100. A max-pooling of three after the convolutional layer decreases the size again. A ReLu activation function activates the neuron. Afterward, a fully connected layer with 100 neurons and a ReLu function uses the feature maps from the convolutional layers to process the data further. And lastly, another fully connected layer with two neurons classifies the data with a softmax activation on top. We train the model with a batch size of 120 and the Adam optimizer [11]. The second architecture is trained on the ElectricDevices data and introduces a residual from the input to the fully connected layers. The original input gets downsampled using a 1D convolution with kernel size seven for the residual addition right before the fully connected layers.

We train our models using the cross-entropy loss for multi-label classification on the datasets for 500 epochs. Our models achieve for FordA an accuracy of 0.99 for the training set and 0.89 for the test set, demonstrating overfitting to the training data. Our models achieve for FordB an accuracy of 0.99 for the training set and 0.70 for the test set, demonstrating overfitting to the training data. Our models achieve for ElectricDevices an accuracy of 0.94 for the training set and 0.64 for the test set, demonstrating overfitting to the training data. As Ismail Fawaz et al. [9] showed, even our simple models are not too far from the state-of-the-art with other more sophisticated models. However, as we want to analyze our model, we look into the attributions of the training data, and thus our overfitting is a nice bonus to investigate spurious correlations and shortcuts [6].

Results – First, we start with the FordA dataset; next, we will present the FordB results, and lastly, the ElectricDevices dataset. FordA demonstrates interesting results regarding the attribution techniques and the perturbation strategies. The best working strategies are setting the perturbed value to an out-of-distribution low [26] on a subsequence [19] as you can see in Fig. 6. Especially, the saliency method [23] achieves the best result regarding the flip of predictions by flipping 2088 of 3601 samples, as also seen in Fig. 2. However, the KernelSHAP method [12] comes close to the flip with just 39 less with 2049 flips. Also, as seen in Fig. 2 on the plot on the right, the perturbation strategy out-of-distribution low changes the class quite late with a lot of perturbed values. Such an effect is unwanted in many cases as the model is, so to say, attacked by an adversarial attack outside of the distribution of the original data. In some cases, we can use such a method to test the model on data shifts, as, for example, the attributions can shift heavily. However, for our focus on the model itself, such an adversarial attack is interesting but does not show internal decision makings for the dataset we are interested in.

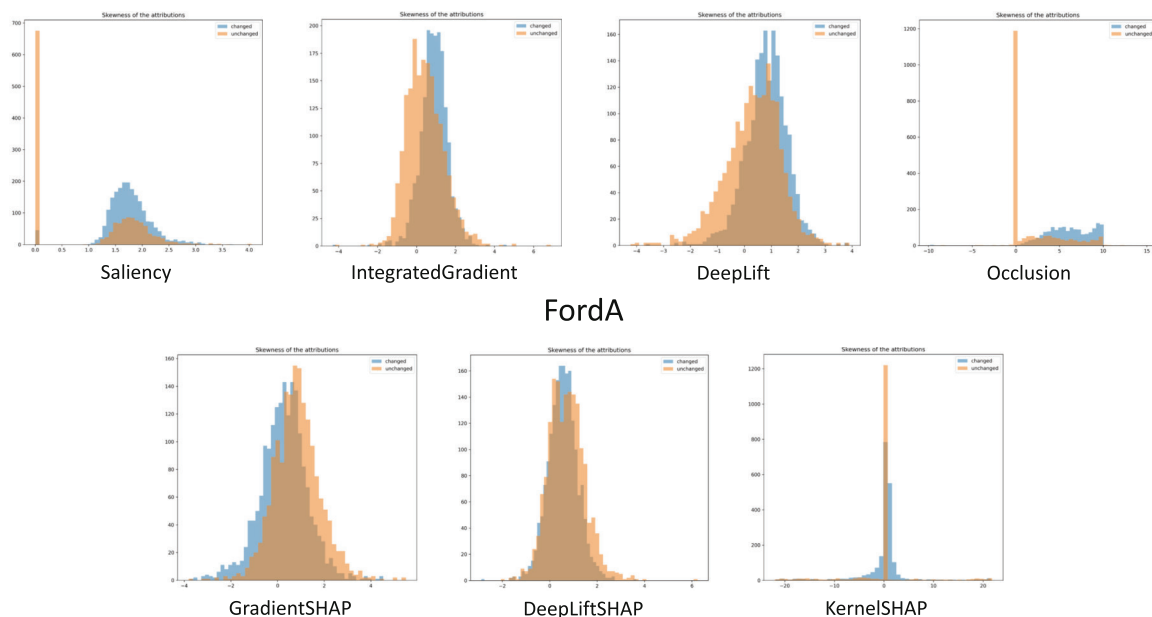


Fig. 3. Skewness distribution of the attribution techniques on the FordA dataset. Clear differences in the distributions of the attributions are visible. Further, the differences between changed and unchanged sample predictions and their attributions in the distributions are observable and show two different patterns in the techniques.

