

# Explaining Fatigue in Runners using Time Series Analysis on Wearable Sensor Data

Bhavathy Kathirgamanathan<sup>1</sup>, Thu Trang Nguyen<sup>1</sup>, Brian Caulfield<sup>2</sup>,  
Georgiana Ifrim<sup>1</sup>, and Pádraig Cunningham<sup>1</sup>

<sup>1</sup> School of Computer Science, University College Dublin, Ireland

<sup>2</sup> School of Public Health, Physiotherapy and Sports Science, University College  
Dublin, Ireland

**Abstract.** Running while fatigued poses an increased risk of injury. Wearable sensors allow the overall running kinematics or running pattern to be captured as time series signals. The changes that happen in the running pattern due to fatigue, although prominent enough to increase the risk of injury, are generally only seen as subtle differences in the signal itself and hence are difficult to differentiate using purely visual inspection. In this paper, we introduce a time series dataset of motion capture data from runners before and after a fatiguing intervention. The total dataset consists of more than 5500 instances and was collected from 19 participants. The evaluation presented in this paper first looks at the effectiveness of a data aggregation technique called time series barycenters which is shown to improve classification performance. We evaluate and compare a set of classifiers and explanation methods for this problem, and select the most informative classifier and explanation for this dataset. We then present feedback from a domain expert on the insights offered by the explanations.

**Keywords:** Time Series · Fatigue · Explanation · Time Series Barycenters

## 1 Introduction

The onset of fatigue in runners can increase their risk of injury due to the higher impact accelerations and alterations in the overall running kinematics [20]. The presence of fatigue can cause a lack of overall control over joint motion and muscle contraction which causes this change in the running kinematics [3]. Furthermore, the changes seen in runners are often specific to the individual as there is known to be no generalizable template of running technique and hence feedback should be individualised to each runner [33].

Running kinematics for an individual can be captured using wearable sensors such as the Shimmer<sup>1</sup> sensors. The data captured from these sensors are time series in nature. The majority of research in this exercise classification domain uses extracted features from these time series as input into the models [23].

---

<sup>1</sup> <http://shimmersensing.com>

This requires a level of understanding regarding which features may be more important for a given classification task. However, many problems in the exercise domain such as the task considered in this paper, where we predict fatigue in runners, have subtle differences between classes which may not be noticeable by the eye alone. Using time series classification techniques allows the signal to be used directly as input into the classifier as raw signals. Additionally, this allows interpretable classification techniques to be used to highlight the particular regions of interest within a time series. This information is highly valuable as it can inform the runner when they are changing their form and hence can help prevent fatigue related injuries.

The data captured from the sensors can be noisy and due to the individualised nature of the running pattern, automated segmentation procedures are not always completely accurate. This can impact the classification performance and the stability of the interpretations made. Hence, we employ a barycenter averaging strategy [5] which aggregates and smoothes multiple strides together which we expect will improve the classification performance and improve the overall stability of the explanations.

The key contributions of this paper are as follows:

- We present the methods used to capture, segment, and classify motion capture data for a running fatigue prediction task. Furthermore, we release this time series dataset for use by the machine learning community. To access the data, see <https://zenodo.org/record/7997851> [10].
- We evaluate the classification performance of various classifiers on this running fatigue prediction dataset using various models that are used for interpretable time series classification. Furthermore, we evaluate the impact on classifier performance when using a barycenter averaging technique as a smoothing strategy on the wearable sensor data.
- We produce explanations for the time series data and discuss the insights into the impact of fatigue in runners that can be deduced from these explanations.

In the next section we give an overview of some related work in the areas of running and fatigue, time series classification, time series barycenters, and time series explanation. Section 3 details the data collection and processing methods used. In Section 4, we evaluate the classification performance of the datasets and compare the performance of using individual strides vs the smoothed barycenter strides. Section 5 presents an evaluation of a few explanation methods and some insights gained in the data from the selected explanation method.

## 2 Related Work

### 2.1 Impact of Fatigue on Runners

Running has become a popular form of exercise over the past few decades and this has led to an increase in the rate of running-related musculoskeletal injuries [36].

There are many factors that can cause these injuries and often they are due to faults in the technique. Faults in a runners technique can start to happen or get exaggerated as they get fatigued as they start to lose control over their running form. This lack of control in the technique is known to be a primary contributor to running related injuries [38]. Altering the kinematics while running is often an attempt by the runner to minimise the overall metabolic cost due to fatigue [8]. Hence, it is important to identify when a runner’s form is beginning to change. Being able to objectively measure these alterations is difficult and often runners are required to rely on their own assessment of running form [4]. Wearable sensors capture these kinematic changes hence allowing an objective analysis of running form to be made.

## 2.2 Time Series Classification

A time series is a time based sequence of observations,  $x_i(t)$ ; [ $i = 1, \dots, m$ ;  $t = 1, \dots, p$ ], where  $x_i(t)$  is the observation for the  $i^{th}$  dimension at time point  $t$ . The time series is univariate when  $m$  is 1 and multivariate when  $m$  is greater than or equal to 2.

Time Series Classification (TSC) techniques are generally classed as distance based, shapelet based, symbol based, deep learning methods and ensembles. Distance based techniques are known to perform well with the common benchmark of 1-Nearest Neighbour (1-NN) with dynamic time warping (DTW) being one of the most popular TSC approaches [1]. 1-NN DTW however has a high computational complexity and sometimes faces challenges to the accuracy in the presence of noise [29]. Time series shapelets are subsequences of the data which capture the portion of the time series which is maximally representative of the class. There have been various implementations of shapelets such as the original work by Ye and Keogh [39], Shapelet Transform [15] and Learning Shapelets [7]. Shapelets, however are not widely used as they can be computationally expensive [26].

