TS-MULE

Local Interpretable Model-Agnostic Explanations For Time Series Forecast Models

Udo Schlegel, Duy Lam Vo, Daniel A. Keim, and Daniel Seebacher
University of Konstanz

Attributions in explainable AI

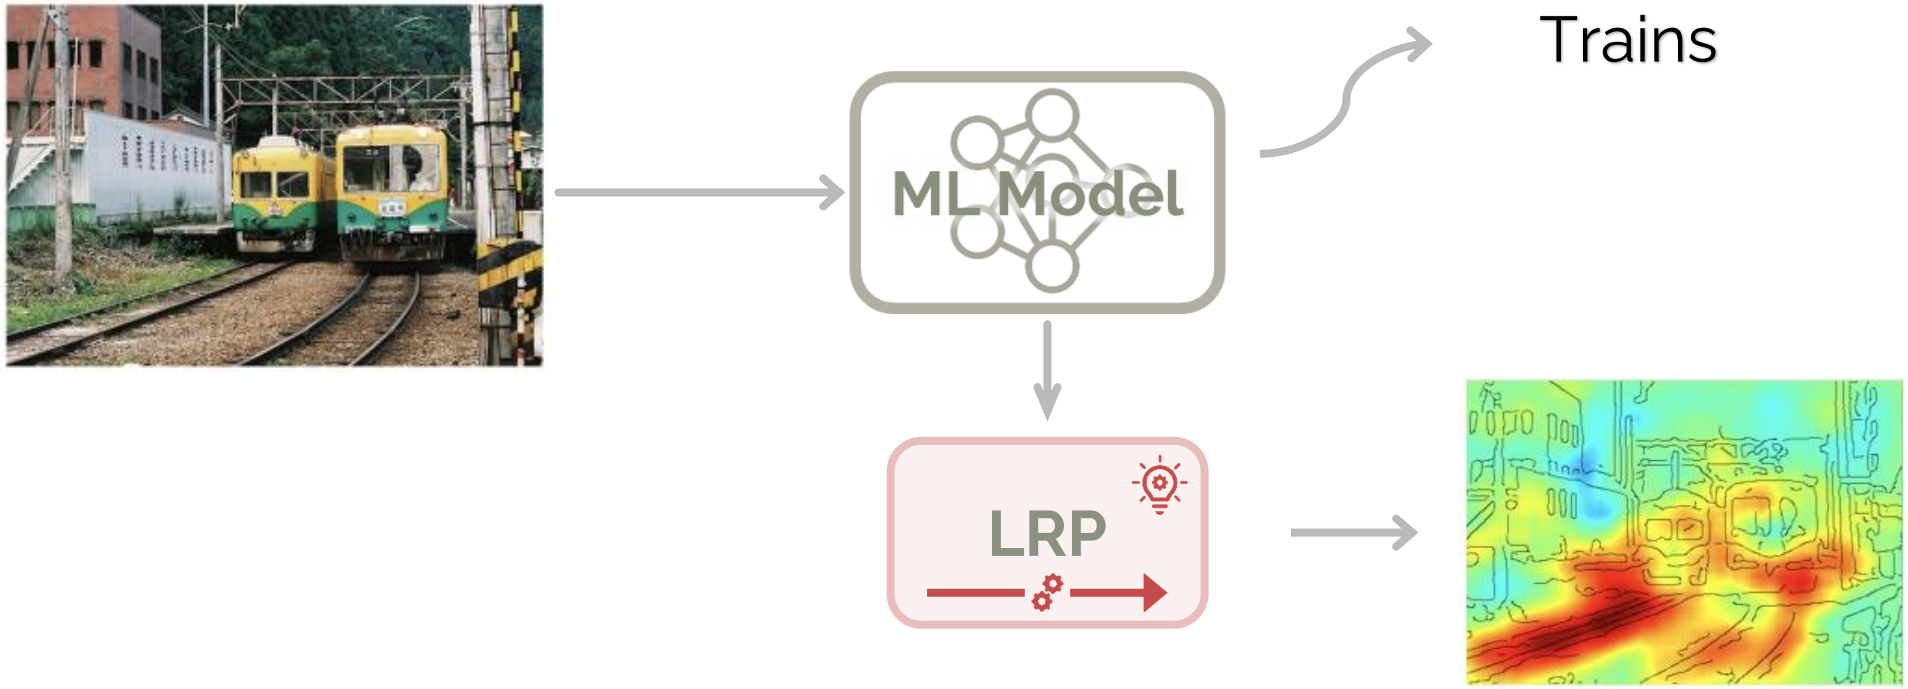

http://xai.unist.ac.kr/static/img/event/ICCV_2019_VXAI_Samek_Talk.pdf

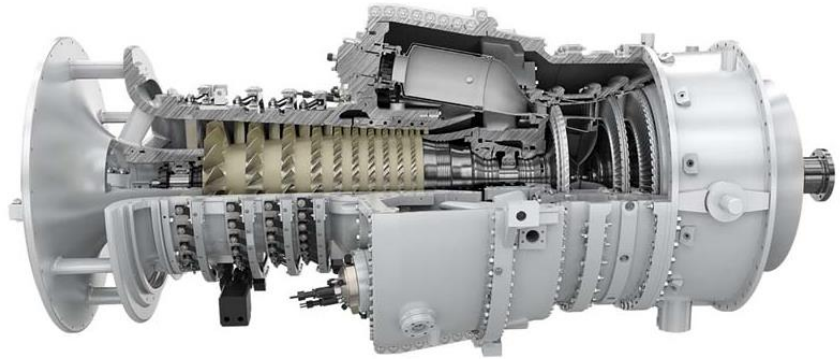

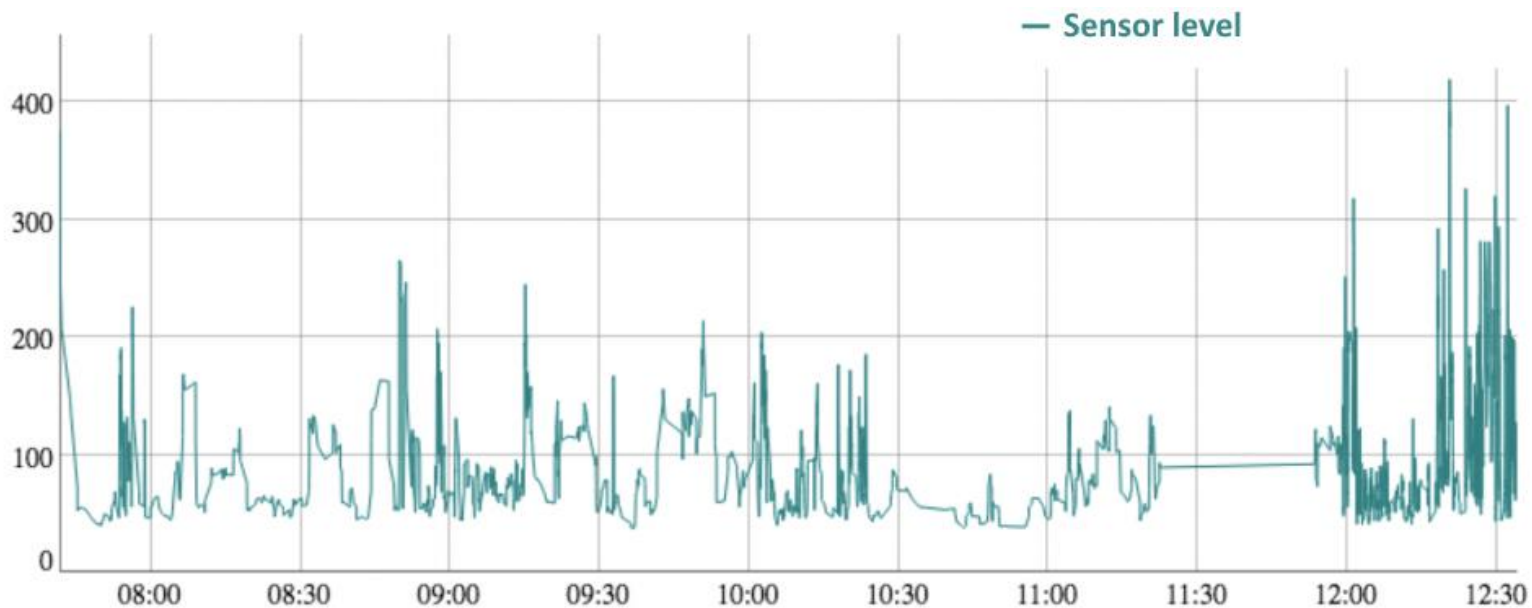
— Sensor level
400
300
200
100
0
08:00
08:30
09:00
09:30
10:00
10:30
11:00
11:30
12:00
12:30

The LIME approach

1. Segment data

Interpretable Components

2. Create masks and get predictions

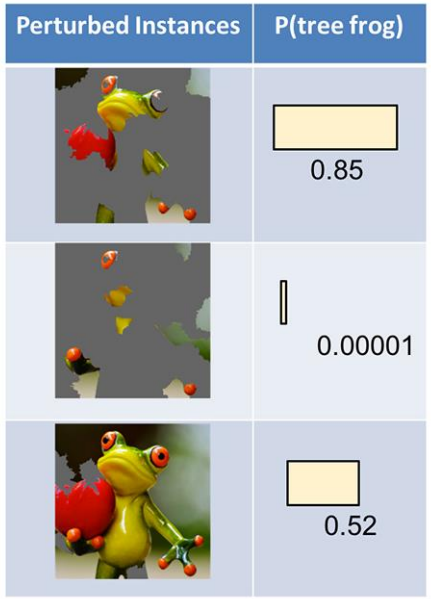

3. Train interpretable model on mask and predictions