However, we also notice that the perturbation strategy heavily influences the best working method. If we switch, for example, to a perturbation to zero, we see Occlusion [28] as the winner in Fig. 6. Such a change in the best working technique demonstrates that the perturbation analysis just with one strategy is not enough to compare attribution techniques. We need multiple strategies to decide on one technique. However, we can also further take a deeper look into the attributions themselves. Focusing on the different skewness of the attributions and their distributions as seen in Fig. 3, we can already see some trends toward techniques enabling an easier inspection of the method and how well the method performs for the perturbation analysis. Especially, KernelSHAP in Fig. 3 demonstrates a nice pattern with two nearly not overlapping distributions. Such a nice distribution can help us to decide on one attribution technique.

The model for the FordB dataset is a bit worse than for the FordA dataset, which leads, in most cases, to worse performance in the perturbation analysis [22]. However, again the KernelSHAP and Saliency generate good working attributions for the change in the prediction for the perturbation to zero strategy. For this dataset, KernelSHAP achieves to change of 3472 from 3636 samples as seen in Fig. 7. Especially interesting is the distribution of the skewness of the attributions. A more in detail analysis of these two peaks could lead to further insights into the model and the attributions, but such an analysis needs another visualization, e.g., projecting the attributions in a scatter plot. However, if we further inspect our corresponding model card in Fig. 4, we can see that on the plot on the right, the change happens if a lot of values are removed from the original sample. Such a result is also observable in the other perturbation cards for

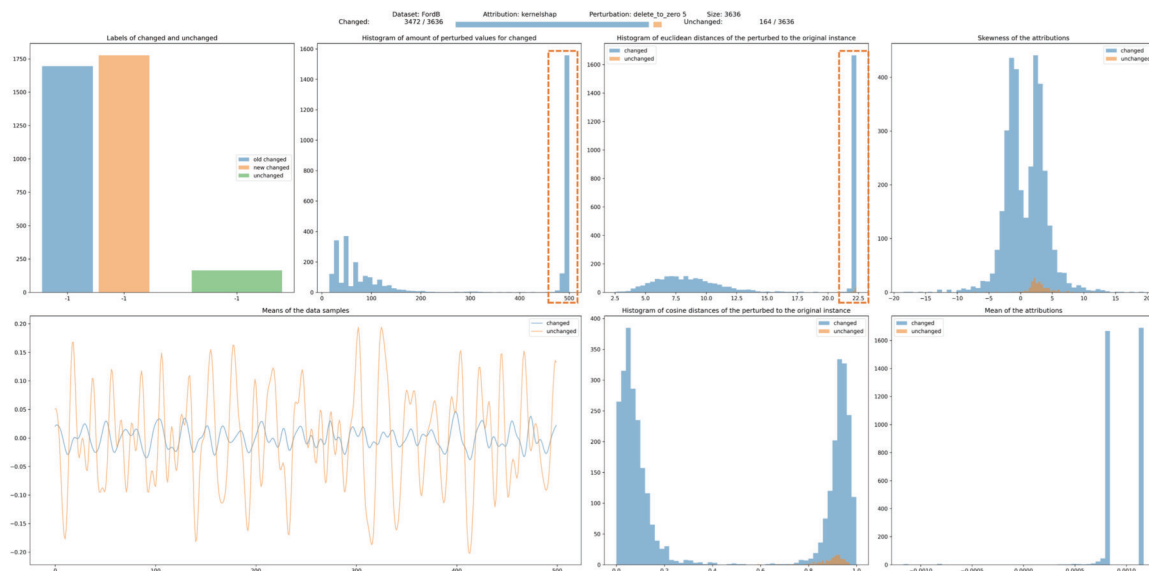


Fig. 4. Perturbation analysis card for the FordB dataset. The visualizations show a very distinct pattern. For the means of the raw samples (R), the few unchanged samples compose quite a diverse pattern, while the changed ones go to zero based on their standardization. However, the plot on the top (C) with the orange marker presents a pattern we do not want to have in a perturbation analysis, as it shows that we need to perturb a lot of data to achieve a change in prediction. (Color figure online)