Symbol based techniques include Mr-SEQL [13], WEASEL+MUSE [31] and BOSS [28]. Symbol based techniques work by passing a sliding window over each time series and representing each window with a word of symbols. This transform is done using techniques such as Symbolic Aggregate Approximation (SAX) [14] or Symbolic Fourier Transform (SFA) [30]. Although these methods are known to work well, they often suffer a long runtime. Most recent ensemble methods work by transforming the time series into a new feature space. From this idea, COTE (Collective of Transformation based ensembles) was developed which works by ensembling different classifiers over different time series representations [2] This technique was later extended as HIVE-COTE (Hierarchical Vote system) which uses a hierarchical structure based on probabilistic voting [16] HIVE-COTE has been shown to achieve high accuracy, however has a large computational complexity and hence is deemed impractical to run on large data mining problems.

Deep learning techniques such as ResNet and Fully Convolutional Networks (FCN) have also been shown to perform well for time series classification. However, the most popular and current state-of-the art technique is ROCKET [6],

which is a method that borrows ideas from deep neural networks where a simple linear classifier is trained on random convolutional kernels. This method has been shown to achieve state-of-the-art performance whilst also maintaining a lower computational load than other state-of-the-art techniques.

For the evaluations in this paper we have selected four techniques, Ridge regression, Rocket, 1-NN-DTW, and Mr-SEQL as detailed in Section 4.

### 2.3 Time Series Barycenter

A time series barycenter is an average measure of a collection of time series. Typically time series averaging strategies are classified as local or global. Local strategies use pairwise averaging where a collection of series are iteratively averaged into a single average series. Local averaging is dependent on the order in which samples are aggregated and hence changing the order can give a different result. Recent advances look at global average strategies which compute the average of the set of time series simultaneously. These averaging strategies generally use a similarity metric to find the distance between the series. Dynamic Time Warping (DTW) [27] is one of the most common similarity metrics used for time series and hence is used in most time series barycenter calculation techniques. DTW maps the time series in a non-linear way and works to find the optimal alignment between the two series. Two popular global averaging strategies are the DTW Barycenter Averaging (DBA) proposed by Petitjean *et al.* [24] and the Soft-DTW Barycenter proposed by Cuturi *et al.* [5].

**DTW Barycenter Averaging (DBA)** computes the optimal average sequence within a group of series in DTW space by minimising the sum of the squared DTW distance between the average sequence and the group of series [24]. To compute the DTW barycenter of a set of time series  $D$  is the optimisation problem outlined in Equation 1.

$$\min_{\mu} \sum_{x \in D} DTW(\mu, x)^2 \quad (1)$$

Where  $x$  is a series belonging to the set  $D$  and  $\mu$  is a candidate barycenter. The DTW distance between each time series and a temporary average sequence (candidate barycenter) is calculated and the temporary average sequence is updated until the optimisation criteria is met [32].

**Soft-DTW Barycenter** computes the average sequence within a group of series by minimising the weighted sum of the Soft-DTW distance between the average sequence and the group of series. Soft-DTW is an extension of the DBA method where the min operator is replaced by the soft-min. This has the advantage of being differentiable with respect to all of its inputs. Soft-DTW also has the advantage where it considers all possible alignments. Soft-min can be computed as

shown by Equation 2 [35].

$$\text{softmin}_\gamma(a_1, \dots, a_n) = -\gamma \log \sum_i e^{-a_i/\gamma} \quad (2)$$

$\gamma$  controls the smoothing and hence as  $\gamma \rightarrow 0$ , the result gets closer to that of DTW [35].

Previous research has shown Soft-DTW barycenter averaging to be an effective way to aggregate time series data as it preserves the key features well [11] and hence this technique will be employed in our evaluations.

## 2.4 Time Series Interpretation

Time Series data often comes from domains such as healthcare where having explainable models is of importance. Many time series interpretation methods are based on explaining through visualisation and feature importance rather than instance based explanations [12]. Instance based techniques such as 1-NN DTW although often used as a benchmark for time series classification may not be ideal for explanation purposes. The nearest labelled neighbour can be found, however it is not known which parts of the time series are more influential in a given classification task [13]. Shapelet based techniques are by nature interpretable as they pick up on the section of the time series that is maximally representative of a class; however they are computationally inefficient. Deep learning methods, although black box by nature, have methods which allow explanations to be made based on the class weights. A popular method for interpretation of time series is the use of Class Activation Maps (CAM) [40] which is a saliency method that highlights discriminating parts of the time series (examples of this are shown in Figure 4).

Many model-agnostic methods used outside the field of time series generate local explanations using a perturbation strategy where features are slightly altered to gain insight into which features are more relevant to a model. For this, mapping functions need to be defined to direct how the perturbations should be done. LIME is one of the first techniques which uses this perturbation strategy [25] which works for image, text, and tabular data. SHapley Additive exPlanations (SHAP) [17] is a current widely used technique which is a unified measure of feature importance. TimeXplain is a recent development which works with SHAP by defining mappings that work on the time and frequency domain [21]. Hence, TimeXplain makes SHAP usable with time series data.

## 3 Data Collection and Processing

### 3.1 Experimental Setup and Protocol

Nineteen recreational runners were recruited to participate in this study. The participants were all free of lower limb injury and were regular runners (at least

2 runs/week). The study protocol was reviewed and approved by the human research ethics committee at the host institution.

