
UKRI
Science and Technology Facilities Council
ISIS Neutron and Muon Source
Welcome

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Machine Learning For Anomaly Detection In Continuous Signals

Alex Saoulis, Accelerator Controls Graduate

alex.saoulis@stfc.ac.uk

Introduction

1 Data Preparation

Generating feature vectors from raw time series data.

2 Data Labelling

Look at a general method for labelling the time series data using clustering techniques.

3 Model Training

Training and evaluating ML models, exploring three different architectures: Feedforward NN, ALSTM and CNN.

4 Deployment

Deployment of models to the live control system, and a brief look at future work.

Science and Technology Facilities Council

ISIS Neutron and Muon Source

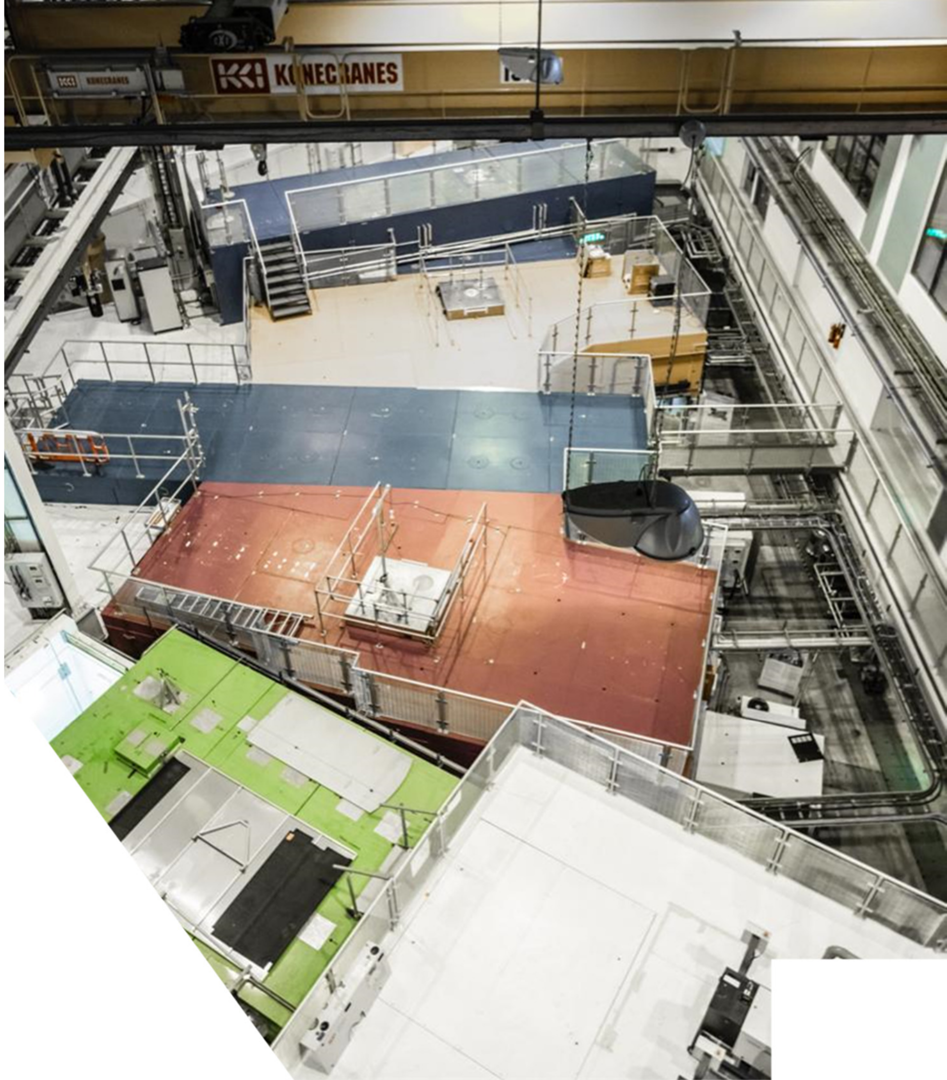

Case study:
ISIS TS1 Methane Moderator

ISIS Target Station 1 is used to generate neutrons through a spallation process – these neutrons must be moderated for suitable energy spectra to perform neutron scattering experiments.

For usable data in downstream instruments, a **very stable temperature** in the methane moderator is required.

The ISIS TS1 liquid methane moderator has operational issues that cause periods of flow and therefore pressure and temperature variability.