the other techniques. In our study, we have identified a possible shortcut [6] that our model has learned from the training data. We speculate that the shortcut consists of a certain range or specific time points which need to be in a certain range of values to be classified as one class or the other class, and if we destroy this property, we change the prediction. So, our model learns a static version or range for one class and classifies everything else into the other class. Such a model does have more in common with an outlier detector than with a wanted classifier. Thus, we identified a shortcut of the model to be able to improve the classification without using all available features [6].

The ElectricDevices dataset is harder for the model as we do not only have a binary classification problem but seven classes the model needs to separate. However, as before, not even the state-of-the-art performance is as accurate as possible [9], which leads to worse attributions and a more diverse perturbation analysis result. Again, KernelSHAP performs best with a change of 8906 from 8926 samples with the values perturbed to the global max as seen in the perturbation card of Fig. 5. However, also IntegratedGradients [24] works well, but only with another perturbation strategy, namely changing the perturbed value to the global mean of the dataset. The dataset demonstrates quite nicely that the attribution techniques need different perturbation strategies to reveal the models' internal decision-making. Some of the techniques focus on different features the model learned as the ranking of the best-performing attribution techniques based on the perturbation analysis changes from strategy to strategy for this dataset. Additionally, when we delve into the labels of the changed and

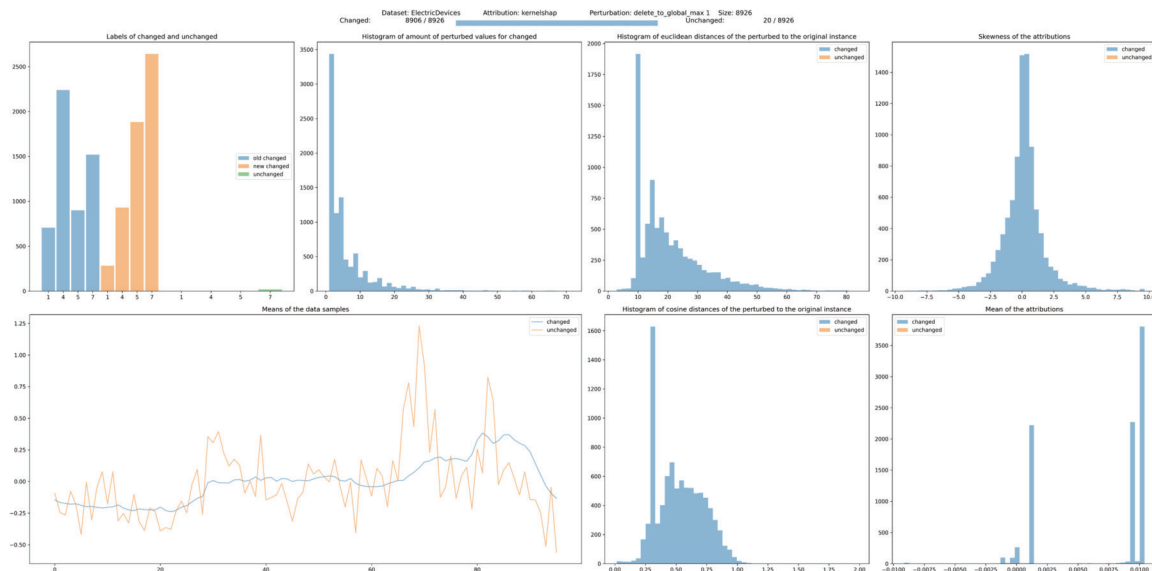


Fig. 5. Perturbation analysis card for the ElectricDevices dataset. The skewness distribution is quite interesting as it nearly presents a Gaussian distribution, with the mean being more sparse and quite focused on only three large pillars.

unchanged predictions, we notice that various attribution methods alter different labels in the perturbation. For example, KernelSHAP seems to modify every class besides seven, whereas Saliency influences classes other than five and six. However, unlike the previous FordB dataset, we do not see an unwanted perturbation pattern in the amount of perturbed values visualization. Such an effect presents that the attribution techniques are more suitable for the dataset and model than for the FordB model.