A single lumbar mounted Shimmer Inertial Measurement Unit (IMU) was mounted on each participant while they completed the task in three segments. The tests were done on an outdoor track typically used for running. In the first segment, the participant was asked to complete a 400m run at a comfortable pace. In the second segment, the participant completed a beep test [19], which is a multi-stage fitness test where the runner is required to run continuously between two points 20m apart following an audio which produces ‘beeps’ which indicate when the person should start running from one end. As the test progresses, the interval between the ‘beeps’ reduces and hence the runner will be required to increase their pace. The test ends when the runner is unable to continue or when they are unable to keep up with the pace of the ‘beeps’. For this study, this was used as the fatigue intervention. In the final stage, the participant was required to repeat the 400m run, this time in their fatigued state. The fatigued run was completed immediately after the completion of the beep test. The sensor was capturing data during all three stages.

The IMU sensor captured acceleration, angular velocity and magnetometer data. The data was collected as three long time series (non-fatigued run, beep test, and fatigued run) and was then segmented into individual strides as described in Section 3.2. The sampling rate of the sensors was set to 256Hz. Altogether from the sensors, there were nine signals extracted: Acceleration (X,Y,Z), Angular Velocity (X,Y,Z), and Magnetometer (X,Y,Z). From the acceleration signals, the magnitude acceleration was derived ( $A_{mag} = \sqrt{X^2 + Y^2 + Z^2}$ ). This signal was primarily used for the analysis presented in this paper. Our analysis focuses on the two 400m runs which we frame as a binary classification task to distinguish between Non-Fatigued and Fatigued.

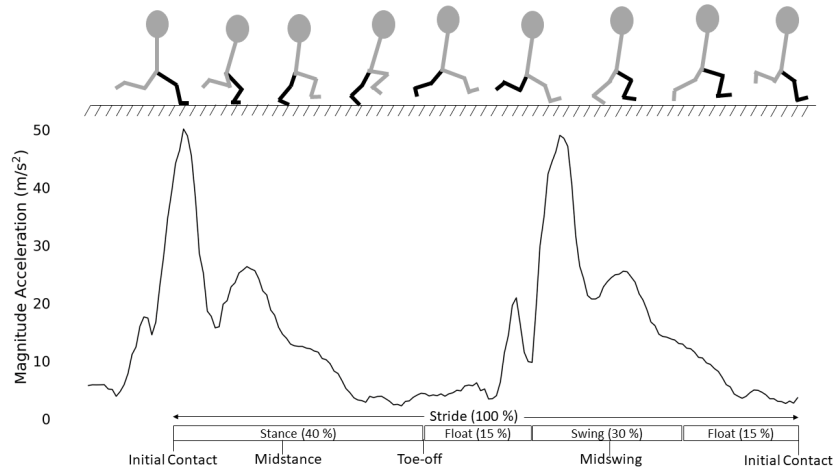


Fig. 1: Breakdown of a single stride

### 3.2 Segmentation and Pre-Processing strategy

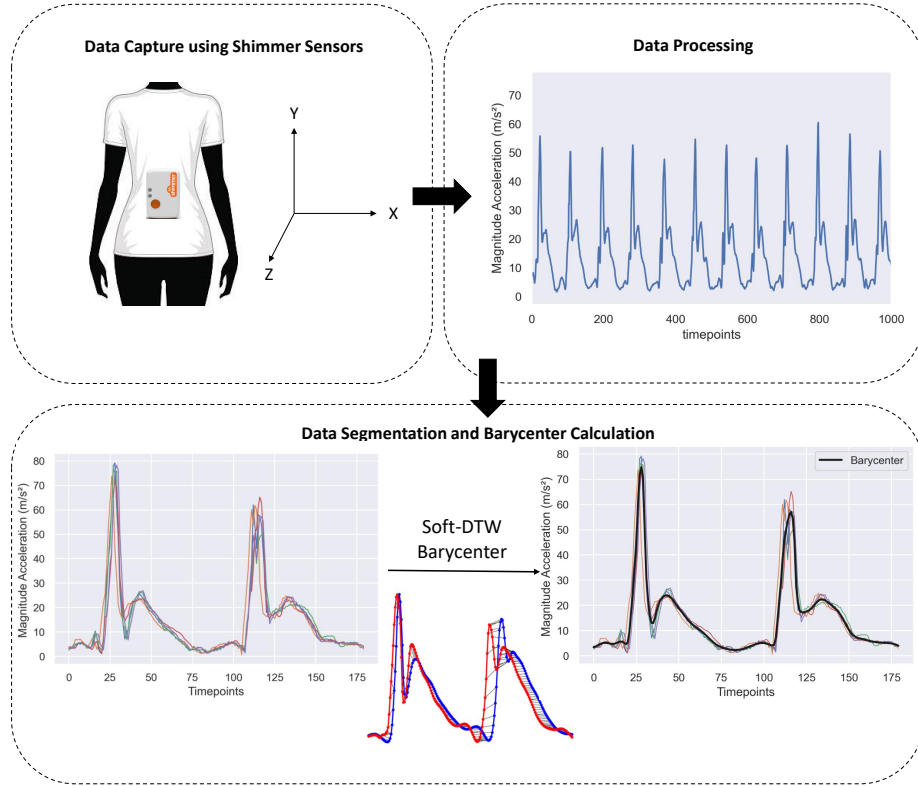


Fig. 2: Data collection and processing protocol; Data was initially collected using Shimmer sensors, The signals were segmented into individual strides and aggregated 5 strides at a time using a rolling Soft-DTW barycenter averaging technique.

The data was segmented into individual strides using a segmentation protocol as follows. Each instance was a stride or two steps (See Figure 1). To segment the data into individual strides, the acceleration in X-direction was first used to identify the left foot so each stride begins with the left limb. The peak point was identified and the minimum before the peak was used as the point for segmentation. The breakdown of a single stride is shown in Figure 1 where the first peak roughly corresponds to the left foot initial contact point. The segmented strides were then resampled to the length of the median stride as the signals are required to be of equal length for some of the interpretation strategies used later in this paper. The segmentation and processing were done using Python semi-automated scripts.

