XAI for Time Series Classification: Evaluating the Benefits of Model Inspection for End-Users

Brigt Håvardstun[1], Cèsar Ferri[2], Kristian Flikka[3], and Jan Arne Telle[1]

[1] Department of Informatics, University of Bergen, Norway
[2] VRAIN, UPV-Polytechnic University of Valencia, Spain
[3] Eviny ASA, Norway

Abstract. We present an XAI tool for time series classification providing model-agnostic instance-based post-hoc explanations, by means of prototypes and counterfactuals. Additionally, our tool allows for model inspection on instances generated by the user, to navigate the boundary between classes. This will allow the user to test and improve their hypotheses when formulating a model of the black box classification. We perform a human-grounded evaluation with forward simulation, to contribute a quantitative end-user evaluation to the field of XAI for time series.

1 Introduction

Temporal data is encountered in many real-world applications ranging from patient data in healthcare (Rajkomar et al., 2018 [14]) to the field of cyber security (Susto et al., 2018 [19]). Deep learning methods have been successful in time series classification (e.g., Fawaz et al., 2019 [6]; Rajkomar et al., 2018 [14]; Susto et al., 2018 [19]), but such methods are not easily interpretable, and often viewed as black boxes, which limits their applications when user trust in the decision process is crucial. To enable the analysis of these black-box models we revert to post-hoc interpretability. Recent research has focused on adapting existing methods to time series, both specific methods like SHAP-LIME (Guillemé et al., 2019) [8], Saliency Methods (Ismail et al., 2020 [10]) and Counterfactuals (Delaney et al., 2021 [4]), and also combinations of these (Schlegel et al [15]).

However, compared to images and text, time series data are not intuitively understandable to humans (Siddiqui et al., 2019 [17]). This makes interpretability of time series extra demanding, both when it comes to understanding how users will react to the provided explanations and also to the evaluation of its usefulness. Nevertheless, as humans learn and reason by forming mental representations of concepts based on examples, and any machine learning model has been trained on data, then data e.g. in the form of prototypes and counterfactuals is indeed the natural common language between the user and this model. In addition, several studies have highlighted the need to rethink new ways of interaction with an XAI algorithm, to allow for a dialogue between explainer and explainee, and to enable model inspection at will, see e.g. [12,1,16].

Hence, we advance an XAI time series tool that on the one hand provides users with instances of both prototype and counterfactual time series, and on the other hand lets the user generate their own instances for classification[4]. This enables active learning and allows the user to test and improve their hypotheses when formulating a mental model of the black box classification.

We perform a quantitative user evaluation, thereby meeting a demand from the XAI research community [20,4,21], to measure how prototypes, counterfactuals and interactivity increases the understanding that the user has of the black box classification. Using the taxonomy of interpretability evaluation of Doshi-Velez and Kim [5] what we do is a Human-Grounded Evaluation with Forward Simulation.

We use 3 datasets having a binary classification (e.g. 24-hour power demand of a household in Winter versus Summer) and for each dataset we train an AI model. Note that it is generally more difficult to explain the classification of an AI model than the classification of a real-world dataset. This since the dataset contains only real instances, whereas the AI model classifies *any* instance, also artifical ones. In this work, for prototypes we use real instances from the dataset, whereas we produce counterfactuals by combining real instances with artifical ones, based on the NativeGuide method of Delaney et al [4]. For tests we have chosen to use real instances.

In the remainder of this paper, we first discuss related work, then we give definitions, followed by a description of the tool and a discussion of the user evaluation results.

2 Related Work

Research on XAI for time series classification has progressed similarly to XAI in general, with earlier work focused on feature-importance rather than on instance-based methods using prototypes and counterfactuals. A comprehensive survey of XAI methods for time series can be found in Theissler et al [20]. Our own work deals with model-agnostic instance-based post-hoc explanations, by means of prototypes and counterfactuals. Poche et al [13] offers a survey of work on instance-based explanations for XAI mainly in the image domain, defining and classifying the various approaches in the litterature. In this paper we evaluate the use of instance-based explanations for time series by end-users. Ford et al [7] studied how non-experts handled post-hoc example-based explanations, however not in the time series domain. They found that even though these do assist users with correct judgement, people have significant difficulties dealing with misclassifications in an unfamiliar domain. Thus we should maybe not expect very high accuracy from our own evaluations on end-users, as time series are notoriously hard for humans to interpret [17].