Summary – As we have seen in our results (Fig. 6, Fig. 7, Fig. 8), KernelSHAP performs quite well but takes a lot of time to calculate the attributions. Due to the sampling-based approach of KernelSHAP, the attributions are not deterministic and can vary from multiple computational runs. Further, in many cases, Saliency (or Vanilla Gradient multiplied by the Input) works surprisingly well and is only sometimes improved by additional extensions on top, such as IntegratedGradients. Thus, Saliency provides a promising variant for future experiments and techniques on top of it. So, if the attribution (explanation) is time-critical, Saliency is a well-suited method. If it is not time-critical, KernelSHAP provides the best-working attributions based on our experiments. Our collected data has even more insights and findings using the proposed perturbation analysis cards, which we look forward to analyzing and publishing with the code. The published source code can be used as a framework to experiment on more datasets, and the perturbation analysis cards can be used to report the results. The GitHub repository can be explored with more perturbation analysis cards and JSON data for the collected results of our experiments.

FordA / 3601	Saliency	IntegratedGradients	DeepLIFT	Occlusion	GradientSHAP	DeepLIFTSHAP	KernelSHAP
<i>Zero</i>	1744	1746	1743	<u>2029</u>	1746	1744	1743
<i>Inverse</i>	1744	1743	1743	1639	1743	1743	1743
<i>Local Mean</i>	1744	1746	1743	<u>2029</u>	1746	1744	1743
<i>Global Mean</i>	1744	1746	1743	<u>2028</u>	1746	1744	1743
<i>Global Max</i>	1743	1743	1743	1624	1743	1743	1743
<i>Global Min</i>	1743	1743	1743	1594	1744	1743	1743
<i>OOD Low</i>	1743	1743	1743	1636	1743	1744	1743
<i>OOD High</i>	1744	1743	1743	1606	1743	1743	1743
<i>Random</i>	1744	1744	1743	1641	1743	1744	1743
<i>Zero Sub</i>	1743	1789	1756	1807	1749	1754	1743
<i>Inverse Sub</i>	1823	1745	1748	1628	1749	1752	1852
<i>Local Mean Sub</i>	1743	1789	1756	1807	1749	1754	1743
<i>Global Mean Sub</i>	1743	1789	1756	1810	1748	1754	1743
<i>OOD Low Sub</i>	<u>2088</u>	1768	1804	1639	1876	1870	<u>2049</u>
<i>OOD High Sub</i>	1997	1760	1776	1445	1835	1815	1985
<i>Random Sub</i>	1796	1787	1774	1625	1780	1793	1771

Fig. 6. Changed samples from the perturbation analysis for the FordA dataset. The higher, the better. Saliency and KernelSHAP as winners, with Occlusion behind.

FordB / 3636	Saliency	IntegratedGradients	DeepLIFT	Occlusion	GradientSHAP	DeepLIFTSHAP	KernelSHAP
<i>Zero</i>	1777	1777	1777	1614	1777	1777	1777
<i>Inverse</i>	1777	1777	1777	1445	1777	1777	1777
<i>Local Mean</i>	1777	1777	1777	1614	1777	1777	1777
<i>Global Mean</i>	1778	1777	1777	1614	1777	1777	1777
<i>Global Max</i>	1777	1777	1777	1444	1777	1777	1777
<i>Global Min</i>	1777	1777	1777	1416	1777	1777	1777
<i>OOD Low</i>	1777	1777	1777	1531	1777	1777	1777
<i>OOD High</i>	1777	1741	1729	1066	1776	1759	1777
<i>Random</i>	1777	1777	1777	1411	1777	1777	1777
<i>Zero Sub</i>	<u>3425</u>	2903	2558	1557	2535	2328	<u>3472</u>
<i>Inverse Sub</i>	1851	1802	1796	1424	1804	1790	1855
<i>Local Mean Sub</i>	<u>3425</u>	2903	2558	1557	2535	2328	<u>3472</u>
<i>Global Mean Sub</i>	3387	2846	2505	1554	2478	2285	<u>3440</u>
<i>OOD Low Sub</i>	2132	2047	2002	1548	2046	1984	2136
<i>OOD High Sub</i>	1904	1039	1009	461	1719	1306	1919
<i>Random Sub</i>	1815	1810	1801	1421	1809	1793	1822

Fig. 7. Changed samples from the perturbation analysis for the FordB dataset. The higher, the better. KernelSHAP is the winner, and Saliency is behind.