Figure 2 summarises the procedure that was used for the data collection and processing. Following data collection and segmentation, we use the barycenter aggregation technique as a smoothing technique for the running data. We employ a rolling barycenter calculation over every five strides. An example of the five strides and its corresponding barycenter can be observed in Figure 2. To check the effectiveness of using barycenters to represent the data, we evaluate the classification performance of the barycenters against the full set of strides in the next section (Section 4).

## 4 Classification Performance

The classification was performed at an individual participant level as previous research has shown that personalised models perform much better than global classifiers in this exercise domain due to variations between the biomechanics of individuals [9, 37]. Hence, we have 19 separate datasets with approximately 290 strides or instances per dataset and as a result 19 stand-alone models that we can evaluate. The task is posed as a binary classification task where we predict fatigued strides against non-fatigued strides. The four time-series classification models considered in the evaluation were selected for accuracy and interpretability:

- **Ridge regression (RidgeCV)**<sup>2</sup> is a linear classifier and is used as one of the simplest models and takes the time series data as tabular data. Ridge regression can work well on time series data where the values rather than the shape of the data are of influence.
- **Rocket** is currently the state-of-the-art technique for time series classification [6]. Rocket works through the generation of random convolutional kernels which are then convolved along the time series to produce a feature map. These features are then input into a simple linear classifier such as a Ridge Classifier or Logistic Regression. In our evaluation, we use the default 10,000 kernels and a random state of 0.
- **1-Nearest Neighbour (1-NN) with Dynamic Time Warping (DTW)** has been used commonly as a benchmark classifier for time series. It is known to be one of the more reliable and simple approaches for time series classification [1].
- **Mr-SEQL** [13] is a linear classifier which uses symbolic features that are extracted from the time series. Mr-SEQL is an interpretable model and has been used by many for explainable AI in Time Series research. Hence, it has been included in this evaluation.

Each dataset was split with the first two-thirds of each of the fatigued and non-fatigued 400m runs being used as the training data and the last one-third used as test data which was done to preserve the time series nature of the problem as instances next to each other are more likely to be similar to each other

<sup>2</sup> [https://scikit-learn.org/stable/modules/generated/sklearn.linear\\_model.RidgeCV.html](https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.RidgeCV.html)



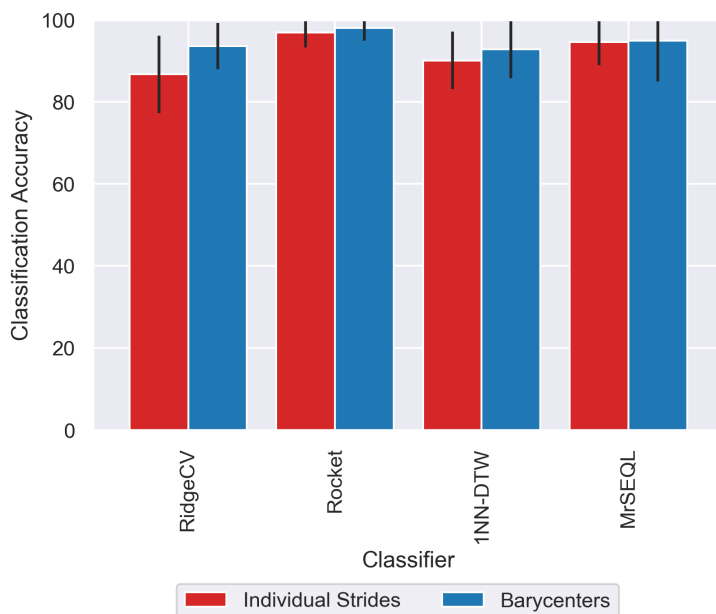


Fig. 3: Average accuracy using individual strides versus barycenters across the 19 subjects for the four selected models with errorbars showing the variation between participants

and hence this technique reduces some bias. On average, across the participants, there were 192 training instances and 98 test instances per participant. The classification models were run on both the individual strides as well as the aggregated barycenter signals. For the barycenter signals, the last four barycenters in the train data were removed to prevent any training data being leaked into the test data.

A summary of the results obtained are shown in Figure 3 with the full set of results presented in the git repository <sup>3</sup>. Overall the Rocket classifier performed the best but all the classifiers were able to produce an average accuracy across the subjects above 0.8. Using the barycenters instead of the individual strides had slightly better performance suggesting that the barycenters could be a good smoothing technique to use.

## 5 Explanation

To provide insights into the impact of fatigue on the runners, we present personalised explanations for selected participants. We aim to investigate if interpretable time series techniques can successfully identify the more discriminative

<sup>3</sup> [https://github.com/bahavathyk/Running\\_Fatigue](https://github.com/bahavathyk/Running_Fatigue)

regions of the time series and if these regions can be explained in a real-world context.

We select four participants and present a personalised explanation for each participant. We assess a suitable explanation technique using a ranking method as detailed in Section 5.1. The participants were selected on the basis of covering a range of scenarios of how fatigue impacts runners. We select three participants who show increased peak accelerations after fatigue which is a commonly seen trend amongst many participants and one participant who does not show this increase.