Interactivity and model inspection have recently been seen as important in XAI. Beretta et al [2] argues that interactivity in XAI is a core value in the

[4] Available at github: Redacted to not show author names

interface between the model and the user, and that a user study is needed for a qualitative evaluation. Wang et al [21] develop an interactive XAI tool for loan applications that allows users to experiment with hypothetical input values and inspect their effect on model outcomes, and perform a user evaluation on MTurk.

User evaluations of XAI systems come in various forms. Doshi-Velez and Kim [5] offer a taxonomy of interpretability evaluation, from the gold standard of Application-oriented evaluations, to Human-grounded evaluations as we perform in this work, to Functionally-grounded evaluations that do not require human experimentations. For time series it seems the latter approach is the more common. Schlegel et al [15] applies several XAI methods previously used on image and text domain to time series, and introduces verification techniques specific to times series, in particular a perturbation analysis and a sequence evaluation, but they do not include any user evaluation of their systems. Likewise, Höllig et al [9] presents a Python package to provide a unified interface to interpretation of time series classification, but no user evaluation.

Delaney et al [4] provides a model-agnostic method for generating counterfactuals for time series classifiers, called Native Guide, that applies Class Activation Mappings [22] to select discriminative areas for modification. They end their paper by arguing that 'Given the ubiquitous nature of time series data and the frequent requirement for explanation, it is clear that experiments with human users and CBR solutions have much to offer in future work.' The survey [20] says about Native Guide that '...evaluating this promising approach involving end users could be promising for future work'. Indeed this is a part of what we do in the current paper as the counterfactuals we show users are exactly the ones generated by Native Guide.

3 Definitions

Let us present formal definitions for Time Series Classification (TSC) and recall basic notions.

Staying consistent with earlier notation [20,4] a time series $T = \{t_1, t_2, \ldots, t_m\}$ is an ordered set of m real-valued observations (or time steps). A time series dataset $D = \{T_1, T_2, ..., T_n\} \in R^{n \times m}$ is a collection of such time series where each time series has a class label c forming a vector of class labels. In this paper we consider only binary classification tasks. Given such a dataset, Time Series Classification is the task of training a mapping b from the space of possible inputs to a probability distribution over the class values. Thus, a black-box classifier $b(T)$ takes a time series T as input and predicts a probability output over the class values. Given a to-be-explained time series T, with predicted label $b(T) = c$ from the black-box classifier, a counterfactual explanation aims to find how T needs to (minimally) change to some T' for the system to classify it alternatively, as $b(T') = c' \neq c$. We refer to T' as a counterfactual explanation for T, without having to specify the target class since we consider only binary classification tasks. The minimality criterion usually refers to a notion of distance (proximity) between time series, but another criteria property can be sparsity (that T and

T' differ on few data points, or on few contiguous sequences of data points) and plausibility (that the instance is not an outlier).

