




Motif-Guided Time Series Counterfactual Explanations

Peiyu Li¹, Soukaïna Filali Boubrahimi¹, and Shah Muhammad Hamdi²

¹ Utah State University, Logan, UT 84322, USA
peiyu.li@usu.edu

² New Mexico State University, Las Cruces, NM 88003, USA

Abstract. With the rising need of interpretable machine learning methods, there is a necessity for a rise in human effort to provide diverse explanations of the influencing factors of the model decisions. To improve the trust and transparency of AI-based systems, the EXplainable Artificial Intelligence (XAI) field has emerged. The XAI paradigm is bifurcated into two main categories: feature attribution and counterfactual explanation methods. While feature attribution methods are based on explaining the reason behind a model decision, counterfactual explanation methods discover the smallest input changes that will result in a different decision. In this paper, we aim at building trust and transparency in time series models by using motifs to generate counterfactual explanations. We propose Motif-Guided Counterfactual Explanation (MG-CF), a novel model that generates intuitive post-hoc counterfactual explanations that make full use of important motifs to provide interpretive information in decision-making processes. To the best of our knowledge, this is the first effort that leverages motifs to guide the counterfactual explanation generation. We validated our model using five real-world time-series datasets from the UCR repository. Our experimental results show the superiority of MG-CF in balancing all the desirable counterfactual explanations properties in comparison with other competing state-of-the-art baselines.

Keywords: Counterfactual explanations · Explainable Artificial Intelligence (XAI) · time series motifs

1 Introduction

Recently, time-series data mining has played an important role in various real-life applications, such as healthcare [10], astronomy [3], and aerospace [11]. The availability of big data provides researchers the unprecedented opportunity to deploy high accurate models in real-life applications. The success of deep learning time series models has been validated by their high accuracy. However, the black-box nature makes the decision-making process of high accuracy models less explained [1]. This represents a big barrier for decision makers in trusting the system. For example, in life-changing decisions such as critical disease management, it is important to know the reasons behind the critical decision of choosing

a treatment line over another one [9]. An inappropriate treatment decision can cost human lives in addition to a substantial monetary loss [26]. Therefore, it is vital to understand the opaque models’ decision process by either prioritizing interpretability during the model development phase, or developing post-hoc explanation solutions.

To build trust between humans and decision-making systems, EXplainable Artificial Intelligence (XAI) methods are increasingly accepted as effective tools to explain machine learning models’ decision-making processes. XAI methods aim at increasing the interpretability of the models whilst maintaining the high-performance levels. There are two XAI dominant paradigms: intrinsic interpretable and post-hoc explanations for opaque models [29]. Examples of intrinsic interpretable models include linear regression, decision trees, rule sets, etc. On the other hand, examples of post-hoc explanation methods include LIME [19], native guide, and SHAP [23,24]. In this paper, we focus only on post-hoc XAI methods. A lot of efforts have been made to provide post-hoc XAI for image and vector-represented data while significantly less attention has been paid to time series data [18]. The complex and time-evolving nature of time series makes the explanation models one of the most challenging tasks [5].

On the light of desirable counterfactual method properties, we propose a Motif-Guided Counterfactual Explanation (MG-CF), a novel model that generates interpretable, intuitive post-hoc counterfactual explanations of time series classification predictions that balance between validity, proximity, interpretability, contiguity and efficiency. The contributions are the following:

1. We incorporate motif mining algorithms for guiding the counterfactual search. In particular, we make full use of the interpretability power of motifs to provide interpretable information in the decision-making process. Our method does not require the use of class activation maps to search for the counterfactual explanation, which makes it model-agnostic.
2. We conduct experiments on six real-life datasets from various domains (image, spectro, ECG, and motion) and show the superiority of our methods over state-of-the-art models.