### 5.1 Strategy for selecting the explanation method

**Explanation Methods.** As we are interested in exploring the difference among fatigued and non-fatigued runners using the raw time series signals, we select local time-based explanation methods that specifically show the critical moments (in time) that differentiate the two classes. This explanation is often presented in the form of a saliency map, highlighting the discriminative areas of the data. It is usually represented as a set of importance weights (from 0 to 100), one weight for each point in the target time series. In this explanation assessment, we evaluate 8 popular methods (Table 1). Of the selected methods, Gradient SHAP [18] and Integrated Gradient [34] are commonly used in explaining deep models for image classification, using the gradients of the trained models. SHAP [17] and LIME [25] are model-agnostic, post-hoc explanation methods that can be used to explain any classifiers. In our evaluation, MrSEQL-LIME, MrSEQL-SHAP, ROCKET-LIME, and ROCKET-SHAP are results of LIME and SHAP explanations for the MrSEQL [13] and ROCKET [6] classifiers. Intrinsic weights from MrSEQL [13] (time-series specific classifier) and Ridge Regression classifier (general classifier) are also added to the explanation methods to be evaluated. A random explanation series (randomly generated importance weights) is used as a sanity check to filter out any ineffective explanation methods.

**Selection Strategy.** We apply the explanation selection strategy presented in [22] to evaluate the selected candidate explanation methods. This strategy aims to estimate the impact of discriminative data areas (identified by an explanation method) by perturbing the data in that area and measuring how a time series classifier (referee classifier) responds to the perturbation. A good explanation method that correctly identifies such discriminative areas will theoretically trigger a more significant drop in accuracy of the referee classifier. This strategy uses a committee of referee classifiers and a combination of different perturbation strategies to reduce possible bias and enhance robustness in the evaluation.

In our assessment, the committee of Referee Classifiers includes ROCKET [6], Mr-SEQL [13], 1NN-DTW [1] and Ridge Regression on datasets perturbed with mean and Gaussian profile from either entire dataset or specific time point [21].

**Results & Selected Explanation** We present the results of our time series explanation evaluation in Table 1. We note that by evaluating and ranking 8

Explanation Method	Rank for Participant ID				Overall Rank	
	<b>07</b>	<b>08</b>	<b>17</b>	<b>18</b>	Average	Number (1-9)
Gradient SHAP	6	3	6	7	5.50	5
Integrated Gradient	3	1	9	8	5.25	4
MrSEQL-LIME	5	5	5	3	4.50	3
ROCKET-LIME	9	6	3	4	5.50	5
<b>MrSEQL-SHAP</b>	<b>1</b>	<b>2</b>	<b>1</b>	<b>1</b>	<b>1.25</b>	<b>1</b>
ROCKET-SHAP	2	4	8	2	4.00	2
MrSEQL-SM	4	7	2	5	4.50	3
RidgeCV-SM	8	9	7	9	8.25	7
Random	7	8	4	6	6.25	6

Table 1: Explanation evaluation ranking results. MrSEQL-SHAP has the best rank and is thus the most informative explanation method for this problem. The rank of each explanation method is computed for each participant dataset, and then averaged across participants.

explanation methods on the 4 participants we analyse in depth (with ids: 07, 08, 17, 18), the most informative explanation method for our data is MrSEQL-SHAP.

## 5.2 Discussion and Feedback from the Explanations

The selected explanation technique, Mr-SEQL-SHAP was used to generate explanations for the four selected participants and these were discussed with our domain expert to gain insights into how fatigue affected these runners. TimeXplain was used to apply Shap to the time series. The aggregated barycenter strides were used for the evaluations due to improved classification performance. Indeed the barycentres may also help to improve the clarity of the explanations. Figure 4 shows saliency maps showing a Mr-SEQL-SHAP explanation for both the individual strides and the aggregated barycenter strides for participant 8 and it can be seen that the barycenters aid in visualising the saliency and potentially improving the stability of the predictions. A reason for this could be that many of these interpretation techniques work well when the data is aligned well. Although the process of segmentation makes the nature of the initial data reasonably well aligned, the barycenter further aids in this overall alignment of the time series.

Figure 5 on the left shows a mapping of all the fatigued and non-fatigued strides against each other and a saliency map of a sample fatigued stride produced from Mr-SEQL-SHAP for each of the four participants. The red highlighted regions in the saliency map suggest that region contributed more towards the classification decision.

The most common effect of fatigue was an increase in the peak acceleration point which corresponds to the initial contact where the foot strikes the




Fig. 4: (a) Mr-SEQL-SHAP explanation on the individual strides (Fatigued Class), (b) Mr-SEQL-SHAP explanation on the rolling barycenters calculated from the individual strides (Fatigued class) for Participant 8.

ground. This suggests that most people lose control during the loading stage. This supports previous research which suggests that the natural damping of runners reduces with fatigue [38] which would affect the loading stage or the weight absorption stage of running which is observed around the peaks in our data.

The insights from the highlighted regions (Figure 5) of each participant can be summarised as follows:

**Participant 8 and 18:** This is the most straightforward case and perhaps the most common in the dataset. These participants have a fairly symmetric gait - the left and right peaks are similar. Perhaps there is a slightly higher peak acceleration on their left (non-dominant) leg. However, there is a clear asymmetry in the fatigued state and this is highlighted in dark red in the saliency maps. For participant 8 the saliency map highlights the second peak but to a lesser extent. So for both participants the feedback is that fatigue shows up particularly in the non-dominant leg resulting in loss of control during the loading stage.

**Participant 17:** This participant is similar to 8 and 18 except that there is already evidence of asymmetry in the non-fatigued strides. The asymmetry shows up as higher peak acceleration on the left leg. This asymmetry is accentuated when fatigued. Both peaks are highlighted but the saliency map gives more weight to the right leg in this case.