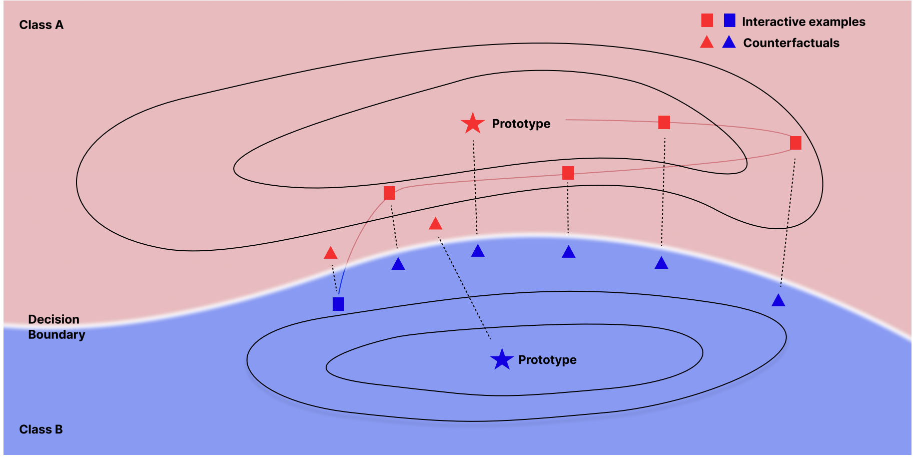

Fig. 1: Explanation types illustrated in a two-class decision space. The two prototypes are representatives of each class, positioned in the middle of the data density contour map. A path of interactive examples and their counterfactuals is shown to illustrate a possible user's journey.

Prototypes are time series exemplifying the main aspects responsible for a classifier's specific decision outcome. It can be a real instance (which is what we opt for in our tool) sampled from the dataset that is important and meaningful because it summarizes the shape of many other similar instances, or a synthetic one, for example a cluster centroid or an instance generated by following some ad-hoc processes. See Figure 1.

4 Prototypes, counterfactuals, and user-model inspection

In this section we introduce our tool for XAI on time series classification providing model-agnostic instance-based post-hoc explanation, by means of prototypes and counterfactuals. Additionally, one of our contributions is to allow for model inspection on instances generated by the explainee. This interactive tool allows the user to themselves navigate the boundary between classes. Starting from a prototype, while simultaneously seeing a counterfactual (see Figure 1) the user can change individual time points at will, and see the resulting classification. This will allow the user to test and improve their hypotheses when formulating a model of the black box classification.

In the rest of this section we describe our tool and start by giving an argument for its motivation, coming from the industry side. We then give a description of the algorithms for generating prototypes plus counterfactuals, and a discussion of how to present these to the user on-screen. We end by reporting on the design choices related to the interactive part, that allows the user to change individual time points and do model inspection.

4.1 The need for trust

When presenting output from black box ML systems to non-trained users, the need for explanations arises from at least two angles: The user must trust the results, and the user must, to a sufficient degree understand the results. To use a real-world example from energy companies [18], if a system based on charging data tells a car owner that their car seems faulty, or tells a home owner that their power consumption deviates from expectations - they would need to trust their correctness, otherwise the message will be ignored. When trust is established, the next step is usually to fix the situation. If the car owner understands the reason for their car charging being classified as faulty, he/she can bring it to the vendor or repair shop with concrete information that can be used to fix it. Similarly, if the home owner understands why their consumption is classified as deviating - they could perhaps fix a broken appliance - or adjust their consumption to a more favorable price pattern. The more complicated the data and classifications are (for example consisting of several series of data measuring power, temperature, etc. - and/or having temporal patterns like slowly decreasing trends), the more challenging will the explanation be. Thus - exploring different approaches and tools for explanations for non-trained users is of great use when companies plan on applying machine learning on complex data. Without trust the models will be ignored, and without an understanding of what the actual problem is no appropriate action can be taken.

4.2 Generating and presenting Prototypes and Counterfactuals

We have opted to use cluster centers as our prototypes. To ensure the prototypes we selected have high plausibility, we use a $k-$medoids algorithm to find the prototypes, see [11]. Specifically we used KMedoids from sklearn_extra.cluster package. We used default configuration of KMedoids, with Euclidean distance as the distance metric.

There is a growing consensus that counterfactuals provide robust and informative explanations to a query time series whose classification is to be explained. In the time series domain the visualization of counterfactuals is straightforward. We have opted to use a novel method for generating counterfactuals for time series called Native Guide, developed recently by Delaney et al [4]. This method extracts counterfactual time series, named Native Guides, starting from initial training data. Starting from the query time series whose classification is to be explained, the Native Guide method starts by finding a counterfactual time series belonging to the dataset that is close to the query. This Native series are then adapted in a Guided way to generate novel counterfactuals, following four identified key properties for good counterfactuals: proximity, sparsity, plausibility, and diversity. The Native Guide counterfactual generation method uses Class Activation Mappings to Guide the counterfactual generation from the Native series. The use of CAM in itself puts some limitations on the AI model used, e.g. having the last layer be global average pooling. To ensure compatibility between the model doing the classification and the Native Guide counterfactual generator, we therefore closely follow the time series classification model implementation of Delaney et al available here.