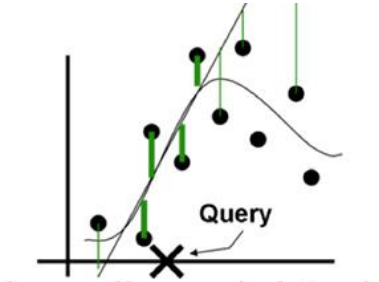

Locally weighted regression

4. Get most important component for prediction

Explanation

https://theblue.ai/wp-content/uploads/2019/04/figure4.jpg

The LIME approach for time series

1. Segment data

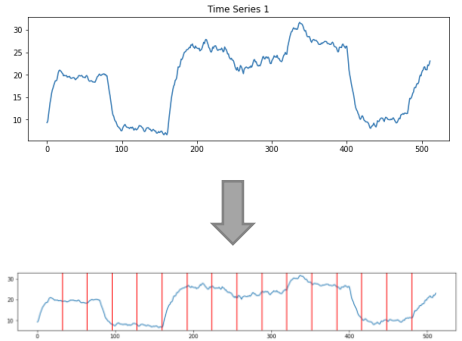

2. Create masks and get predictions

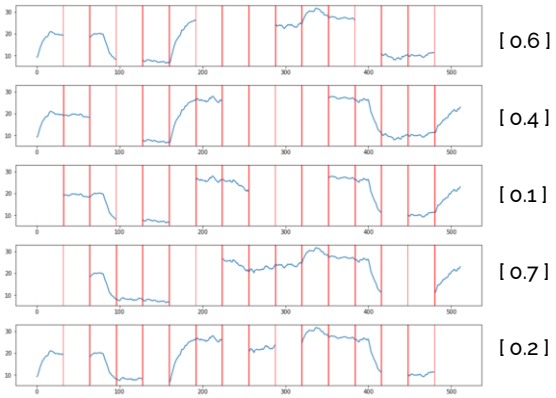

3. Train interpretable model on mask and predictions

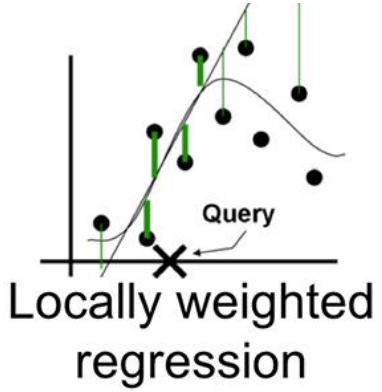

4. Get most important component for prediction

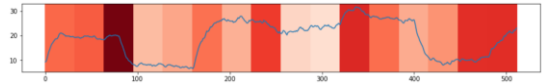

https://theblue.ai/wp-content/uploads/2019/04/figure4.jpg

Time series segmentation

- What are meaningful segments for time series?