**Participant 7:** This is the most interesting participant in some respects. We can see in Figure 5e that there is no difference in peak impact between the fatigued and non-fatigued states. So there is no evidence of loss of control during the loading stage. Nevertheless the classifiers are able to distinguish between the two classes with very high accuracy. The saliency map in Figure 5f indicates that the discriminative region is later in the stride; it is during the unloading phase

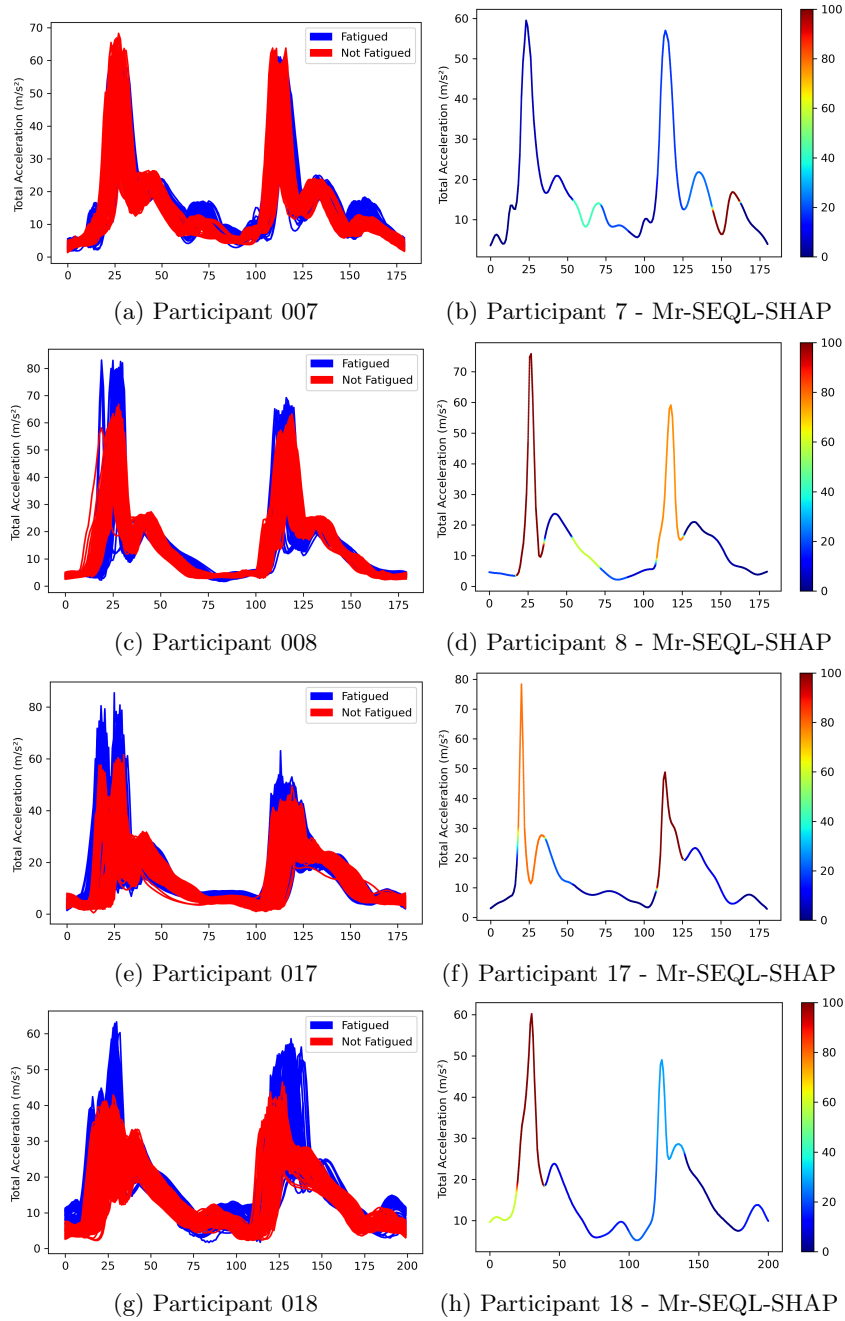


Fig. 5: Left: A map of the fatigued versus not-fatigued instances for the selected four participants. Right: Mr-SEQL-SHAP generated saliency maps for a single sample of a fatigued stride

where the leg is preparing to go into flight. This is interesting as although the participant was able to maintain their impact acceleration which is known as an indicator of fatigue, they are still modifying their overall running gait in a way that is less obvious to an observer.

Overall the discriminatory regions that were identified by the Mr-SEQL-SHAP technique supported the overlay plots in Figure 5 as the regions which were clearly different in the overlay plots were highlighted in the explanations. However, the saliency plots also provide more information as they were able to highlight the particular limb that was contributing more to the classification task. Furthermore, it should be noted that there were some minor discrepancies in the segmentation due to the automated nature and this can be seen in the bi-modal behaviour in participant 17 (Figure 5e) where there are two apparent points that were interchangeably selected during segmentation. Despite this, the models were easily able to classify accurately and select meaningful regions as discriminatory. This adaptability is important due to the personalised nature of running, where a one-fit-all solution is not feasible.

## 6 Conclusions and Future Work

In this paper, we present the protocol for collecting and processing a dataset of wearable sensor data from runners in their normal and fatigued states. We then evaluate the utility of using a barycenter averaging strategy as a smoothing and aggregation strategy to improve overall classification and explanation. Finally, we identify a suitable explanation technique and apply this to our data to gain insights into how fatigue impacts runners.