When presenting time series to the user we have opted to use two colors for the two classes, namely Blue and Pink. Since a counterfactual, say Blue, will be used to explain the classification of a given query Pink time series, we have chosen to present both at the same time to the user. Thus, we plot the query with Pink lines, and then show only the deviation of the counterfactual by Blue dotted lines. See Figure 2 which shows also that the user is allowed to make interactive changes, as described in the next subsection.

4.3 Allowing interaction and model inspection

A central aspect of our tool is that it allows the user to alter individual data points and do model inspection. In real-time the tool will update the model classification by changing the color of the time series if the classification changes, showing the model confidence in this classification. It will also update to a new counterfactual. This enables active learning and allows the user to test and improve their hypotheses when formulating a mental model of the black box classification. In Figure 2 we see 4 screen shots of an actual session with the tool. Note how the user can explore their understanding of the classification by progressive changes to the current time series. Compare also with Figure 1.

5 End-user Evaluation with Forward Simulation

We have not before seen such an interactive tool for time series in the research literature. As time series are notoriously difficult for human users, most of the human evaluations of XAI methods done so far have been qualitative and involving domain experts. Our goal in this work has been to add the new dimension of quantitative user evaluations involving end-users, i.e. non-experts, on XAI for time series. In this section we present this user evaluation. We start by presenting the discussion leading up to the 3 time series datasets chosen for the evaluation. We then discuss the 3 distinct survey groups that will be given different abilities in the training stage. We end by giving the results of the evaluation, on each pairing of dataset and survey group, and a discussion of their significance.

Our tool is designed for univariate datasets. Moreover, since our test users are not domain experts and time series are notoriously non-intuitive to us humans we had some desiderata when choosing datasets for our survey. Firstly, we wanted univariate datasets with a binary classification, Secondly, we wanted datasets on time series with not too many data points. Thirdly, we wanted datasets where it is fairly easy for a non-expert to understand the domain. Datasets satisfying