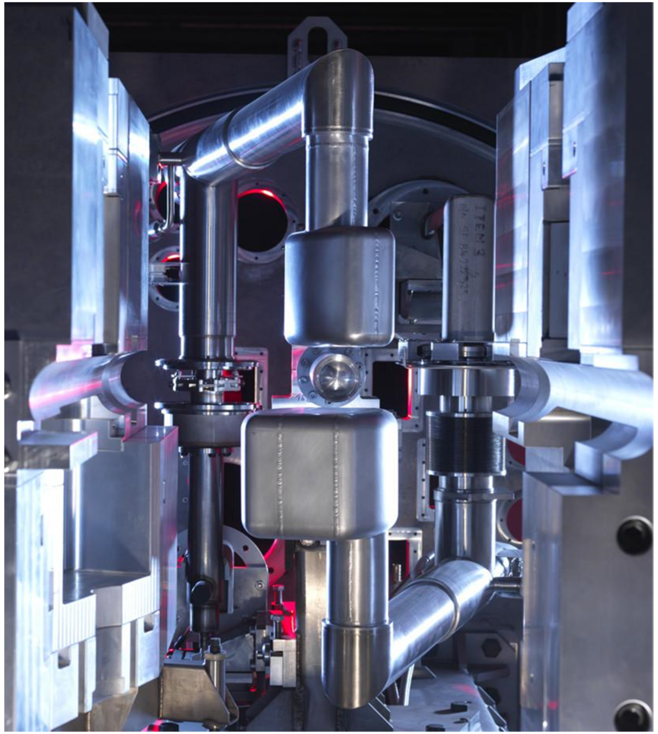

Science and
Technology
Facilities Council

ISIS Neutron and
Muon Source

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Data Preparation

The Raw Data

ISIS Accelerator Controls have recently undergone upgrades to store a large number of time series signals across the facility in InfluxDB.

Signals are recorded every 2-3s as long as the value of the signal as changed.

We focus on a single signal: TC68M thermocouple located by the liquid methane.

Pandas was used to filter periods of bad/out-of-cycle data, downsample the raw data and account for long-term signal drift.

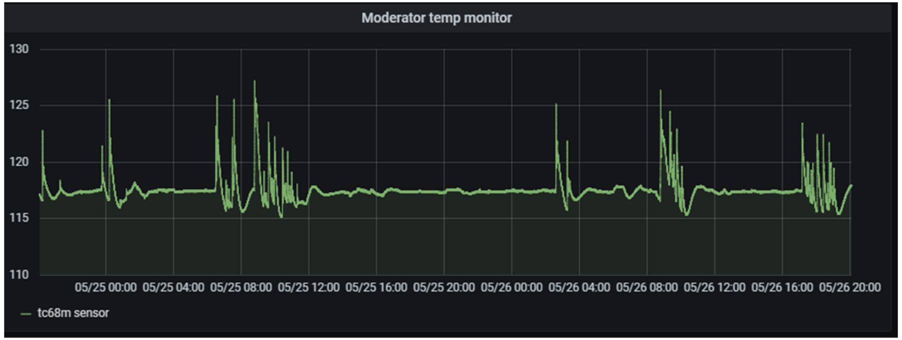

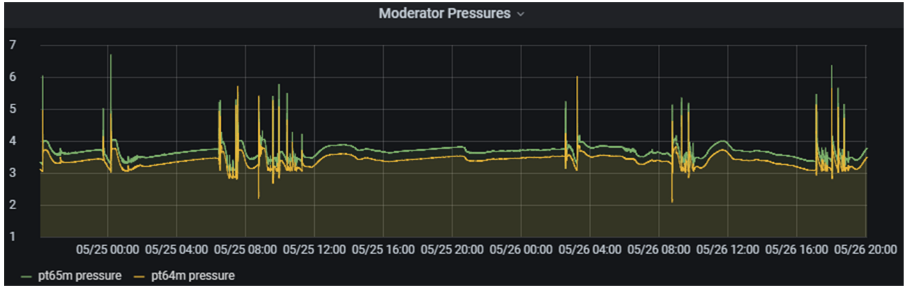

Feature Vector Generation

Anomalies tended to occur over 1-2 hours, but features associated with anomalies such as rapid oscillation/spikes in temperature over much shorter timespan

Settled on bin lengths of 20s, with each feature vector lasting 90 minutes: this gave feature vectors of dimension (270,1).

Want to evaluate models **constantly during live operation**, so feature vectors should be randomly sampled from the time series so that the training data is representative of real-time evaluation.