The use of barycenters to aggregate the time series may help improve the classification performance as well as improve the visual aspect of the explanation. A variety of explanation techniques were also investigated and Mr-SEQL-SHAP came out to be the most effective way to explain this particular dataset. Combining the explanations generated from the Mr-SEQL-SHAP method with domain specific knowledge, some insights into how fatigue impacts runners were made. It was observed that many runners lost control during the loading stage, which is around the peak impact point. However, this was not the case for every participant and hence shows the importance of having personalised feedback systems. These feedback systems could be very helpful to runners in providing personalised feedback on the changes they are making to their overall running kinematics.

The results showed that Mr-SEQL with SHAP was able to highlight domain meaningful discriminative regions of the time series for the magnitude of acceleration data. However, the sensors have multivariate data, and looking at the other signals may be useful in providing further insights into the data. Furthermore, there is limited research that looks at explanations for multivariate time series data and hence an interesting avenue to investigate would be to extend this work to a multivariate case.

## Acknowledgements

This work has emanated from research conducted with the financial support of Science Foundation Ireland under the Grant number 18/CRT/6183.

## References

1. Bagnall, A., Lines, J., Bostrom, A., Large, J., Keogh, E.: The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery* **31**(3), 606–660 (2017). <https://doi.org/10.1007/s10618-016-0483-9>
2. Bagnall, A., Lines, J., Hills, J., Bostrom, A.: Time-Series Classification with COTE: The Collective of Transformation-Based Ensembles. *IEEE Transactions on Knowledge and Data Engineering* **27**(9), 2522–2535 (2015). <https://doi.org/10.1109/TKDE.2015.2416723>
3. Borgia, B., Dufek, J.S., Silvernail, J.F., Radzak, K.N.: The effect of fatigue on running mechanics in older and younger runners. *Gait and Posture* **97**(July), 86–93 (2022). <https://doi.org/10.1016/j.gaitpost.2022.07.249>, <https://doi.org/10.1016/j.gaitpost.2022.07.249>
4. Buckley, C., O’Reilly, M., Whelan, D., Farrell, A.V., Clark, L., Longo, V., Gilchrist, M., Caulfield, B.: Binary classification of running fatigue using a single inertial measurement unit. In: 2017 IEEE 14th International Conference on Wearable and Implantable Body Sensor Networks (BSN). pp. 197–201 (2017). <https://doi.org/10.1109/BSN.2017.7936040>
5. Cuturi, M., Blondel, M.: Soft-DTW: A differentiable loss function for time-series. 34th International Conference on Machine Learning, ICML 2017 **2**, 1483–1505 (2017)
6. Dempster, A., Petitjean, F., Webb, G.I.: ROCKET: exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery* **34**(5), 1454–1495 (2020). <https://doi.org/10.1007/s10618-020-00701-z>
7. Grabocka, J., Schilling, N., Wistuba, M., Schmidt-Thieme, L.: Learning time-series shapelets. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* pp. 392–401 (2014). <https://doi.org/10.1145/2623330.2623613>
8. Hunter, I., Smith, Æ.G.A.: Preferred and optimal stride frequency , stiffness and economy : changes with fatigue during a 1-h high-intensity run pp. 653–661 (2007). <https://doi.org/10.1007/s00421-007-0456-1>
9. Kathirgamanathan, B., Buckley, C., Caulfield, B., Cunningham, P.: Feature Subset Selection for Detecting Fatigue in Runners Using Time Series Sensor Data. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **13363 LNCS**, 541–552 (2022). [https://doi.org/10.1007/978-3-031-09037-0\\_44](https://doi.org/10.1007/978-3-031-09037-0_44)
10. Kathirgamanathan, B., Caulfield, B., Cunningham, P.: Multivariate Time Series data of Fatigued and Non- Fatigued Running from Inertial Measurement Units (Jun 2023). <https://doi.org/10.5281/zenodo.7997851>, <https://doi.org/10.5281/zenodo.7997851>

11. Kathirgamanathan, B., Davenport, J., Caulfield, B., Cunningham, P.: Using Barycenters as Aggregate Representations of Repetition-Based Time-Series Exercise Data. *Communications in Computer and Information Science* **1571 CCIS**, 178–188 (2022). [https://doi.org/10.1007/978-3-031-02044-5\\_15](https://doi.org/10.1007/978-3-031-02044-5_15)
12. Kenny, E.M., Delaney, E.D., Greene, D., Keane, M.T.: Post-hoc explanation options for xai in deep learning: The insight centre for data analytics perspective. In: Del Bimbo, A., Cucchiara, R., Sclaroff, S., Farinella, G.M., Mei, T., Bertini, M., Escalante, H.J., Vezzani, R. (eds.) *Pattern Recognition. ICPR International Workshops and Challenges*. pp. 20–34. Springer International Publishing, Cham (2021)
13. Le Nguyen, T., Gsponer, S., Ilie, I., O’Reilly, M., Ifrim, G.: Interpretable time series classification using linear models and multi-resolution multi-domain symbolic representations. *Data Mining and Knowledge Discovery* **33**(4), 1183–1222 (2019). <https://doi.org/10.1007/s10618-019-00633-3>, <https://doi.org/10.1007/s10618-019-00633-3>
14. Lin, J., Keogh, E., Wei, L., Lonardi, S.: Experiencing SAX: A novel symbolic representation of time series. *Data Mining and Knowledge Discovery* **15**(2), 107–144 (2007). <https://doi.org/10.1007/s10618-007-0064-z>
15. Lines, J., Davis, L.M., Hills, J., Bagnall, A.: A shapelet transform for time series classification. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* pp. 289–297 (2012). <https://doi.org/10.1145/2339530.2339579>
16. Lines, J., Taylor, S., Bagnall, A.: Time Series Classification with HIVE-COTE. *ACM Transactions on Knowledge Discovery from Data* **12**(5), 1–35 (2018). <https://doi.org/10.1145/3182382>
17. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. p. 4768–4777. NIPS’17, Curran Associates Inc., Red Hook, NY, USA (2017)
18. Lundberg, S.M., Nair, B., Vavilala, M.S., Horibe, M., Eisses, M.J., Adams, T., Liston, D.E., Low, D.K.W., Newman, S.F., Kim, J., others: Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature biomedical engineering* **2**(10), 749–760 (2018)
19. Léger, L.A., Mercier, D., Gadoury, C., Lambert, J.: The multistage 20 metre shuttle run test for aerobic fitness. *Journal of Sports Sciences* **6**(2), 93–101 (1988). <https://doi.org/10.1080/02640418808729800>, <https://doi.org/10.1080/02640418808729800>, PMID: 3184250
20. Mizrahi, J.: Fatigue Effect on Running Kinematics. *Human Movement Science* **19**, 139–151 (2000)
21. Mujkanovic, F., Doskoč, V., Schirneck, M., Schäfer, P., Friedrich, T.: timeXplain – A Framework for Explaining the Predictions of Time Series Classifiers pp. 1–17 (2020), <http://arxiv.org/abs/2007.07606>
22. Nguyen, T.T., Nguyen, T.L., Ifrim, G.: Amee: A robust framework for explanation evaluation in time series classification. *arXiv preprint arXiv:2306.05501* (2023)
23. O’Reilly, M.A., Johnston, W., Buckley, C., Whelan, D., Caulfield, B.: The influence of feature selection methods on exercise classification with inertial measurement units. *2017 IEEE 14th International Conference on Wearable and Implantable Body Sensor Networks, BSN 2017* pp. 193–196 (2017). <https://doi.org/10.1109/BSN.2017.7936039>