- Uniform?

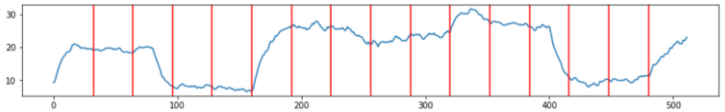

- Does not look very meaningful

Segmentations for time series

Uniform segmentation

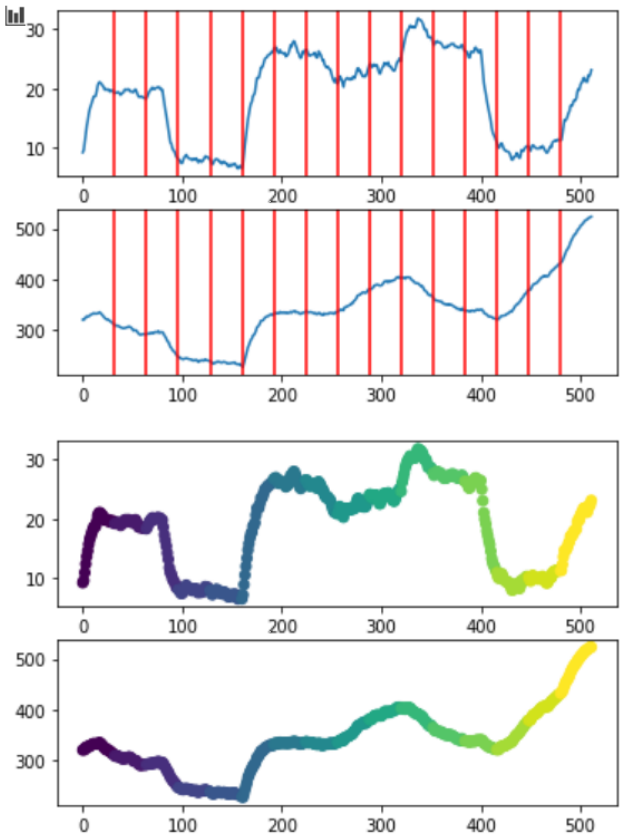

Exponential segmentation

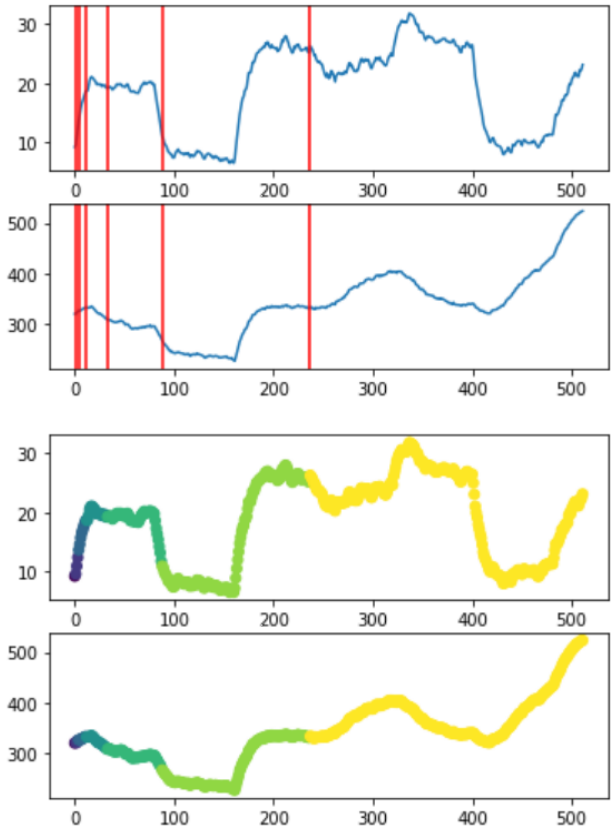

Segmentations for time series

Uniform segmentation

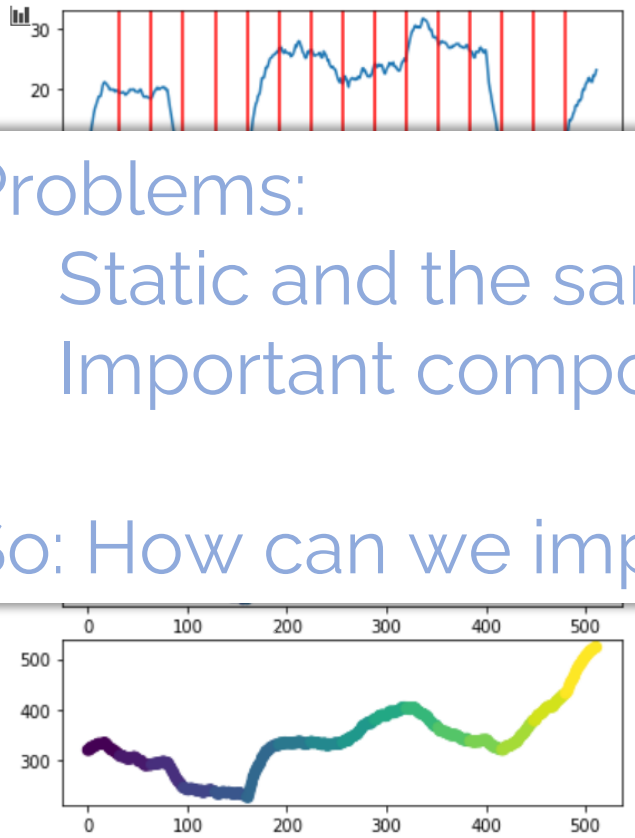

Exponential segmentation

Problems:

- Static and the same for every feature
- Important components can be split into two segments

So: How can we improve these segments?