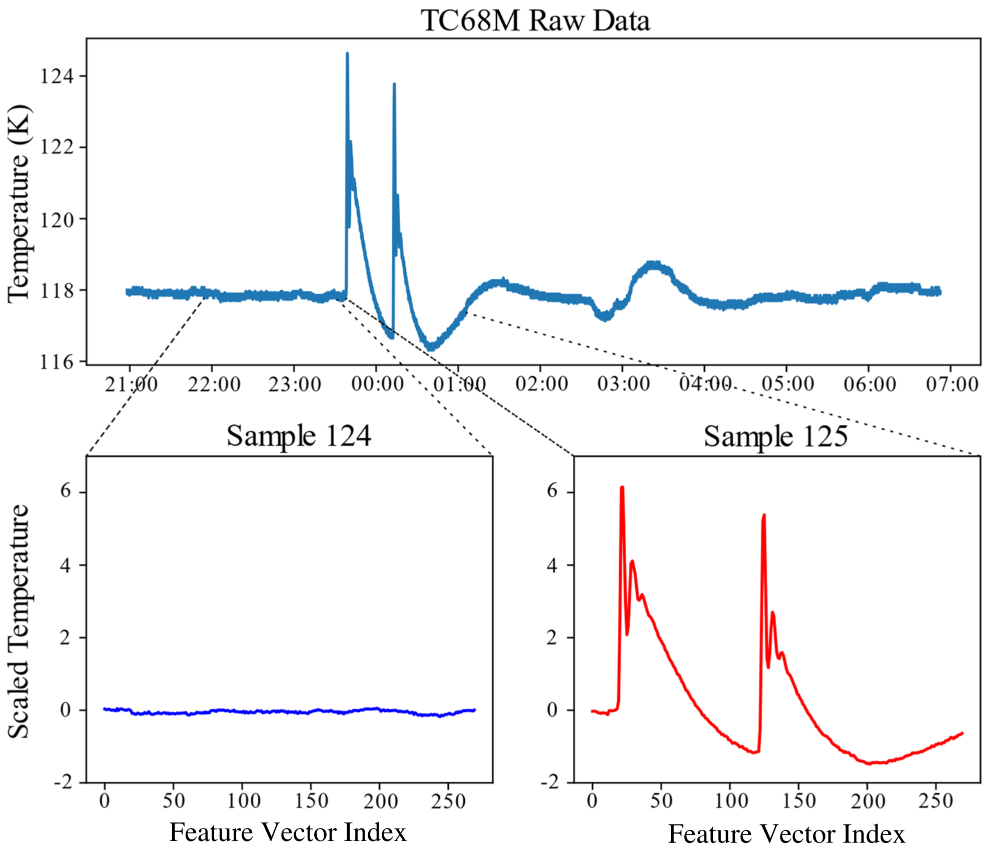

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Labelling The Data

Labelling the Data

In order to do supervised learning we need a class label associated with each feature vector.

Generated ~4300 feature vectors from the raw temperature data, but no class labels!

Manual labelling would be extremely tedious and require a lot of expert time and effort.

Settled on using t-SNE as a dimensionality reduction + clustering algorithm, with DTW as a “metric” for calculating similarities.