ElectricDevices / 8926	Saliency	IntegratedGradients	DeepLIFT	Occlusion	GradientSHAP	DeepLIFTSHAP	KernelSHAP
<i>Zero</i>	7227	8566	8782	5249	8100	7959	7944
<i>Inverse</i>	7263	8632	8596	5281	8461	8486	8159
<i>Local Mean</i>	7227	8566	8782	5249	8100	7959	7944
<i>Global Mean</i>	7823	8495	8734	5185	8228	8118	8219
<i>Global Max</i>	8627	<u>8844</u>	<u>8884</u>	5606	<u>8861</u>	8591	<u>8906</u>
<i>Global Min</i>	6636	6968	7129	5337	6791	6675	6719
<i>OOD Low</i>	8261	7926	7955	5581	8096	7791	7893
<i>OOD High</i>	6563	5669	6067	4070	6152	5815	6194
<i>Random</i>	7531	7930	8032	5313	7789	7672	7728
<i>Zero Sub</i>	8502	8786	8829	5162	8628	8799	8433
<i>Inverse Sub</i>	6743	7067	6781	4907	6971	6892	6944
<i>Local Mean Sub</i>	8502	8786	8829	5162	8628	8799	8433
<i>Global Mean Sub</i>	8451	<u>8874</u>	<u>8865</u>	5182	8749	<u>8831</u>	8687
<i>OOD Low Sub</i>	7742	7283	7089	5317	7455	7383	7367
<i>OOD High Sub</i>	6034	5416	5553	3760	5718	5718	5943
<i>Random Sub</i>	7601	7751	7707	5254	7786	7608	7826

Fig. 8. Changed samples from the perturbation analysis for the FordB dataset. The higher, the better. KernelSHAP is the winner, and DeepLIFT and IntegratedGradients are behind on different perturbation strategies.

6 Conclusion and Future Work

After reviewing related work, we presented an in-depth analysis of perturbation strategies for attributions on time series. With the analysis, we dug into a CNN trained on time series classification to investigate attributions, perturbation strategies, and shortcuts the network learned. We presented our results in perturbation analysis cards to enable users to analyze the results in detail by inspecting the aggregated data in visualizations and comparing them easily with other techniques based on the provided cards. We identified SHAP as a suitable method to generate working attributions in all experimented datasets. Other gradient-based methods also work quite well but do not perform as well as, e.g., KernelSHAP. However, depending on the perturbation strategy, the best working attribution technique changes quite drastically also for some techniques. We advise not only focusing on a single strategy but to using multiple strategies and aggregating the results of these, and looking at the distribution of the skewness to enhance the comparability. In our experiments, we also found a shortcut or spurious correlation for the FordB dataset, which our model learned to classify one class and to classify everything else as the other class.

Future Work – We want to extend the experiment to other attribution techniques and compare the results with the already collected experiment results. Also, we want to compare the attributions even in more detail by, e.g., aggregating the attributions and comparing them on a higher level to find matching patterns. Different trends and subsequences are further patterns to analyze to gain knowledge into the attribution techniques. With such an approach, we also

want to include the *local Lipschitz estimate* [3] to rank consistent attributions higher. Last, we want to extend the *Perturbation Effect Size* [22] and use our gained knowledge to combine perturbation strategies, switching predictions, and distances to generate a measure to evaluate attributions on time series classification models more robust and fully automatically to make it easier for users to decide which attributions to use for explanations. We also want to enhance our perturbation analysis cards further to be more easily readable and comfortable for non-experts to be able to gain insights at a single glance.

Acknowledgment. This work has been partially supported by the Federal Ministry of Education and Research (BMBF) in VIKING (13N16242).