Matrix Profile

a vector that stores the z-normalized Euclidean distance between any subsequence within a time series and its nearest neighbour

Pairwise Euclidean Distance

0	1	3	2	9	1	14	15	1	2	2	10	7

#DistanceProfile

https://stumpy.readthedocs.io/en/latest/Tutorial_The_Matrix_Profile.html

Matrix Profile

a vector that stores the z-normalized Euclidean distance between any subsequence within a time series and its nearest neighbour

Pairwise Euclidean Distance

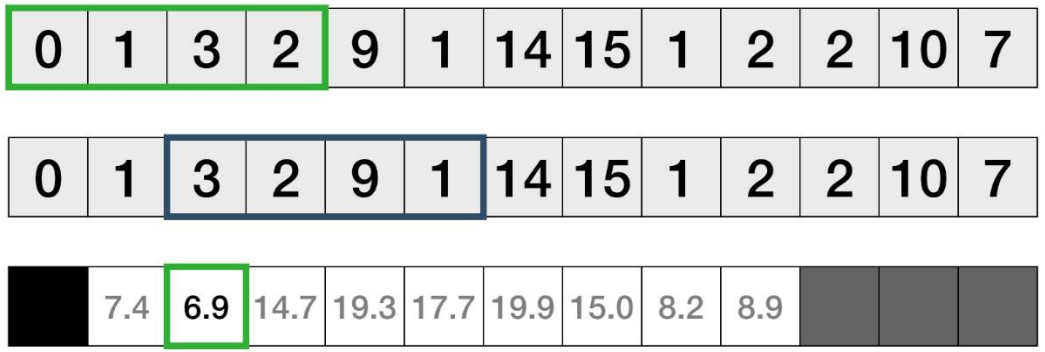

https://stumpy.readthedocs.io/en/latest/Tutorial_The_Matrix_Profile.html

Matrix Profile

a vector that stores the z-normalized Euclidean distance between any subsequence within a time series and its nearest neighbour

Pairwise Euclidean Distance

0	1	3	2	9	1	14	15	1	2	2	10	7
0	1	3	2	9	1	14	15	1	2	2	10	7
	7.4	6.9	14.7	19.3	17.7	19.9	15.0	8.2	8.9			

#BestMatch

Distance Matrix

	*	6.9	*	*	*	*	*	*	*			
*		*	*	*	*	*	*	1.4	*			
*	*		*	*	*	*	*	*	6.2			
*	7.9	*		*	*	*	*	*	*			
*	*	*	*		*	*	*	*	11.4			
*	*	13.6	*	*		*	*	*	*			
*	*	*	*	*	*		14.1	*	*			
*	*	14.0	*	*	*	*		*	*			
*	1.4	*	*	*	*	*	*		*			
*	*	6.2	*	*	*	*	*	*				

https://stumpy.readthedocs.io/en/latest/Tutorial_The_Matrix_Profile.html

Matrix Profile

a vector that stores the z-normalized Euclidean distance between any subsequence within a time series and its nearest neighbour

Distance Matrix

■	*	6.9	*	*	*	*	*	*	*			
*	■	*	*	*	*	*	*	1.4	*			
*	*	■	*	*	*	*	*	*	6.2			
*	7.9	*	■	*	*	*	*	*	*			
*	*	*	*	■	*	*	*	*	11.4			
*	*	13.6	*	*	■	*	*	*	*			
*	*	*	*	*	*	■	14.1	*	*			
*	*	14.0	*	*	*	*	■	*	*			
*	1.4	*	*	*	*	*	*	■	*			
*	*	6.2	*	*	*	*	*	*	■			

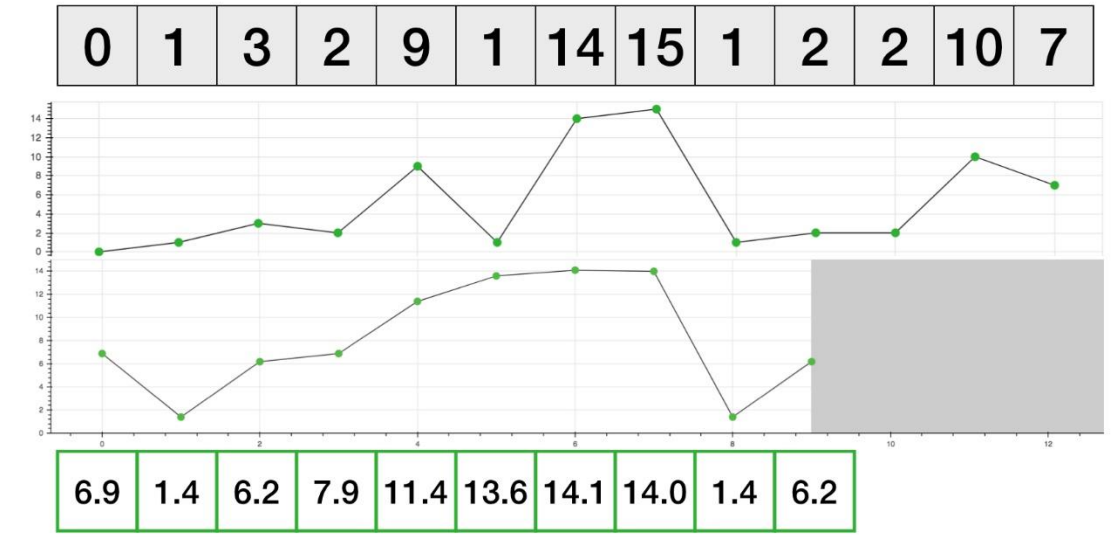

#MatrixProfileAnnotation

https://stumpy.readthedocs.io/en/latest/Tutorial_The_Matrix_Profile.html

Proposed slopes segmentation

Using the matrix profile

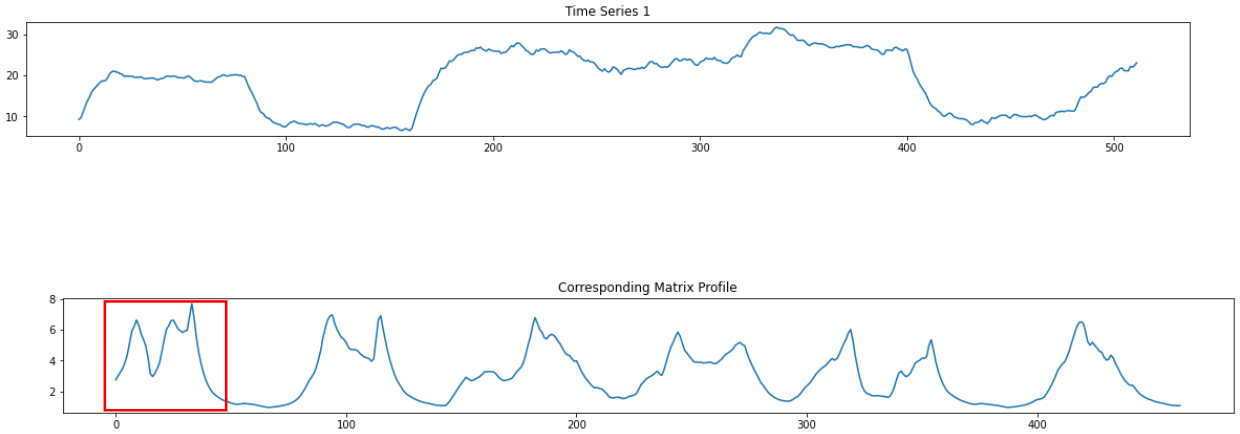

Take slopes of matrix profile
-> largest jump leads to change in nearest neighbours