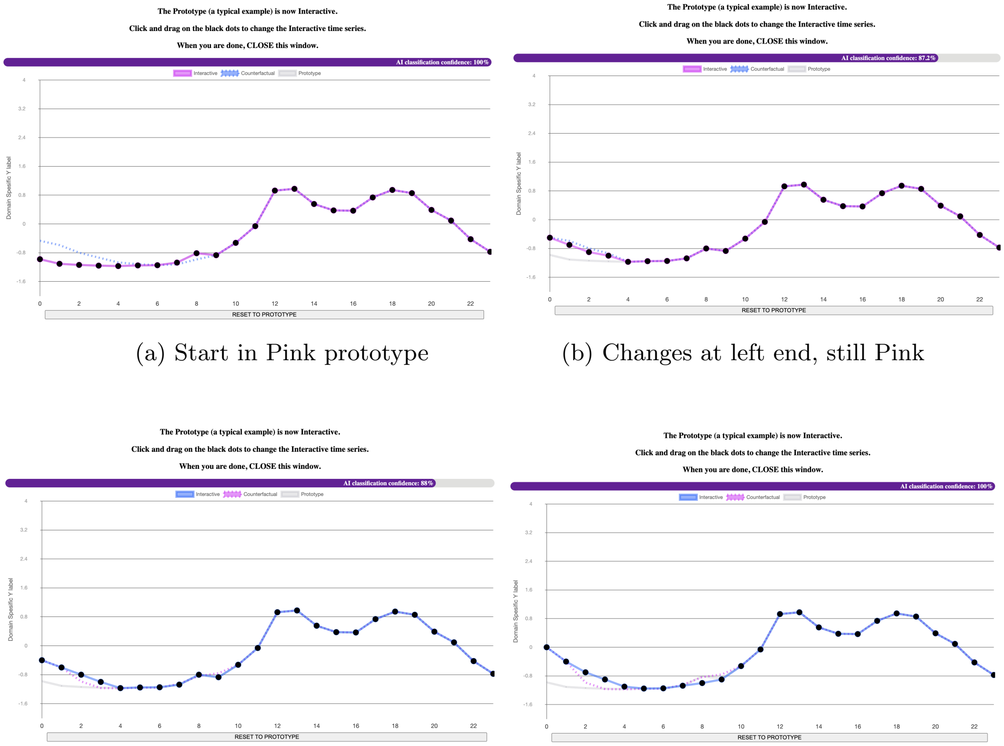

The Prototype (a typical example) is now Interactive.

Click and drag on the black dots to change the Interactive time series.

When you are done, CLOSE this window.

AI classification confidence: 100%

(a) Start in Pink prototype

The Prototype (a typical example) is now Interactive.

Click and drag on the black dots to change the Interactive time series.

When you are done, CLOSE this window.

AI classification confidence: 87.2%

(b) Changes at left end, still Pink

The Prototype (a typical example) is now Interactive.

Click and drag on the black dots to change the Interactive time series.

When you are done, CLOSE this window.

AI classification confidence: 88%

(c) Further changes, now Blue

The Prototype (a typical example) is now Interactive.

Click and drag on the black dots to change the Interactive time series.

When you are done, CLOSE this window.

AI classification confidence: 100%

(d) Last changes, still Blue

Fig. 2: Series of model inspections. (a) User starts with a Pink prototype and a counterfactual. (b) Makes changes to left end of series so that confidence (top bar) of model classification drops, plus new counterfactual, but same color/class. (c) Makes further changes and now the model classification switches. (d) Last changes made by user and confidence of model classification increases.

these criteria will increase the prospective of providing constructive quantitative feedback, and also they form the more common situation where an explanation would potentially be provided for end-users. However, it is important to note that apart from the above constraints we did not want simple datasets where the binary classification is particularly easy or could be described (e.g. by ourselves) in any straightforward way. We chose the following three datasets satisfying the above criteria:

– From UCR [3]: Italy Power Demand. This dataset shows power demand in Italian households over a 24-hour time period, and classifies these into Winter (October-March) and Summer (April-September).

- From UCR [3]: Chinatown. This dataset shows the number of pedestrians on a particular street corner of Chinatown in Melbourne over a 24-hour time period, and classifies these into Weekend (Sat-Sun) and Weekdays (Mon-Fri).
- From Eviny: Car Charging. This dataset shows the power demand at a particular charging station for electric vehicles over a 24-hour period, and classifies these into Weekend (Sat-Sun) and Weekdays (Mon-Fri).

As mentioned earlier, the three main components used in our interactive XAI system for time series are prototypes (PT), counterfactuals (CF), and model inspection (MI) with user-generated instances. Our survey groups will enter a 3-stage process, as follows:

- Intro: A short introduction is given to the relevant components, the dataset, training stage, and testing stage.
- Training: Group dependent.
- Testing: 10 randomly selected time series from the dataset are shown, and for each one the user must guess the AI model classification. See Figure below.

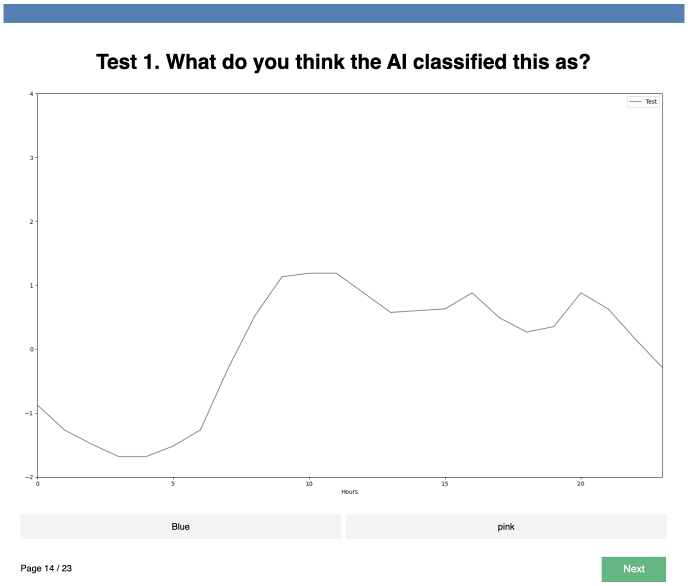

Test 1. What do you think the AI classified this as?