To the best of our knowledge, this is the first effort to leverage the a-priori mined motifs to produce counterfactual explanations for time series classification. The rest of this paper is organized as follows: in Sect. 2, background and related works are described. Section 3 introduces the preliminary concepts. Section 4 describes our proposed method in detail. We present the experimental results and evaluations in comparison to other baselines in Sect. 5. Finally, we conclude our work in Sect. 6.

2 Background and Related Work

In the post-hoc interpretability paradigm, one of the most prominent approaches is to identify important features given a prediction through local approximation, such as LIME [19], LORE [7], TS-MULE [21], and SHAP [14]. In particular,

LIME is a feature-based approach that operates by fitting an interpretable linear model to the classifier’s predictions of random data samples, weighted based on their distance to the test sample. LORE is an extension work based on LIME, which is a local black box model-agnostic explanation approach based on logic rules. TS-MULE is also an extension to LIME with novel time series segmentation techniques and replacement methods to enforce a better non-informed values exchange. SHAP is a unified framework that operates by calculating feature importance values using model parameters. In addition, visualizing the decision of a model is also a common technique for explaining model predictions [16]. In the computer vision domain, visualization techniques have been widely applied to different applications successfully, such as highlighting the most important parts of images to class activation maps (CAM) in convolutional neural networks [15]. Schlegel et al. [20] tested the informativeness and robustness of different feature-importance techniques in time series classification. LIME was found to produce poor results because of the large dimensionality by converting time to features; in contrast, saliency-based approaches and SHAP were found to be more robust across different architectures.

An alternative method to feature-based methods has been proposed by Wachter et al. [25], which aims at minimizing a loss function and using adaptive Nelder-Mead optimization to encourage the counterfactual to change the decision outcome and keep the minimum Manhattan distance from the original input instance. Similarly, another method that is used to deal with the plausibility and feasibility issues of the generated counterfactual explanation has been proposed, which is called GeCo. The model achieves the desirable counterfactual properties by introducing a new plausibility-feasibility language (PLAF) [22]. Both GeCo and wCF focus on structured tabular datasets. Since these two methods explore a complete search space, they are not adequate to be used in a high dimensional feature space such as in the case of time-series data.

Recently, an instance-based counterfactual explanation for time series classification has been proposed [5]. The instance-based counterfactual explanation uses the explanation weight vector (from the Class Activation Mapping) and the in-sample counterfactual (NUN) to generate counterfactual explanations for time series classifiers. The proposed technique adapts existing training instances by copying segments from other instances that belong to the desired class. More recently, MAPIC, an interpretable model for TSC based on Matrix Profile, shapelets, and decision tree has been proposed by [6]. Finally, a counterfactual solution for multivariate time series data called CoMTE has been proposed by Etes et al. [2], which focuses on observing the effect of turning off one variable at a time.

3 Preliminary Concepts

Notation. We define a time series $T = \{t_1, t_2, \dots, t_m\}$ as an ordered set of real values, where m is the length, then we can define a time series dataset $\mathbf{T} = \{T_0, T_1, \dots, T_n\}$ as a collection of such time series where each time series has

mapped to a mutually exclusive set of classes $C = \{c_1, c_2, \dots, c_l\}$ and used to train a time series classifier $f : \mathbf{T} \rightarrow C$, where $\mathbf{T} \in \mathbb{R}^k$ is the k -dimensional feature space. For each T_i in time series dataset \mathbf{T} associated with a class $f(T_i) = c_i$, a counterfactual explanation model \mathcal{M} generates a perturbed sample T'_i with the minimal perturbation that leads to $f(T'_i) = c'_i$ such that $c_i \neq c'_i$.