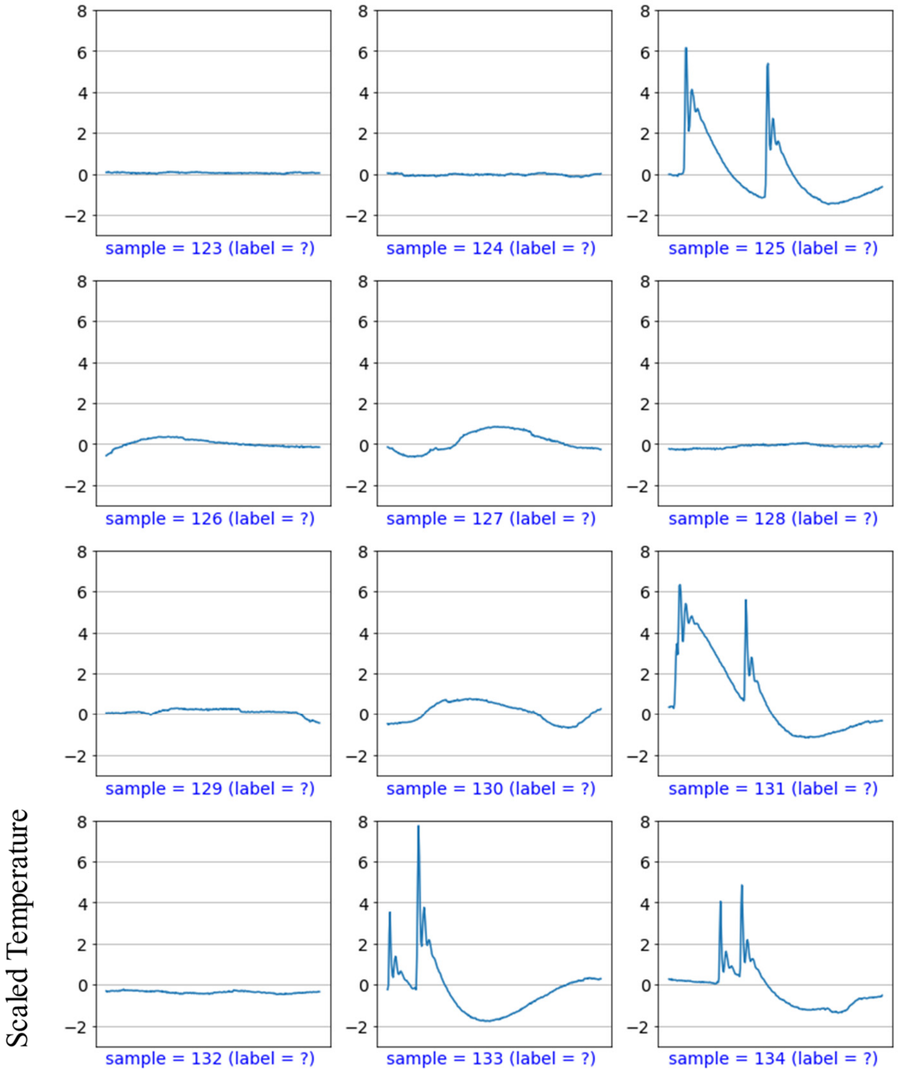

Science and Technology Facilities Council

ISIS Neutron and Muon Source

t-SNE

t-distributed Stochastic Neighbourhood Embedding is a popular dimensionality reduction/data visualisation/ clustering technique.

It works by taking a pairwise similarity matrix between all high-dimensional feature vectors and reproducing those similarities in a low dimensional embedding.

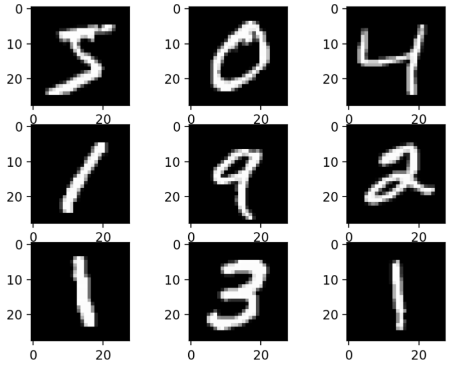

Figure from:
Van Der Maaten, L., 2009, April. Learning a parametric embedding by preserving local structure. *In Artificial Intelligence and Statistics (pp. 384-391).* PMLR.

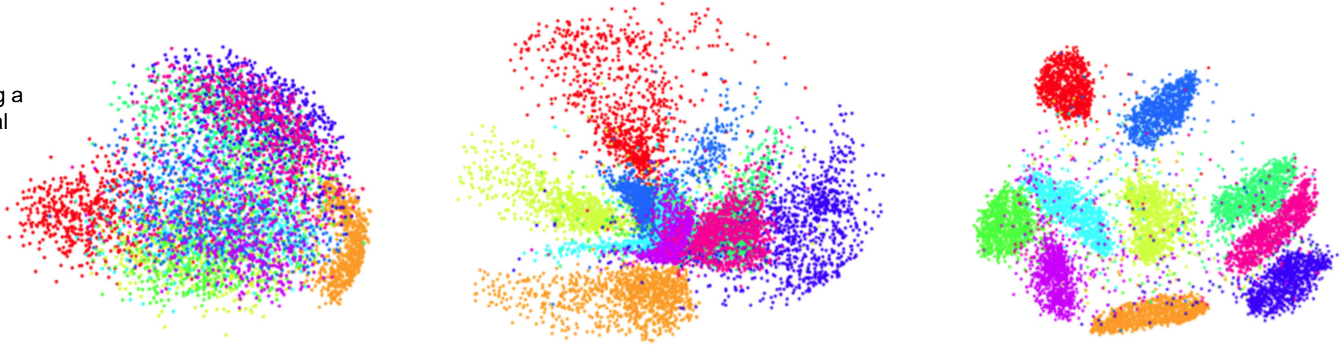

(a) Visualization by PCA. (b) Visualization by an autoencoder. (c) Visualization by parametric t-SNE.

Science and Technology Facilities Council

ISIS Neutron and Muon Source

t-SNE

In order to apply t-SNE to our time series data, we need a similarity metric that can distinguish between anomalies and normal operation.

For **misaligned** time series, the standard metrics such as Euclidean distance won't work – need something tailored for time series

Use Dynamic Time Warping algorithm, which can give estimates of similarity between misaligned time series.