Slopes segmentation

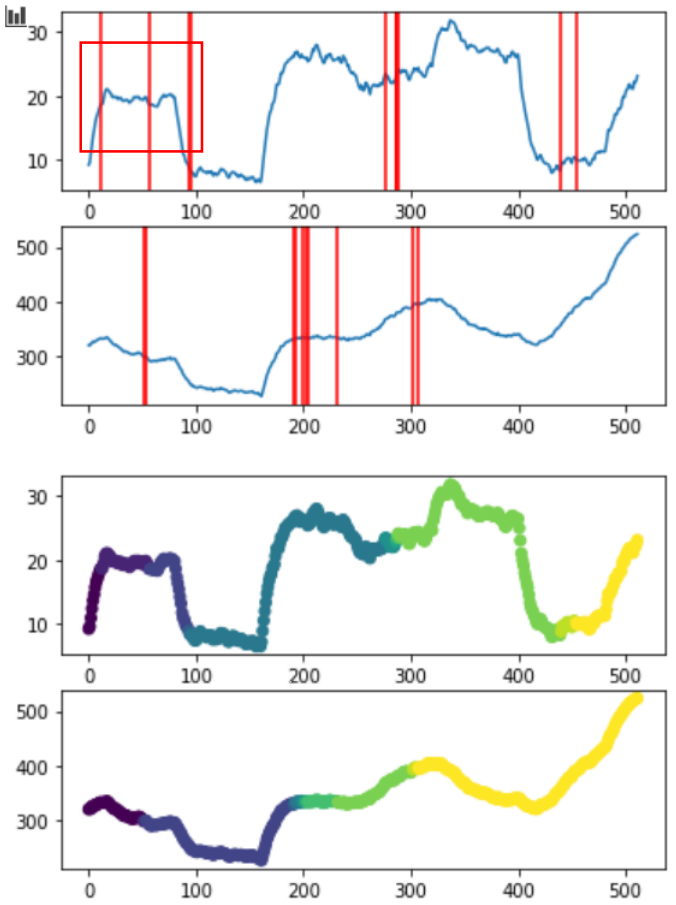

Proposed bins segmentation

Using the matrix profile

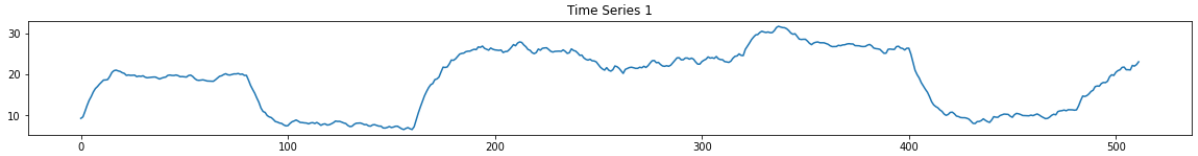

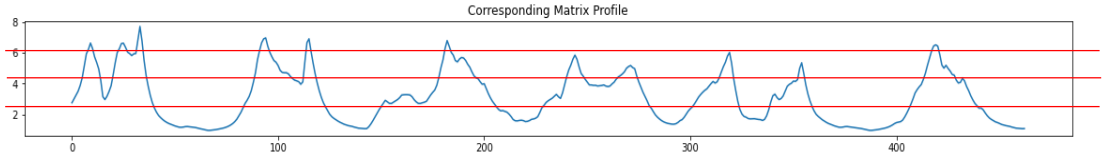

Take horizontal bins and assign segments based on the corresponding bin in the matrix profile