According to [5, 13], a desire counterfactual instance x_{cf} should obey the following initial properties :

- **Validity:** The prediction of the to-be-explained model f on the counterfactual instance x_{cf} needs to be different from the prediction of the to-be-explained model f on the original instance x (i.e., if $f(x) = c_i$ and $f(x_{cf}) = c_j$, then $c_i \neq c_j$).
- **Proximity:** The to-be-explained query needs to be close to the generated counterfactual instance, which means the distance between x and x_{cf} should be minimal.
- **Sparsity:** The perturbation δ changing the original instance x_0 into $x_{cf} = x_0 + \delta$ should be sparse (i.e., $\delta \approx \epsilon$).
- **Model-agnosticism:** The counterfactual explanation model should produce a solution independent of the classification model f , high-quality counterfactuals without prior knowledge of the gradient values derived from optimization-based classification models should be generated.
- **Contiguity:** The counterfactual instance $x_{cf} = x + \delta$ needs to be perturbed in a single contiguous segment which makes the solution semantically meaningful.
- **Efficiency:** The counterfactual instance x_{cf} needs to be found fast enough to ensure it can be used in a real-life setting.

4 Motif-Guided Counterfactual Explanation (MG-CF)

In this section, we describe our proposed Motif-Guided Counterfactual Explanation method. The general architecture is illustrated in Fig. 1, there are two main steps: motifs mining and counterfactual explanation generation.

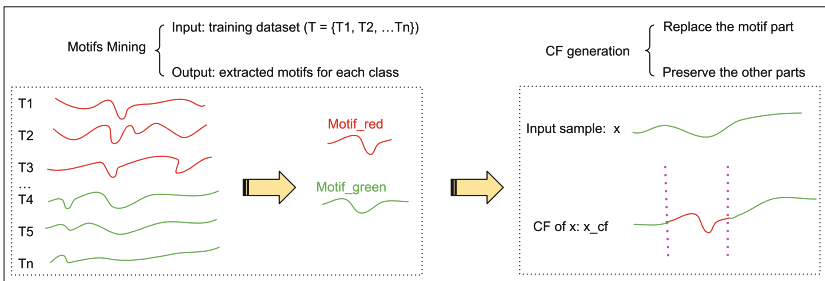


Fig. 1. Motif-Guided Counterfactual Explanation

4.1 Motif Mining

Motifs, also known as shapelets, are time series sub-sequences that are maximally representative of a class [27]. Motifs provide interpretable information beneficial for domain experts to better understand their data. To mine the most prominent motifs, we apply the Shapelet Transform (ST) algorithm proposed by [12]. To generate the most prominent motifs, a set of candidate motifs need to be generated, a distance measure between a motif and each time series instance should be defined, and a measure of the discriminatory power of the mined motifs [12] need to be defined. The details of the algorithm is shown in Algorithm 1. The algorithm starts by sampling specific-length intervals from the time series instance to generate candidate motifs (line 7). The motifs' discriminatory power is then estimated by first computing their distance to all the instances from each class (line 9), and then computing the distance correlation to the target class (line 10). Finally, the motifs are sorted and only the best motif from each class is retained (lines 14, 16). Figure 2a and 2b show the two extracted motifs of different classes from ECG 200 dataset.

Candidates Generation. Given a time series T with length m , a subsequence S of length l of time series T is a contiguous sequence of l points in T , where $l \leq m$. For a given m -lengthed time series, there are $(m-l)+1$ distinct subsequences of length l . The set of all subsequences of length l for a dataset is described as:

$$W_l = \{W_{1,l}, \dots, W_{n,l}\}, \quad (1)$$

where $W_{i,l}$ is defined as the set of all subsequences of length l for time series T_i . To speed up the process of finding motifs, instead of considering all the lengths that range in $(3, m)$, similar to [12], we only consider 0.3, 0.5 and 0.7 percent length of the whole time series for each dataset. In this case, the set of all candidate motifs for data set T is:

$$W = \{W_{0.3m}, W_{0.5m}, W_{0.7m}\}, \quad (2)$$

We denote the candidates generation as *generateCandidates* in Algorithm 1.

Measuring Distances. The distance between two subsequences S and R of length l is defined as

$$dist(S, R) = \sum_{i=1}^l (s_i - r_i)^2 \quad (3)$$

To calculate all distances between a candidate motif S and all series in T , we denote $d_{i,s}$ as the distance between a subsequence S of l and time series T_i , then

$$d_{i,s} = \min_{R \in W_{i,l}} dist(S, R) \quad (4)$$

$$D_S = \langle d_{1,S}, d_{2,S}, \dots, d_{n,S} \rangle \quad (5)$$

We denote the process of measuring distance as *findDistances* in Algorithm 1.