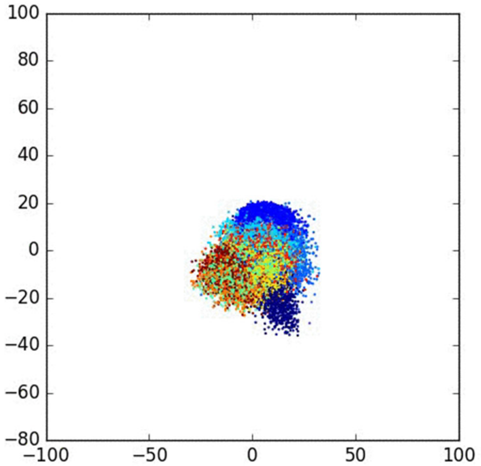

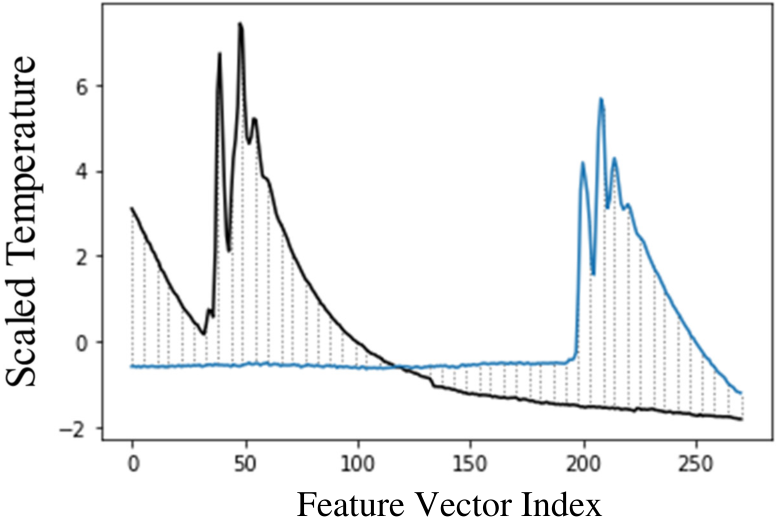

Science and Technology Facilities Council
ISIS Neutron and Muon Source

Dynamic Time Warping (DTW)

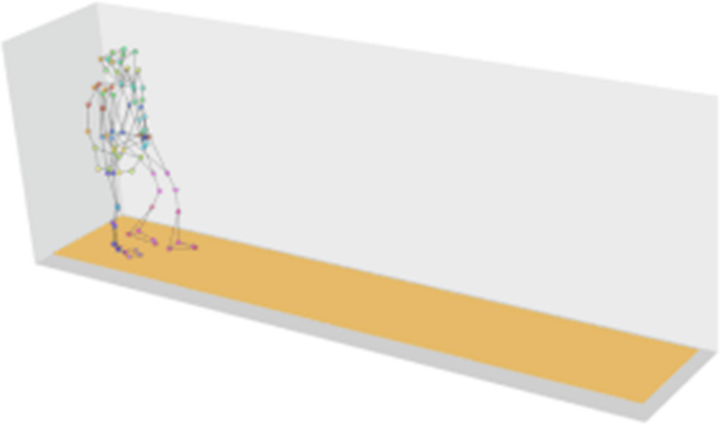

DTW searches for an optimal match between two time series, given some constraints, by warping the time series non-linearly along the time axis.

There are many adaptations of DTW that can be tailored for a particular application/set of constraints, but the standard algorithm sufficed for the temperature data in this problem.

N.B. Multivariate time series data must be treated slightly differently, but DTW can still be applied.

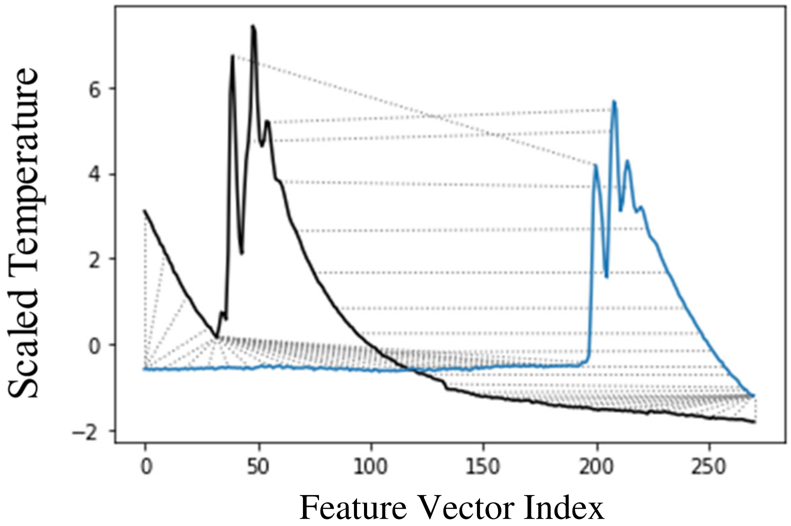

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Giorgino, T., 2009. Computing and visualizing dynamic time warping alignments in R: the dtw package. *Journal of statistical Software*, 31(7), pp.1-24.

Applying t-SNE + DTW

Generate the pairwise similarity for the whole dataset once – this can be a very computationally expensive step.

Run the t-SNE algorithm with many different hyperparameters, and explore the resulting low dimensional embedding with an interactive plotter.

Repeat with different hyperparameters for the algorithm until two large-scale, distinct clusters are generated.