Bins-max segmentation

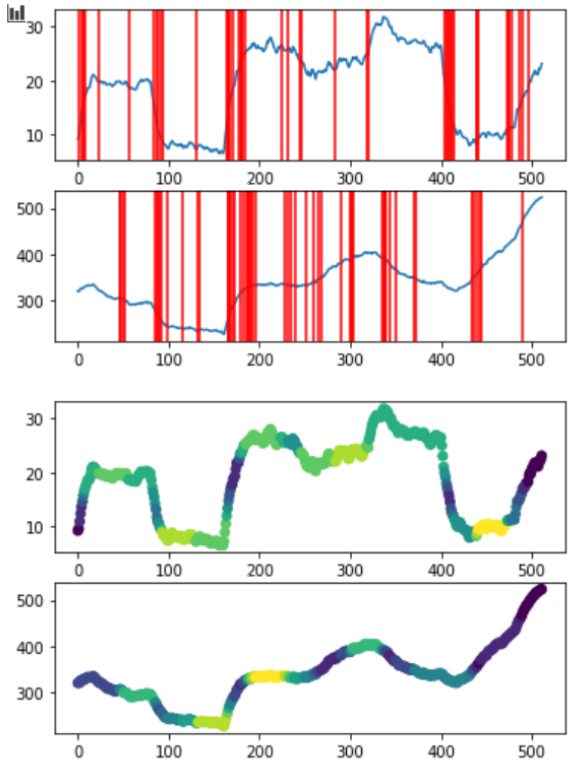

Bins-min segmentation

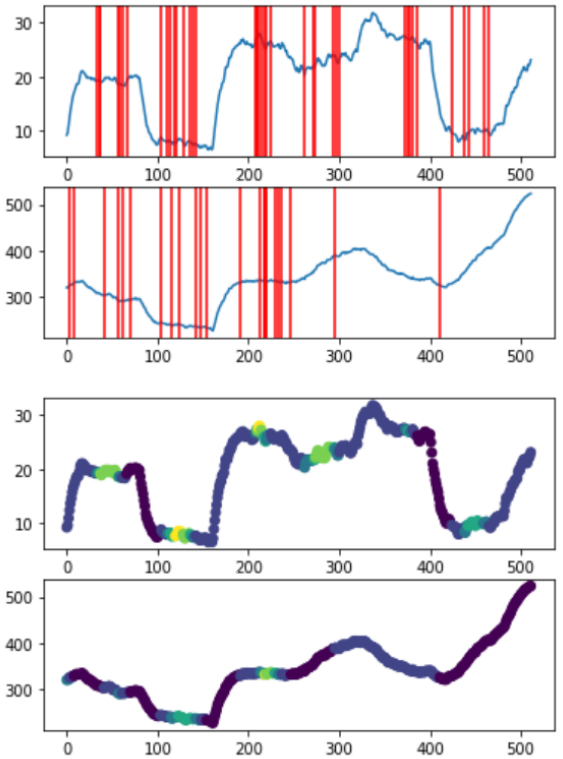

Proposed SAX segmentation

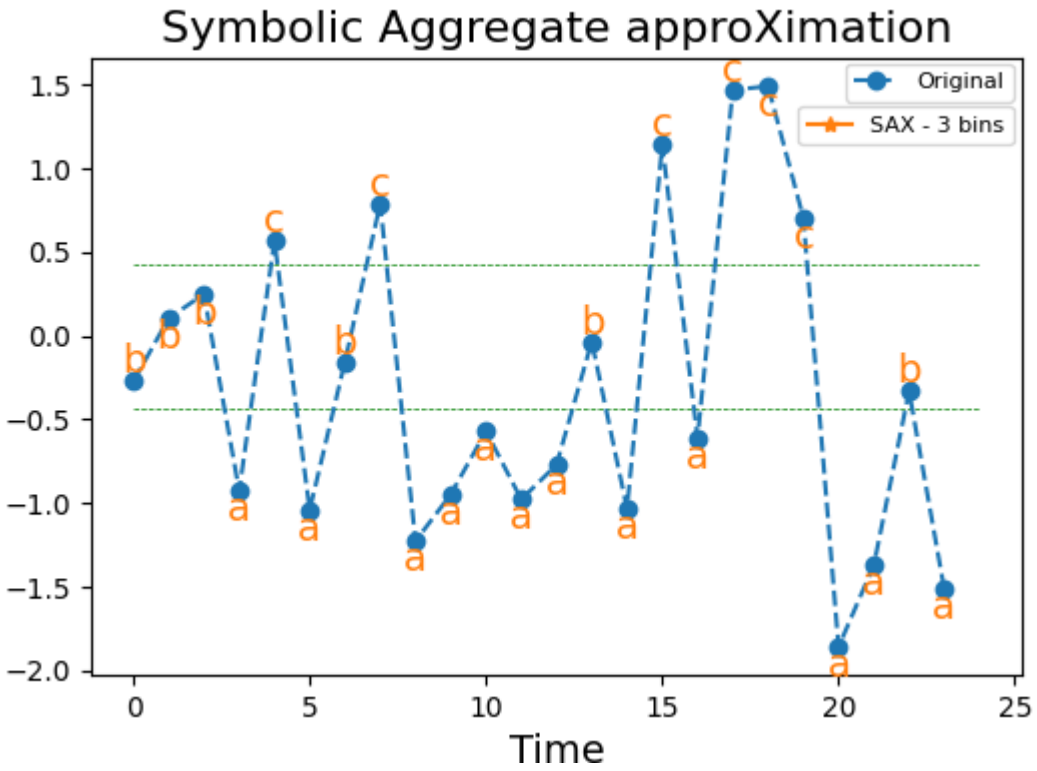

https://pyts.readthedocs.io/en/stable/_images/sphx_glr_plot_sax_001.png

Using the SAX transformation

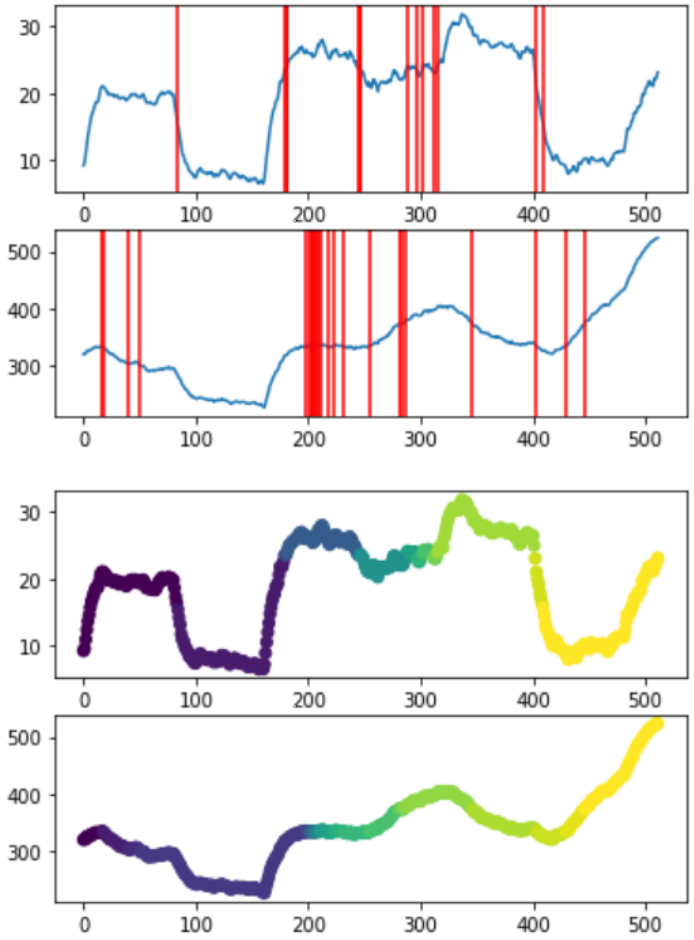

Comparison

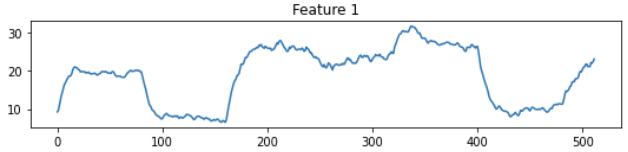
Feature 1

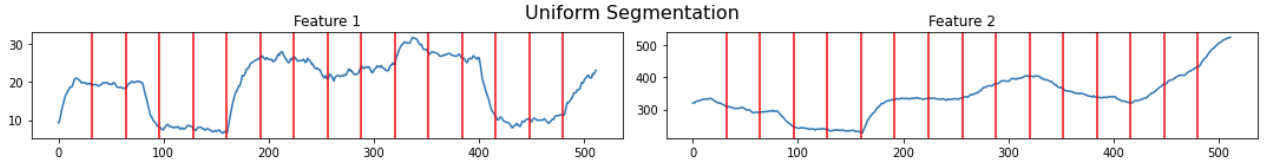
Uniform Segmentation
Feature 1
Feature 2