Fig. 3: Test case, for ItalyPowerDemand.

To cover the distribution of time series within each class, we wanted to show users more than one prototype for each class. However, we did not want to overload the user with information, hence we opted to have 6 prototypes in total, three for each class. We will have three survey groups (PT, PT+CF, PT+CF+MI) depending on what is made available to users in the Training stage:

- PT: One at a time, the user is shown 3 prototypes from class A, then 3 from class B, and finally shown a screen with all 6 prototypes.

- PT+CF: One at a time, the user is shown 3 prototypes from class A together with a counterfactual from class B, then 3 converse pairs, and finally shown a screen with all 6 pairs.
- PT+CF+MI: The user is shown same as PT+CF but model inspection is permitted, with prototypes being interactive to allow iterative changes of any chosen time points, and the ensuing classification and also new counterfactual continually updated.

In Figure 2 we see how a user in group PT+CF+MI is shown a pair consisting of a prototype and a counterfactual and is allowed to modify data points. A user in group PT+CF is shown a single such pair consisting of prototype and counterfactual but without the ability to make modifications, while a user in group PT is shown only the prototype. In the last part of the training stage the users are shown all 6 prototypes/counterfactuals on one screen, to easily make a comparison, as in Figure 4 for group PT+CF+MI.

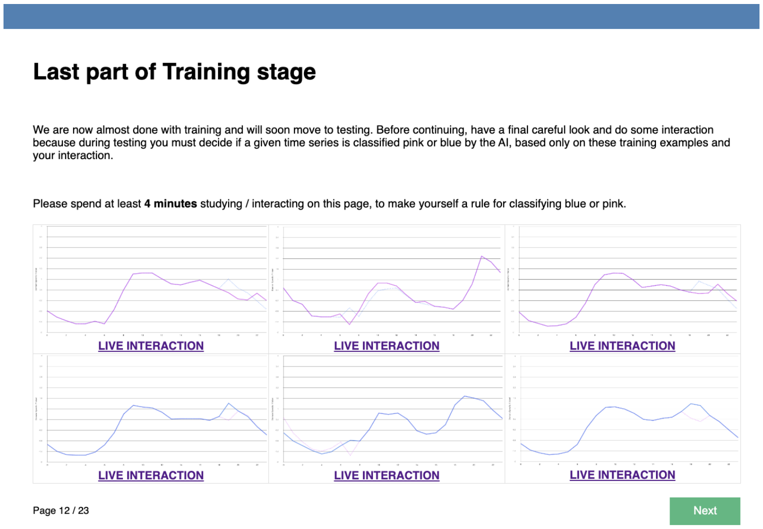

Last part of Training stage

We are now almost done with training and will soon move to testing. Before continuing, have a final careful look and do some interaction because during testing you must decide if a given time series is classified pink or blue by the AI, based only on these training examples and your interaction.

Please spend at least **4 minutes** studying / interacting on this page, to make yourself a rule for classifying blue or pink.

Fig. 4: Blue and Pink prototypes with counterfactuals, for ItalyPowerDemand

5.1 Evaluation results

Let us present the results of the user survey. In total, we had 65 voluntary participants, who were either students in a university-level informatics course, or researchers in informatics, none of whom received compensation. The participants were presented with the survey and freely chose to participate. To our knowledge none of them had any previous background with time series classification. The participants were divided into the 3 survey groups. After introduction and training they were asked to answer 10 test questions from each of up to 3 datasets. Those who spent less than a minimum amount of time (less than 1 minute for training and testing combined on a single dataset) were discarded.