Used scikit-learn python package for similarity matrix and t-SNE fit algorithm

Raw 2D t-SNE embedding

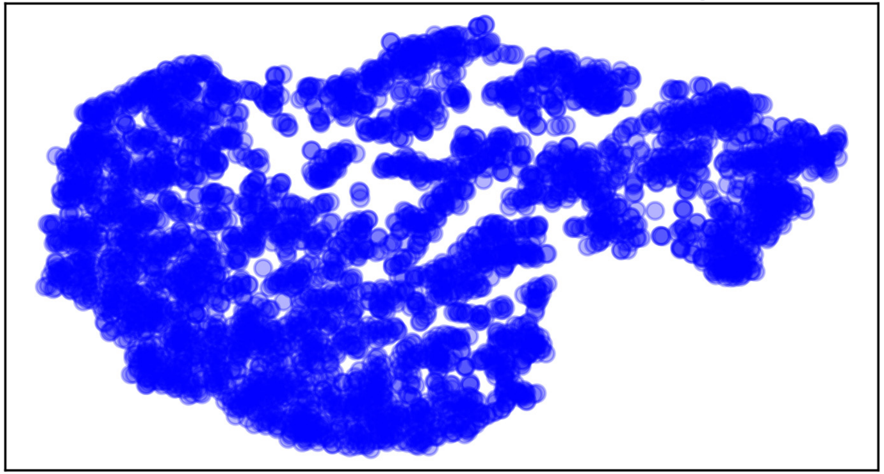

All ~4300 time series feature vectors (270-D) have been embedded into the 2-D map. Two large clusters can be discerned.

Applying t-SNE + DTW

Once the two clusters have been generated, we can manually define cut-off boundaries for the two classes.

For more complicated problems, could use unsupervised clustering techniques like k-means to define class boundaries.

Use the interactive tool to make corrections to the labels, particularly along the class boundaries.

A simple model or heuristic can then be used to sanity check/look for very obvious misclassifications.

Labelled 2D t-SNE embedding

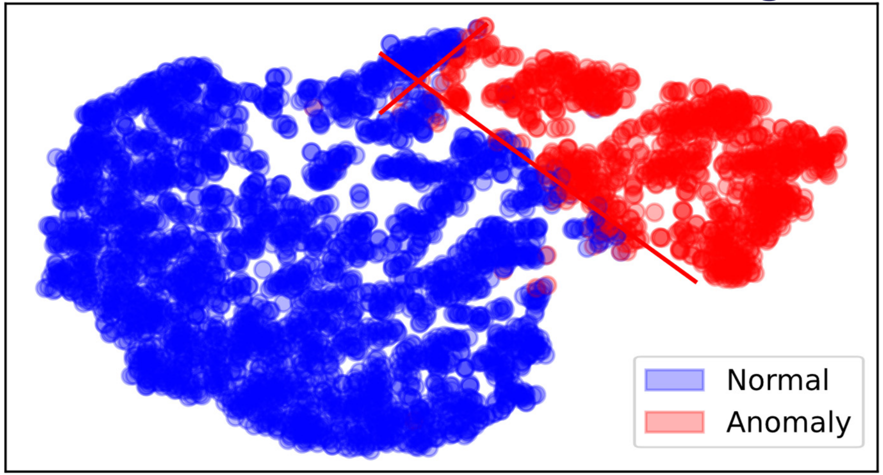

The generated class labels for each data point. Two approximate lines were used as a first pass for the class boundaries, and then a manual labelling process along the boundaries was used to make corrections.