24. Petitjean, F., Ketterlin, A., Gançarski, P.: A global averaging method for dynamic time warping, with applications to clustering. *Pattern Recognition* **44**(3), 678–693 (2011). <https://doi.org/10.1016/j.patcog.2010.09.013>
25. 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*. p. 1135–1144. KDD '16, Association for Computing Machinery, New York, NY, USA (2016). <https://doi.org/10.1145/2939672.2939778>, <https://doi.org/10.1145/2939672.2939778>
26. Ruiz, A.P., Flynn, M., Bagnall, A.: *Benchmarking Multivariate Time Series Classification Algorithms*. arXiv (2020)
27. Sakoe, H., Chiba, S.: Dynamic Programming Algorithm Optimization for Spoken Word Recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing* **26**(1), 43–49 (1978). <https://doi.org/10.1109/TASSP.1978.1163055>
28. Schäfer, P.: The BOSS is concerned with time series classification in the presence of noise. *Data Mining and Knowledge Discovery* **29**(6), 1505–1530 (2015). <https://doi.org/10.1007/s10618-014-0377-7>, <http://dx.doi.org/10.1007/s10618-014-0377-7>
29. Schäfer, P.: Scalable time series classification. *Data Mining and Knowledge Discovery* **30**(5), 1273–1298 (2016). <https://doi.org/10.1007/s10618-015-0441-y>
30. Schäfer, P., Höggqvist, M.: SFA: A symbolic fourier approximation and index for similarity search in high dimensional datasets. *ACM International Conference Proceeding Series* pp. 516–527 (2012). <https://doi.org/10.1145/2247596.2247656>
31. Schäfer, P., Leser, U.: *Multivariate time series classification with Weasel+MUSE*. arXiv (2017)
32. Shi, K., Qin, H., Sima, C., Li, S., Shen, L., Ma, Q.: Dynamic Barycenter Averaging Kernel in RBF Networks for Time Series Classification. *IEEE Access* **7**, 47564–47576 (2019). <https://doi.org/10.1109/ACCESS.2019.2910017>
33. Strohrmann, C., Harms, H., Kappeler-setz, C., Tr, G.: Monitoring Kinematic Changes With Fatigue in Running Using Body-Worn Sensors. *IEEE Transactions on Information Technology In Biomedicine* **16**(5), 983–990 (2012)
34. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks. In: *International conference on machine learning*. pp. 3319–3328 (2017)
35. Tavenard, R., Faouzi, J., Vandewiele, G., Divo, F., Androz, G., Holtz, C., Payne, M., Yurchak, R., Rußwurm, M., Kolar, K., Woods, E.: *tslearn: A machine learning toolkit dedicated to time-series data* (2017), <https://github.com/rtavenar/tslearn>
36. Tonoli, D.C., Cumps, E., Aerts, I., Verhagen, E., Meeusen, R.: Incidence, risk factors and prevention of running related injuries in long-distance running: a systematic review. *Sport & Geneeskunde* **43**(5) (2010)
37. Whelan, D.F., O'Reilly, M.A., Ward, T.E., Delahunt, E., Caulfield, B.: Technology in rehabilitation: Comparing personalised and global classification methodologies in evaluating the squat exercise with wearable IMUs. *Methods of Information in Medicine* **56**(5), 361–369 (2017). <https://doi.org/10.3414/ME16-01-0141>
38. Willems, T., De Clercq, D., Delbaere, K., Vanderstraeten, G., De Cock, A., Witvrouw, E.: A prospective study of gait related risk factors for exercise-related lower leg pain. *Gait & Posture* **23**(1), 91–98 (2006). <https://doi.org/https://doi.org/10.1016/j.gaitpost.2004.12.004>, <https://www.sciencedirect.com/science/article/pii/S0966636204002747>

39. Ye, L., Keogh, E.: Time series shapelets: A new primitive for data mining. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining pp. 947–955 (2009). <https://doi.org/10.1145/1557019.1557122>
40. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)