In Table 1 we see the accuracy and number of tests (i.e. number of participants times 10) for each of the three survey groups, on each of the three datasets. Accuracy is the percentage of correct test answers. The rightmost column gives the aggregated information for each survey group, and the bottom row the aggregated information for each dataset. The value in the bottom right corner shows that the overall accuracy was 66.7%, thus clearly better than a random guess, and satisfactory given the complicated nature of these time series classifications.

Let us compare the performance between the three survey groups. Perhaps surprisingly, on aggregate we see that survey group PT that was only shown the prototypes, did slightly better than the other two, with accuracy of 70.2% versus 65.9% and 64%. This could indicate that non-expert users are not necessarily able to use the extra information provided by counterfactuals and model inspection in a meaningful way. It could also mean that a 'rule-of-thumb' that appears useful based only on prototypes (or prototypes + counterfactuals) actually works well on a high percentage, say 80%, of test instances. However, after close model inspection a user in group PT+CF+MI may discover that this rule is not precise (as it fails on several instances) leaving the user to discard this rule-of-thumb, and subsequently actually performing worse on the tests. A final possible explanation for why group PT+CF+MI did not achieve higher accuracy is the fact that the average time spent by users in total on all 3 datasets was about the same, around 15 minutes: group PT 954 seconds, group PT+CF 934 seconds, group PT+CF+MI 966 seconds. However, it seems natural that a user doing model inspection to learn a better rule should have spent more time than one who cannot do model inspection, making us question how careful the users in group PT+CF+MI were. On the other hand, there is not a clear trend when we compare accuracy versus time spent for users in this group.

Let us turn to comparing between the 3 datasets. Interestingly, all 3 survey groups had the highest accuracy on Chinatown and the lowest accuracy on CarCharging. We asked some users about rules they had been using for their own mental classification, and the most mentioned rule was for Chinatown, something like the following: 'if there is a sharp dip at the beginning of the series then it is Blue, and otherwise Pink'. Compare this rule to Figure 2 from the Chinatown dataset. See also Figures 4 and 3 which for ItalyPowerDemand gives all prototypes and counterfactuals, plus an example of a test, to get an impression of how difficult the classification indeed is for this dataset. For the test shown in Figure 3 the average accuracy was 59.3%. Note the users could not navigate back to see the prototypes when answering the tests.

6 Conclusion

As companies and controllers must cope with the EU regulations in form of GDPR and the AI Act, explainability that allows model inspection by end-users

Survey ID	Data	ItalyPowerDemand	Chinatown	CarCharging	All 3 datasets
PT	accuracy	65.9%	79.4%	65.3%	70,2%
	tests	170	170	170	510
PT+CF	accuracy	61.9%	81.3%	54.4%	65.9%
	tests	160	160	160	480
PT+CF+MI	accuracy	60.2%	76.2%	55.7%	64.0%
	tests	240	210	230	680
All groups	accuracy	62.7%	79.0%	58,4%	66.7%
	tests	570	540	560	1670

Table 1: Overview of accuracy and number of tests by survey group and dataset.

may very well become the target for developers. In domains where we humans have poor intuition, such as time series, this may pose several challenges. In this work we have contributed a tool for XAI on time series classification that allows such model inspection. The evaluation results show that users can then successfully perform a difficult forward simulation test. However, to attain the full benefits from model inspection for such a complicated domain, it seems necessary to have highly motivated users that are willing to spend more time with such a tool, in order to form better mental models of the black-box classification. As future work, we propose to explore the use of simplified versions of time series prototypes generated by Machine Teaching techniques to explain time series classification models.

References

1. Ashraf M. Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan S. Kankanhalli. Trends and trajectories for explainable, accountable and intelligible systems: An HCI research agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018*, page 582. ACM, 2018.
2. Isacco Beretta, Eleonora Cappuccio, and Marta Marchiori Manerba. User-driven counterfactual generator: A human centered exploration. In *1st World Conference on eXplainable Artificial Intelligence (xAI-2023)*, volume 3554 of *CEUR Workshop Proceedings*, pages 83–88, 2023.
3. Hoang Anh Dau, Eamonn Keogh, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, Yanping, Bing Hu, Nurjahan Begum, Anthony Bagnall, Abdullah Mueen, Gustavo Batista, and Hexagon-ML. The UCR time series classification archive, 2018.
4. Eoin Delaney, Derek Greene, and Mark T. Keane. Instance-based counterfactual explanations for time series classification. In *Case-Based Reasoning Research and Development - 29th International Conference, ICCBR 2021*, volume 12877 of *LNCS*, pages 32–47. Springer, 2021.
5. Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *CoRR*, abs/1702.08608, 2017.
6. Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Deep learning for time series classification: a review. *Data Min. Knowl. Discov.*, 33(4):917–963, 2019.