Measuring the Discriminatory Power of a Motif. Information gain has been used to evaluate the quality of a motif in the previous shapelet papers [17, 27, 28]. Firstly, the distance list D_S , defined in Eq. 5, is sorted, then the information gain is evaluated on the class values for each possible split value. Ye et al. propose to use an early abandon technique on calculating D_S through maintaining an upper bound on the quality of the candidate whilst generating D_S [28]. If the upper bound falls below the best found so far, the calculation of D_S can be abandoned. After each $d_{i,S}$ is found, the current best information gain split is calculated and the upper bound is found, assuming the most optimistic division of the remaining distances. We denote the process of measuring the discriminatory power of a motif as *assessCandidate* in Algorithm 1.

Algorithm 1. Motifs mining

Input: Training set samples \mathbf{T} with labeled binary classes $C = [0, 1]$

Output: Extracted motifs for each class

```

1: Motifs =  $\emptyset$ 
2:  $N = \text{length}(\mathbf{T}[0])$  ▷ Number of time series samples
3:  $m = \text{length}(\mathbf{T}[0][0])$  ▷ Length of time series
4: for  $\mathbf{T}_i \leftarrow \mathbf{T}_1$  to  $\mathbf{T}_N$  do
5:   Motifs  $\leftarrow \emptyset$ 
6:   for  $l$  in  $[0.3m, 0.5m, 0.7m]$  do
7:      $W_{i,l} \leftarrow \text{generateCandidates}(\mathbf{T}_i, l)$ 
8:     for all subsequences  $S$  in  $W_{i,l}$  do
9:        $D_S \leftarrow \text{findDistances}(S, W_{i,l})$ 
10:      quality  $\leftarrow \text{assessCandidate}(S, D_S)$ 
11:      Motifs.add( $i, \text{start\_idx}, \text{end\_idx}, S, \text{quality}$ ) ▷ The index of time series,
the start idx and end idx of motifs will be stored
12:    end for
13:  end for
14:  sortByQuality(Motifs)
15: end for
16: return  $[[i_0, \text{start\_idx}_0, \text{start\_idx}_0], [i_1, \text{start\_idx}_1, \text{start\_idx}_1]]$  ▷ return the
index information for motifs of different classes

```

4.2 CF Generation

After motifs are extracted from the training dataset, the next step is to generate CF explanations. In this section, we introduce the counterfactual explanation generation process as outlined at Algorithm 2. The CF generation starts by using the pre-trained model f to predict the classification result of the test dataset *samples* (line 1). We set the target classes of the test dataset as the opposite label of the prediction results (line 2). To generate an explanation of a query time series X , we consider the motif’s existing segment from the target class as the most prominent segment (lines 4–7). Therefore, we replace the motif

existing part with the target class motif and preserve the remaining parts (line 9). In this case, the counterfactual explanation for the input query time series is generated (line 10). It is important to note that MG-CF does not require the use of class activation maps which makes it a model-agnostic model. Figure 2 shows examples of generated counterfactual explanations for the ECG 200 dataset of the Myocardial Infarction and Normal Heartbeat class.