Interactive plotter

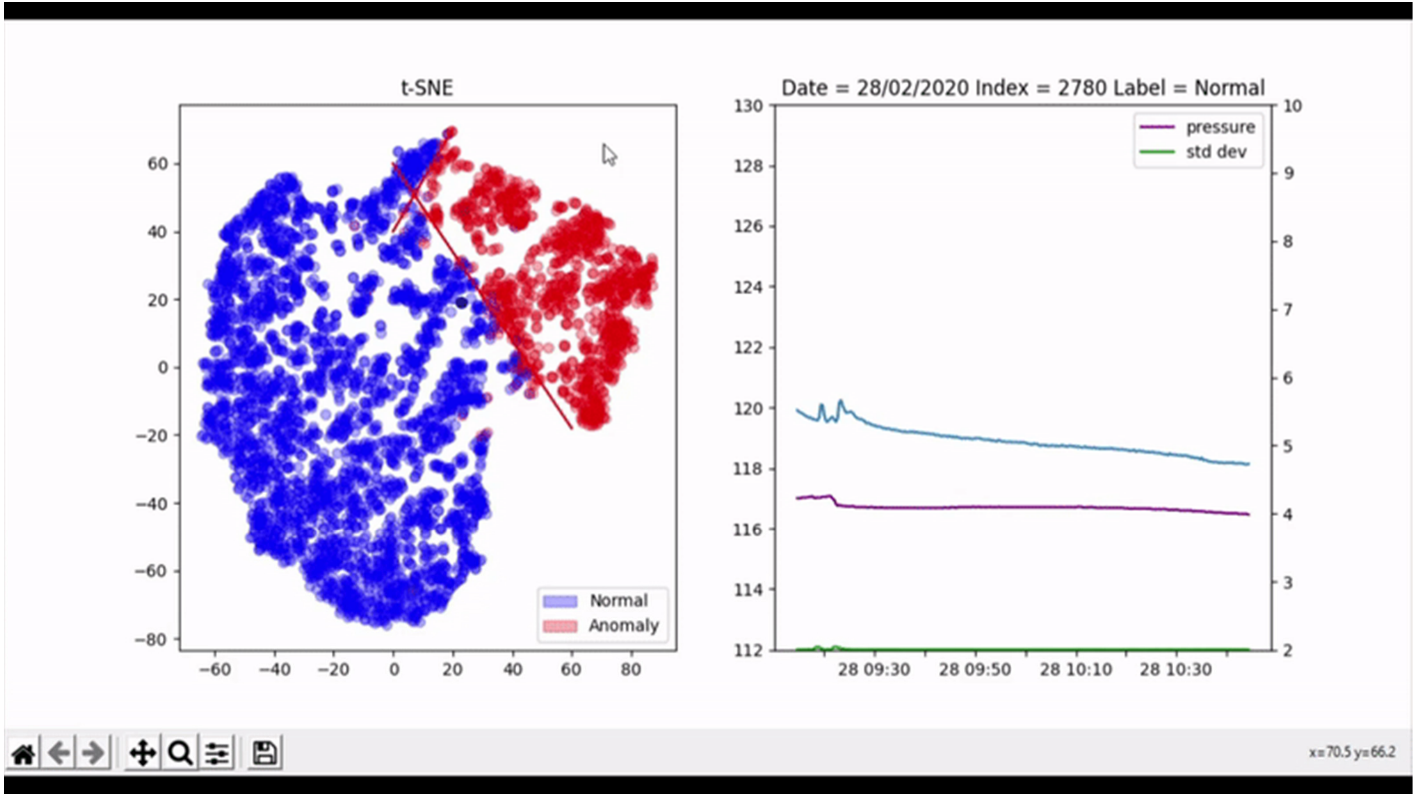

Science and
Technology
Facilities Council

ISIS Neutron and
Muon Source

Labelled Data

Once the labelling process is complete, we have a labelled dataset suitable for use with supervised ML algorithms.

Note: no such labelling process is going to be perfect, so there will be some degree of noisiness in the class labels – the model you train should be robust against noisy labels!