7. Courtney Ford and Mark T Keane. Explaining classifications to non experts: An xai user study of post hoc explanations for a classifier when people lack expertise, 2022. `arXiv:2212.09342`.
8. Maël Guillemé, Véronique Masson, Laurence Rozé, and Alexandre Termier. Agnostic local explanation for time series classification. In *31st IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2019, Portland, OR, USA, November 4-6, 2019*, pages 432–439. IEEE, 2019. `doi:10.1109/ICTAI.2019.00067`.
9. Jacqueline Höllig, Cedric Kulbach, and Steffen Thoma. Tsinterpret: A python package for the interpretability of time series classification. *J. Open Source Softw.*, 8(87):5220, 2023.
10. Aya Abdelsalam Ismail, Mohamed K. Gunady, Héctor Corrada Bravo, and Soheil Feizi. Benchmarking deep learning interpretability in time series predictions. In *Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
11. Leonard Kaufman and Peter J. Rousseeuw. *Finding Groups in Data: An Introduction to Cluster Analysis.* John Wiley, 1990. `doi:10.1002/9780470316801`.
12. Tim Miller. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267:1–38, 2019.
13. Antonin Poché, Lucas Hervier, and Mohamed Chafik Bakkay. Natural example-based explainability: A survey. In *Explainable Artificial Intelligence - 2023*, volume 1902, pages 24–47. Springer, 2023.
14. Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M. Dai ..., and Jeffrey Dean. Scalable and accurate deep learning with electronic health records. *npj Digit. Medicine*, 1, 2018.
15. Udo Schlegel, Hiba Arnout, Mennatallah El-Assady, Daniela Oelke, and Daniel A. Keim. Towards a rigorous evaluation of XAI methods on time series. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pages 4321–4325, Piscataway, NJ, 2019. IEEE. `doi:10.1109/ICCVW.2019.00516`.
16. Ben Shneiderman. Human-centered artificial intelligence: Reliable, safe & trustworthy. *Int. J. Hum. Comput. Interact.*, 36(6):495–504, 2020.
17. Shoaib Ahmed Siddiqui, Dominique Mercier, Mohsin Munir, Andreas Dengel, and Sheraz Ahmed. Tsviz: Demystification of deep learning models for time-series analysis. *IEEE Access*, 7:67027–67040, 2019. `doi:10.1109/ACCESS.2019.2912823`.
18. Laisuo Su, Shuyan Zhang, Alan JH McGaughey, B Reeja-Jayan, and Arumugam Manthiram. Battery charge curve prediction via feature extraction and supervised machine learning. *Advanced Science*, 10(26):2301737, 2023.
19. Gian Antonio Susto, Angelo Cenedese, and Matteo Terzi. Chapter 9 - time-series classification methods: Review and applications to power systems data. In *Big Data Application in Power Systems*, pages 179–220. Elsevier, 2018.
20. Andreas Theissler, Francesco Spinnato, Udo Schlegel, and Riccardo Guidotti. Explainable AI for time series classification: A review, taxonomy and research directions. *IEEE Access*, 10:100700–100724, 2022.
21. Zijie J. Wang, Jennifer Wortman Vaughan, Rich Caruana, and Duen Horng Chau. GAM coach: Towards interactive and user-centered algorithmic recourse. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, CHI 2023, Hamburg, Germany, April 23-28, 2023*, pages 835:1–835:20. ACM, 2023.
22. Bolei Zhou, Aditya Khosla, Àgata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2921–2929. IEEE Computer Society, 2016.