Algorithm 2. CF generation

Input: Test set samples $samples$, pretrained time series classifier f , extracted motifs for each class from Motif mining algorithm

Output: **CF**, Counterfactual explanation for each sample in test set

```

1: preds = f.predict(samples)    ▷ Applying the test data on f to get the prediction
  results
2: targets= the opposite label of the preds    ▷ Get the target labels based on the
  prediction results
3: for sample ← samples do
4:   if target(sample) == 0 then
5:     i, start_idx, end_idx = [i_0, start_idx_0, end_idx_0]    ▷ extracted index
  information of the motif from class 0
6:   else
7:     i, start_idx, end_idx = [i_1, start_idx_1, end_idx_1]    ▷ extracted index
  information of the motif from class 1
8:   end if
9:   sample[start_idx, end_idx] = T[i][start_idx, end_idx]
10:  cf_sample = sample
11:  CF.append(cf_sample)
12: end for
13: return CF

```

5 Experimental Evaluation

In this section, we outline our experimental design and discuss our findings. We conduct our experiments on the publicly-available univariate time series data sets from the University of California at Riverside (UCR) Time Series Classification Archive [4]. We test our model on five real-life datasets from various domains (image, spectro, ECG, and motion). The number of labels of each dataset is binary and the time series classes are evenly sampled across the training and testing partitions. The length of the time series in the data sets varies between 96 and 512. Table 1 shows all the dataset details (number of classes, time series length, training size, testing size, and the type of time series).

5.1 Baseline Methods

We evaluate our proposed SG-CF model with two other baselines, namely, Alibi [8] and Native guide counterfactual explainers [5].

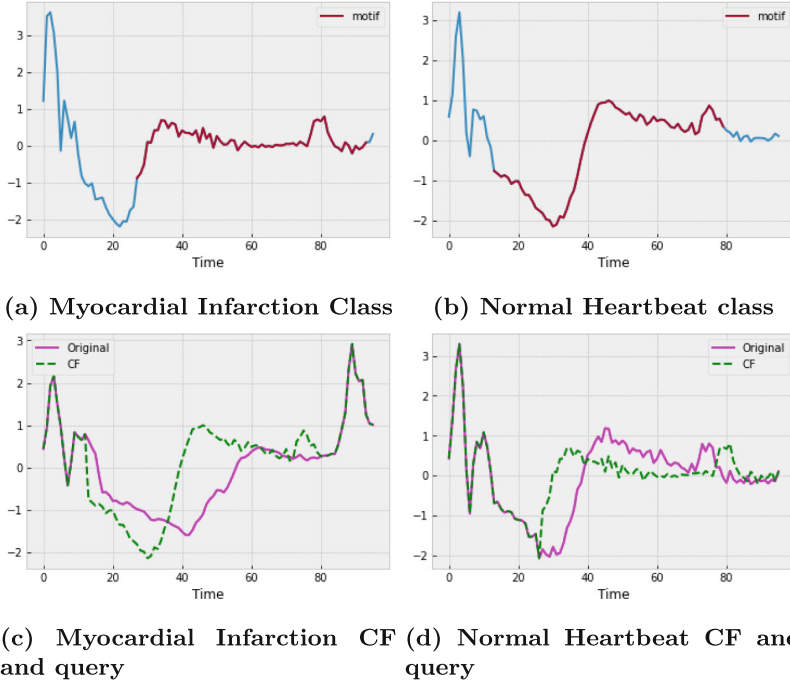


Fig. 2. Example MG-CF explanations for the ECG200 dataset

- **Alibi Counterfactual (Alibi):** The Alibi generates counterfactuals by optimizing a new objective function that includes class prototypes generated from encoders [8]. The method leverages the resulting autoencoder reconstruction error after decoding the counterfactual instance x_{cf} as a measure of fitness to the training data distribution. The additional loss term L_{AU} penalizes out-of-distribution counterfactual instances.
- **Native guide counterfactual (NG-CF):** NG-CF is the latest time series counterfactual method that perturbs the query time series by inserting segments from the opposing class [5]. Although the method is model agnostic, the method reaches its full potential when class activation weights are available. NG-CF uses Dynamic Barycenter (DBA) averaging of the query time series x and the nearest unlike neighbor from another class to generate the counterfactual example.

The source code of our model, competing baselines, and the experimental dataset are available on the MG-CF project website¹.

¹ <https://sites.google.com/view/naive-shapelets-guided-cf>.