References

1. Adadi, A., Berrada, M.: Peeking inside the black-box: a survey on explainable artificial intelligence (XAI). *IEEE Access* **6**, 52138–52160 (2018)
2. Aghabozorgi, S., Shirkhorshidi, A.S., Wah, T.Y.: Time-series clustering—a decade review. *Inf. Syst.* **53**, 16–38 (2015)
3. Alvarez-Melis, D., Jaakkola, T.S.: On the robustness of interpretability methods. arXiv preprint [arXiv:1806.08049](https://arxiv.org/abs/1806.08049) (2018)
4. Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W.: On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE* **10**(7), e0130140 (2015)
5. Dau, H.A., et al.: The UCR time series archive. *IEEE/CAA J. Automatica Sinica* **6**(6), 1293–1305 (2019)
6. Geirhos, R., et al.: Shortcut learning in deep neural networks. *Nat. Mach. Intell.* **2**(11), 665–673 (2020)
7. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. *ACM Comput. Surv. (CSUR)* **51**(5), 1–42 (2018)
8. Hooker, S., Erhan, D., Kindermans, P.J., Kim, B.: A benchmark for interpretability methods in deep neural networks. In: Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc. (2019)
9. Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Deep learning for time series classification: a review. *Data Min. Knowl. Disc.* **33**(4), 917–963 (2019). <https://doi.org/10.1007/s10618-019-00619-1>
10. Jeyakumar, J.V., Noor, J., Cheng, Y.H., Garcia, L., Srivastava, M.: How can i explain this to you? An empirical study of deep neural network explanation methods. In: *Advances in Neural Information Processing Systems*, vol. 33 (2020)
11. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
12. Lundberg, S., Lee, S.I.: A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems* (2017)
13. Mercier, D., Bhatt, J., Dengel, A., Ahmed, S.: Time to focus: a comprehensive benchmark using time series attribution methods. arXiv preprint [arXiv:2202.03759](https://arxiv.org/abs/2202.03759) (2022)

14. Mohseni, S., Zarei, N., Ragan, E.D.: A multidisciplinary survey and framework for design and evaluation of explainable AI systems. *ACM Trans. Interact. Intell. Syst. (TIIS)* **11**(3–4), 1–45 (2021)
15. Ribeiro, M.T., Singh, S., Guestrin, C.: “Why should i trust you?” Explaining the predictions of any classifier. In: *International Conference on Knowledge Discovery and Data Mining* (2016)
16. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **1**(5), 206–215 (2019)
17. Schlegel, U., Arnout, H., El-Assady, M., Oelke, D., Keim, D.A.: Towards a rigorous evaluation of XAI methods on time series. In: *ICCV Workshop on Interpreting and Explaining Visual Artificial Intelligence Models* (2019)
18. Schlegel, U., Keim, D.A.: Time series model attribution visualizations as explanations. In: *TREX: Workshop on TRust and EXpertise in Visual Analytics* (2021)
19. Schlegel, U., Oelke, D., Keim, D.A., El-Assady, M.: An empirical study of explainable AI techniques on deep learning models for time series tasks. In: *Pre-Registration Workshop NeurIPS* (2020)
20. Schlegel, U., Vo, D.L., Keim, D.A., Seebacher, D.: TS-MULE: local interpretable model-agnostic explanations for time series forecast models. In: Kamp, M., et al. (eds.) *ECML PKDD 2021. CCIS*, vol. 1524, pp. 5–14. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93736-2_1
21. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: visual explanations from deep networks via gradient-based localization. In: *International Conference on Computer Vision* (2017)
22. Šimić, I., Sabol, V., Veas, E.: Perturbation effect: a metric to counter misleading validation of feature attribution. In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pp. 1798–1807 (2022)
23. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: visualising image classification models and saliency maps. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2014)
24. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: *International Conference on Machine Learning. JMLR.org* (2017)
25. Theissler, A., Spinnato, F., Schlegel, U., Guidotti, R.: Explainable AI for time series classification: a review, taxonomy and research directions. *IEEE Access* **10**, 100700–100724 (2022)
26. Turbé, H., Bjelogrić, M., Lovis, C., Mengaldo, G.: InterpretTime: a new approach for the systematic evaluation of neural-network interpretability in time series classification. *arXiv preprint arXiv:2202.05656* (2022)
27. Yeh, C.K., Hsieh, C.Y., Suggala, A., Inouye, D.I., Ravikumar, P.K.: On the (in) fidelity and sensitivity of explanations. In: *Advances in Neural Information Processing Systems*, vol. 32 (2019)
28. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014. LNCS*, vol. 8689, pp. 818–833. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10590-1_53
29. Zhou, Y., Booth, S., Ribeiro, M.T., Shah, J.: Do feature attribution methods correctly attribute features? In: *AAAI Conference on Artificial Intelligence* (2022)