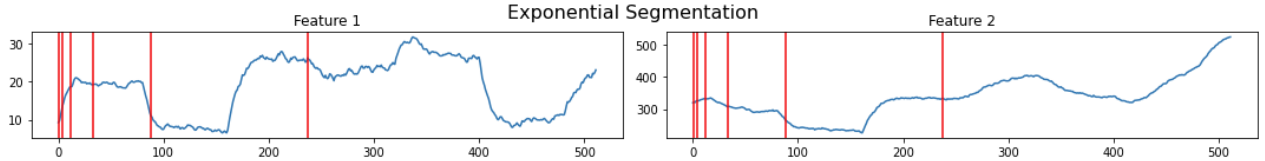
Exponential Segmentation
Feature 1
Feature 2

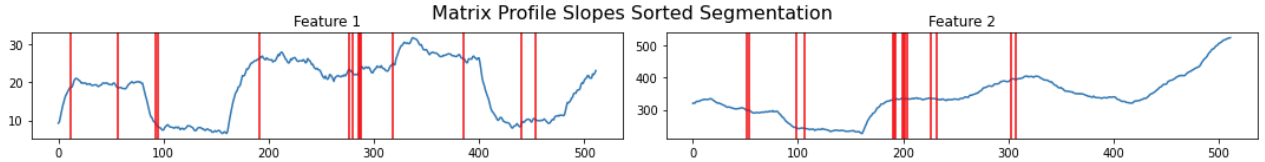
Matrix Profile Slopes Sorted Segmentation
Feature 1
Feature 2

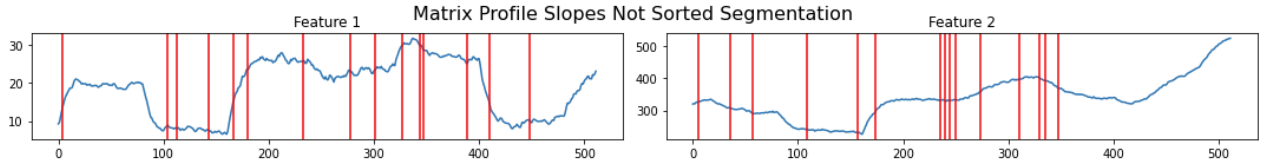
Matrix Profile Slopes Not Sorted Segmentation
Feature 1
Feature 2

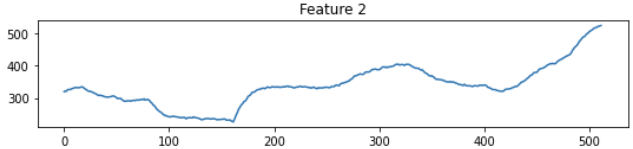
Feature 2

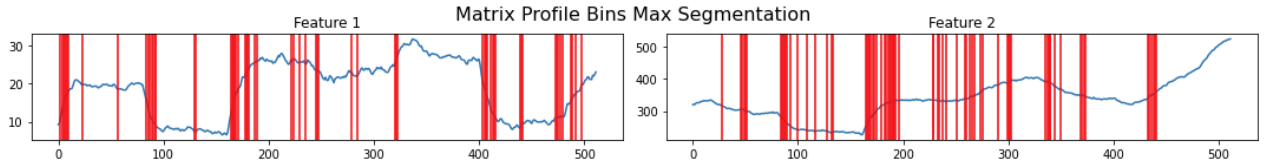
Matrix Profile Bins Max Segmentation
Feature 1
Feature 2

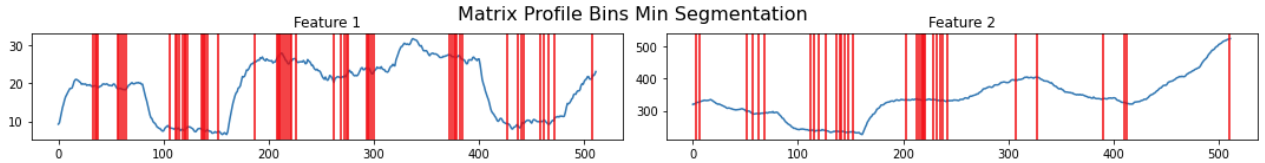
Matrix Profile Bins Min Segmentation
Feature 1
Feature 2