Table 1. UCR datasets metadata (each dataset is binary labeled)

ID	Dataset Name	TS length	DS train size	DS test size	Type
0	ECG200	96	100	100	ECG
1	Coffee	286	28	28	SPECTRO
2	GunPoint	150	50	150	MOTION
3	BeetleFly	470	20	20	IMAGE
4	BirdChicken	512	20	20	IMAGE

5.2 Experimental Results

The goal of our experiments is to assess the performance of the baseline methods with respect to all the desired properties of an ideal counterfactual method. To evaluate our proposed method, we compare our method with the other two baselines in terms of several metrics: L1 distance (proximity), sparsity level, number of independent segments is perturbing (contiguity), the label flip rate (validity), and the run time (efficiency).

L1 distance, which is defined in Eq. 6, measures the distance between the counterfactuals and the original samples, a smaller L1 distance is desired.

$$proximity = \|x_{cf} - x\| \quad (6)$$

Sparsity level, which indicates the level of time series perturbations. A high sparsity level that is approaching 100% is desirable, which means the time series perturbations made in x to achieve x_{cf} is minimal. We computed the sparsity level using a new metric that we formulate in Eqs. 7 and 8.

$$sparsity = 1 - \frac{\sum_{i=0}^{len(x)} g(x_{cf_i}, x_i)}{len(x)} \quad (7)$$

$$g(x, y) = \begin{cases} 1, & \text{if } x \neq y \\ 0, & \text{otherwise} \end{cases} \quad (8)$$

The number of independent non-contiguous segments is also investigated to show the contiguity, which is shown in Fig 3. The lower the number of independent non-contiguous segments the better. Finally, we define the validity metric by comparing the label flip rate for the prediction of the counterfactual explanation result, we computed the flip rate following the formula of Eq. 9.

$$flip_rate = \frac{num_flipped}{num_testsample}, \quad (9)$$

where we denote the `num_flipped` as the number of generated counterfactual explanation samples that flip the label, and the `num_testsample` is the total number of samples in the test dataset that we used to generate counterfactual explanation samples. The closer the label flip rate is to 1, the better.

Figure 3 and 4 show the evaluation results on the CF desired properties assessed using the aforementioned metrics. Since our proposed method relies on an instance-based counterfactual explanation, which generates a CF for each time series in the dataset, we calculate the mean value and the standard deviation for all CF samples for each dataset to show the overall performances. We note that Alibi achieves the lowest L1 distance (proximity) in the four datasets, which is highly desirable. However, Alibi minimizes proximity in the cost of sparsity level, contiguity (number of independent segments), and efficiency (running time). Alibi achieves the lowest sparsity level and run time, which entails that the method generates CF explanations with less proximity, contiguity, and efficiency.

We also note that NG-CF results in highly valid counterfactual explanations that are guaranteed to have the desired class. Since the method does not rely on any optimization function to generate new shapes, it simply copies fragments from existing training time series instances that the prediction model f has already learned. Therefore, it is expected that the prediction model f recognizes the copied segments with high confidence. Although highly valid, NG-CF sacrifices the other metrics (sparsity, contiguity) for maximizing validity. Figures 3 and 4 show that our proposed MG-CF results in a higher sparsity level,

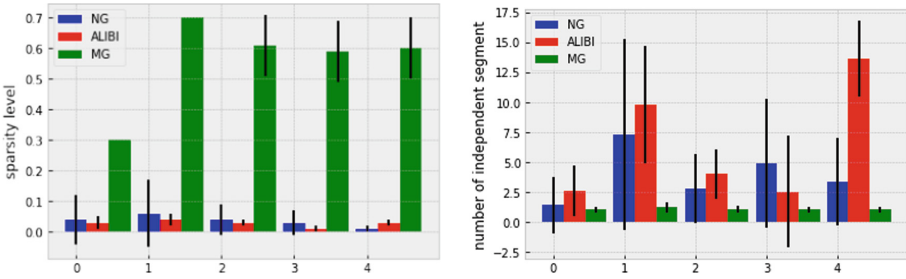


Fig. 3. The sparsity level (the higher the better) and number of independent segments (the lower the better) of the CF explanations