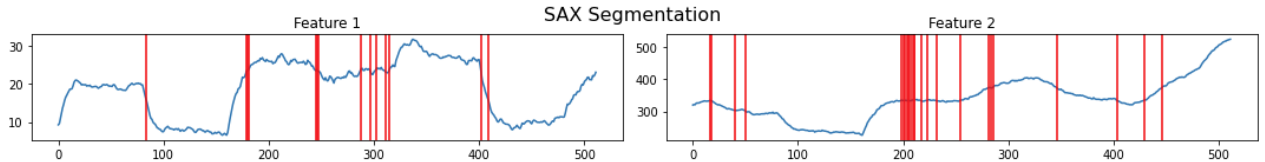
SAX Segmentation
Feature 1
Feature 2

Evaluation through fidelity perturbation

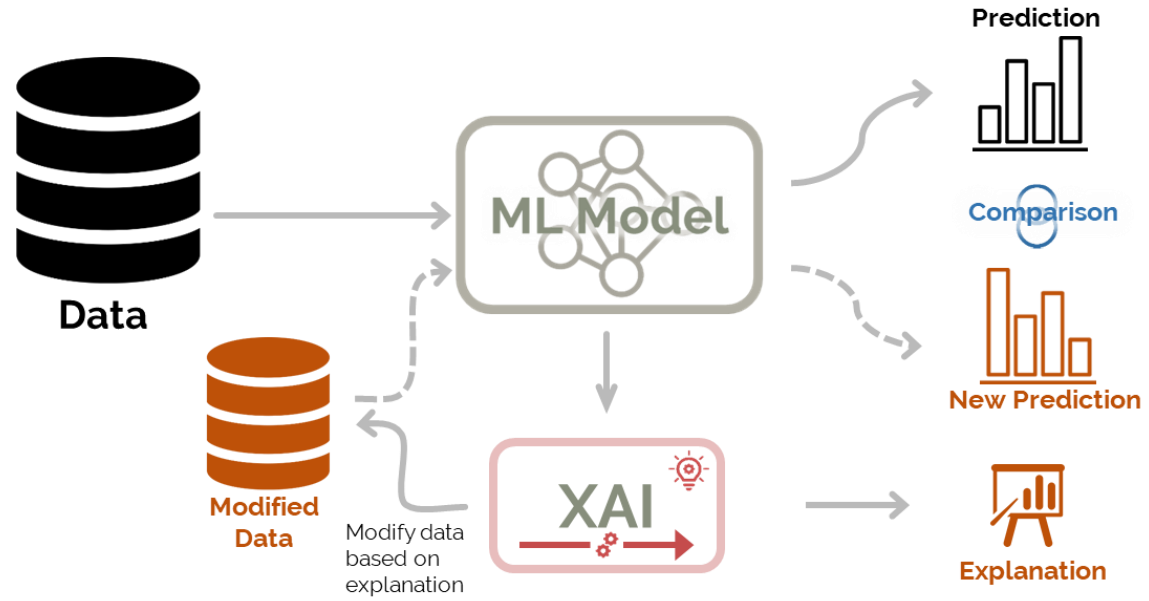

- Assumption:
  - Change data according to attribution
    => get worse accuracy of model for changed data

Preliminary results

Zero	Beijing Air Quality 2.5	CNN	DNN	RNN	Beijing Air Quality Multi Site	CNN	DNN	RNN	Metro Interstate Traffic	CNN	DNN	RNN
Uniform		2.31	4.24	2.32		1.50	9.00	7.67		2.43	0.22	6.55
Exponential		0.56	1.12	1.41		0.62	0.16	11.52		0.55	0.01	0.62
Slopes		1.31	2.11	1.95		1.30	6.76	3.97		3.39	0.18	9.29
Bins Min		0.35	3.43	3.60		0.41	10.46	5.71		1.25	0.40	7.38
Bins Max		1.69	1.22	2.38		1.52	1.68	2.67		1.44	0.44	2.68
SAX		1.24	2.58	2.23		1.10	8.00	4.15		1.55	1.16	7.34

Input length of dataset

Take accuracy change of attributions and scale by random perturbation accuracy change
=> Larger than 1 shows working explanations

Conclusion

- Improved segmentations improve explanations
- Different architectures work better with different segmentations
- Evaluate parameters (e.g., window size)
- Improve presented algorithms to better handle close splits to get less segments

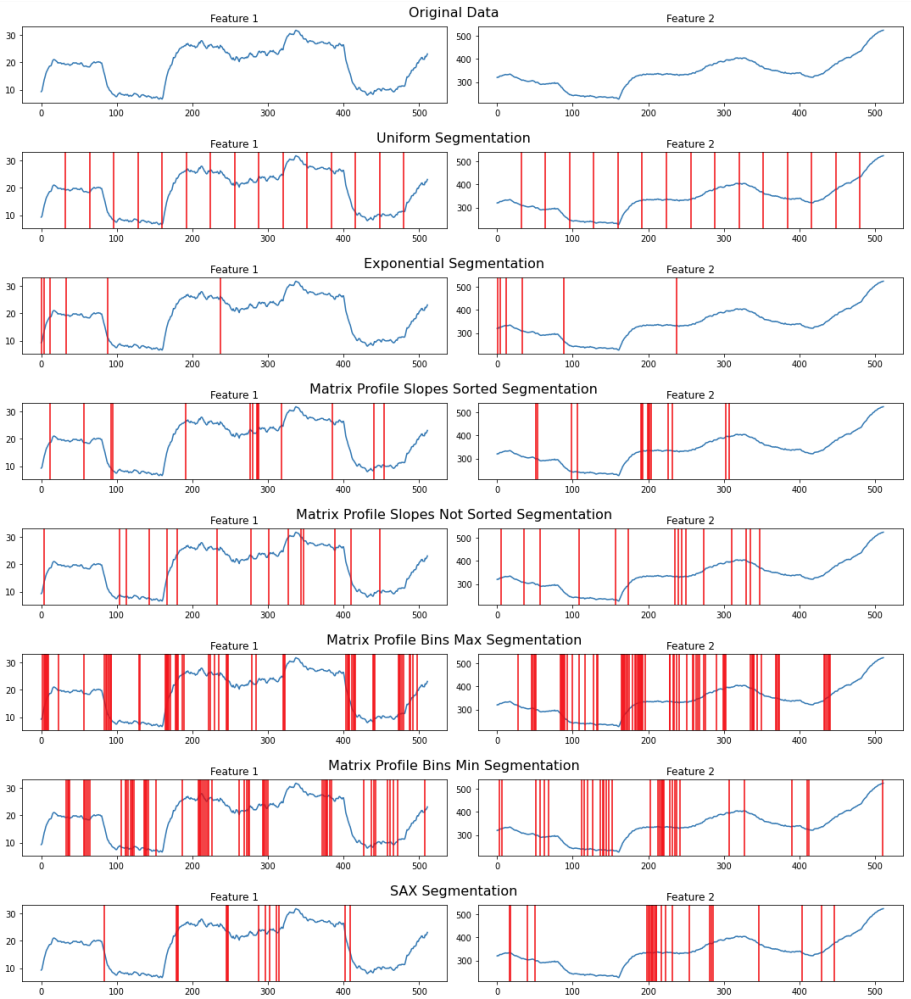

Source Code can be found at:
https://github.com/dbvis-ukon/ts-mule