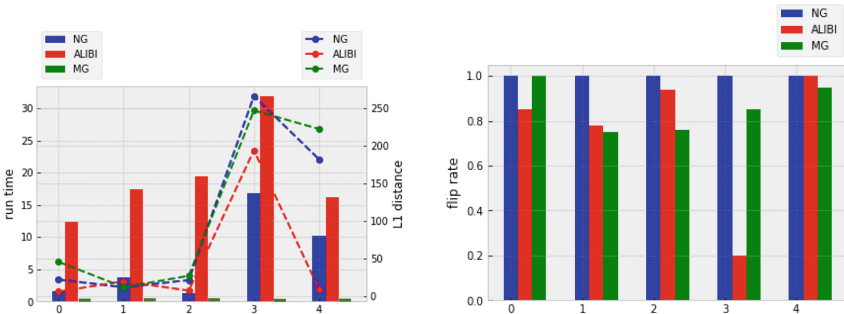


Fig. 4. The average running time (the lower the better), average L1 distance (the lower the better), and the label flip rate (the higher the better) of the CF explanations

a lower number of independent perturbed segments (contiguity), and a lower running time (efficiency) compared to NG-CF and Alibi. Finally, although Alibi and NG-CF can lead to high proximity and validity counterfactuals, they are not valid and sparse respectively. Thus, MG-CF shows a good balance across all the metrics without maximizing one property and compromising the other.

6 Conclusion

In this paper, we propose a novel model that derives intuitive, interpretable post-hoc counterfactual explanations of time series classification models that finds balance between validity, proximity, sparsity, contiguity, and efficiency desirable properties. The main limitation of current time series explainers is that they produce perturbations that maximize a metric on the cost of others. This shortcoming limits the interpretability of their resulting counterfactual explanations. We address these challenges by proposing MG-CF, a motif-based model that produces high-quality counterfactual explanations that are contiguous and sparse. The MG-CF method guides the perturbations on the query time series resulting in significantly sparse and more contiguous explanations than other state-of-the-art methods. This is the first effort to leverage the a-priori mined motifs to produce high-quality counterfactual explanations. There are spaces for extensions of our work with time series datasets with high dimensionality and multiple-classes case. As a future direction of this work, we would like to extend our work to multiple-class case time series dataset and also focus on generating the counterfactual explanations for high dimensional time series data.

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. Ates, E., Aksar, B., Leung, V.J., Coskun, A.K.: Counterfactual explanations for multivariate time series. In: 2021 International Conference on Applied Artificial Intelligence (ICAPAI), pp. 1–8. IEEE (2021)
3. Boubrahimi, S.F., Aydin, B., Schuh, M.A., Kempton, D., Angryk, R.A., Ma, R.: Spatiotemporal interpolation methods for solar event trajectories. *APJs* **236**(1), 23 (2018)
4. Dau, H.A., et al.: The UCR time series archive. *IEEE/CAA J. Automat. Sin.* **6**(6), 1293–1305 (2019)
5. Delaney, E., Greene, D., Keane, M.T.: Instance-based counterfactual explanations for time series classification. In: Sánchez-Ruiz, A.A., Floyd, M.W. (eds.) ICCBR 2021. LNCS (LNAI), vol. 12877, pp. 32–47. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86957-1_3
6. Guidotti, R., D’Onofrio, M.: Matrix profile-based interpretable time series classifier. *Front. Artif. Intell.* **4** (2021)
7. Guidotti, R., Monreale, A., Ruggieri, S., Pedreschi, D., Turini, F., Giannotti, F.: Local rule-based explanations of black box decision systems. arXiv preprint [arXiv:1805.10820](https://arxiv.org/abs/1805.10820) (2018)

8. Klaise, J., Looveren, A.V., Vacanti, G., Coca, A.: Alibi explain: algorithms for explaining machine learning models. *J. Mach. Learn. Res.* 22(181), 1–7 (2021), <http://jmlr.org/papers/v22/21-0017.html>
9. Kundu, S.: AI in medicine must be explainable. *Nat. Med.* **27**(8), 1328–1328 (2021)
10. Lin, J., Keogh, E., Fu, A., Van Herle, H.: Approximations to magic: finding unusual medical time series. In: 18th IEEE Symposium on Computer-Based Medical Systems (CBMS'05), pp. 329–334. IEEE (2005)
11. Lin, J., Keogh, E., Lonardi, S., Lankford, J.P., Nystrom, D.M.: Visually mining and monitoring massive time series. In: Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 460–469 (2004)
12. Lines, J., Davis, L.M., Hills, J., Bagnall, A.: A shapelet transform for time series classification. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 289–297 (2012)
13. Van Looveren, A., Klaise, J.: Interpretable counterfactual explanations guided by prototypes. In: Oliver, N., Pérez-Cruz, F., Kramer, S., Read, J., Lozano, J.A. (eds.) ECML PKDD 2021. LNCS (LNAI), vol. 12976, pp. 650–665. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-86520-7_40
14. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **30** (2017)
15. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5188–5196 (2015)
16. Mothilal, R.K., Sharma, A., Tan, C.: Explaining machine learning classifiers through diverse counterfactual explanations. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, pp. 607–617 (2020)
17. Mueen, A., Keogh, E., Young, N.: Logical-shapelets: an expressive primitive for time series classification. In: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1154–1162 (2011)
18. Nguyen, T.T., Le Nguyen, T., Ifrim, G.: A model-agnostic approach to quantifying the informativeness of explanation methods for time series classification. In: Lemaire, V., Malinowski, S., Bagnall, A., Guyet, T., Tavenard, R., Ifrim, G. (eds.) AALTD 2020. LNCS (LNAI), vol. 12588, pp. 77–94. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-65742-0_6
19. Ribeiro, M.T., Singh, S., Guestrin, C.: “Why should i trust you?” explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1135–1144 (2016)
20. Schlegel, U., Arnout, H., El-Assady, M., Oelke, D., Keim, D.A.: Towards a rigorous evaluation of XAI methods on time series. In: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), pp. 4197–4201. IEEE (2019)
21. Schlegel, U., Vo, D.L., Keim, D.A., Seebacher, D.: Ts-mule: local interpretable model-agnostic explanations for time series forecast models. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases, pp. 5–14. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-93736-2_1
22. Schleich, M., Geng, Z., Zhang, Y., Suciu, D.: Geco: quality counterfactual explanations in real time. arXiv preprint [arXiv:2101.01292](https://arxiv.org/abs/2101.01292) (2021)
23. Slack, D., Hilgard, S., Jia, E., Singh, S., Lakkaraju, H.: How can we fool lime and shap? adversarial attacks on post hoc explanation methods (2019)
24. Verma, S., Dickerson, J., Hines, K.: Counterfactual explanations for machine learning: a review. arXiv preprint [arXiv:2010.10596](https://arxiv.org/abs/2010.10596) (2020)

25. Wachter, S., Mittelstadt, B., Russell, C.: Counterfactual explanations without opening the black box: automated decisions and the GDPR. *Harv. JL Tech.* **31**, 841 (2017)
26. Xu, Y., et al.: Deep learning predicts lung cancer treatment response from serial medical imaging. *Clin. Cancer Res.* **25**(11), 3266–3275 (2019)
27. Ye, L., Keogh, E.: Time series shapelets: a new primitive for data mining. In: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 947–956 (2009)
28. Ye, L., Keogh, E.: Time series shapelets: a novel technique that allows accurate, interpretable and fast classification. *Data Mining Knowl. Discov.* **22**(1), 149–182 (2011)
29. Zhou, Z., Hooker, G., Wang, F.: S-lime: stabilized-lime for model explanation. In: *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 2429–2438 (2021)