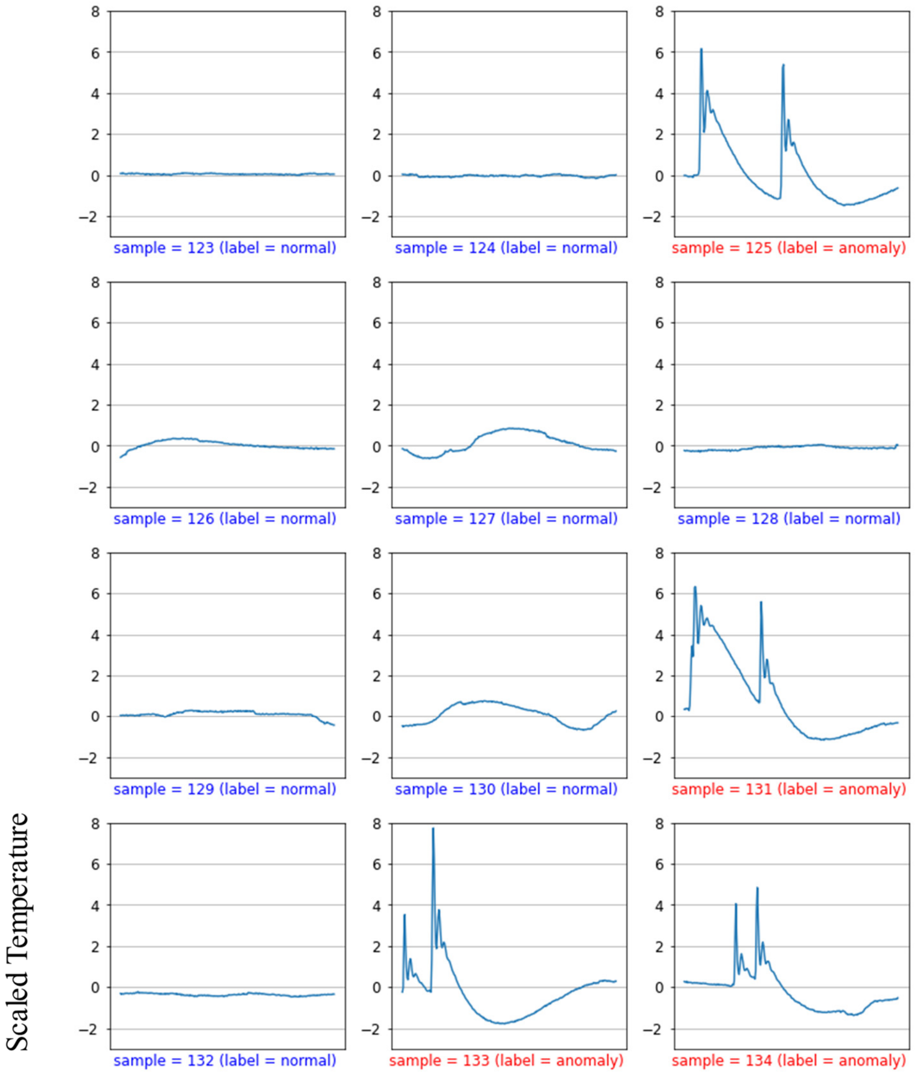

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Supervised Learning

Supervised Learning

Now we have labels corresponding to each feature vector, we can use supervised learning algorithms to train a model to complete this classification task.

We evaluate three different neural network architectures, all done in Python using Keras and Tensorflow.

For hyperparameter tuning, the labelled data was shuffled and split into a training set (65%), a validation set (15%) and a holdout test set (20%).

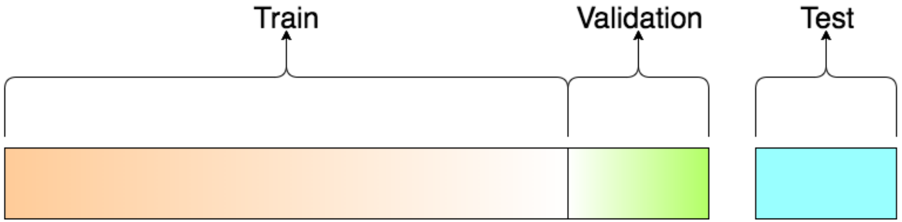

TensorFlow

Keras

Neural Network Models

We tune, train and evaluate three different types of neural networks:

- Feedforward Neural Network
- Attention-based LSTM
- Convolution Neural Network

Every network has a fixed input layer of length 270, and a softmax output layer of length 2 that outputs the probability of an input feature vector belonging to each class.

Trained using Adam optimiser + categorical cross entropy loss.

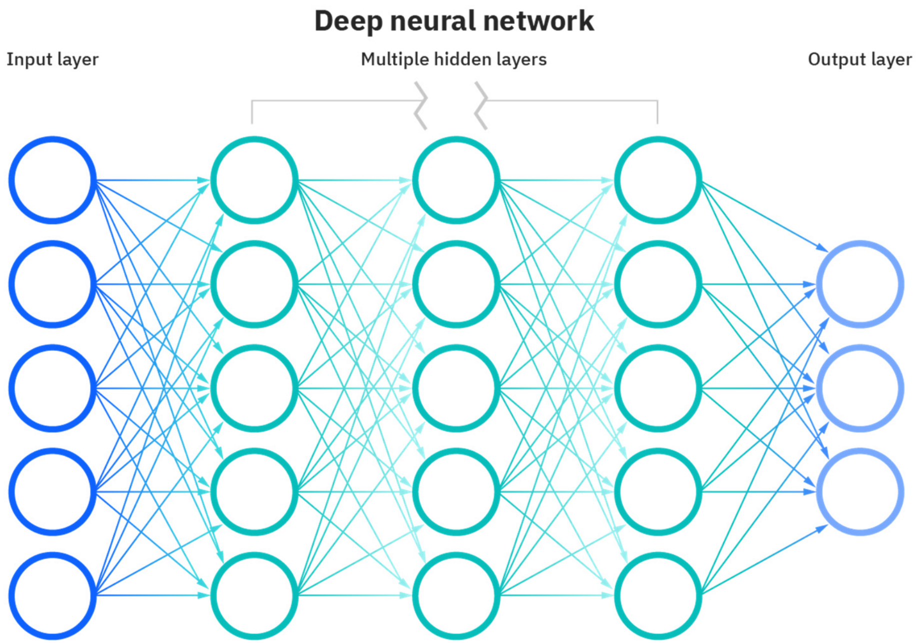

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Fawaz, H.I., Forestier, G., Weber, J., Idoumghar, L. and Muller, P.A., 2019. Deep learning for time series classification: a review. *Data mining and knowledge discovery*, 33(4), pp.917-963.

Feedforward Neural Network

Simplest neural network architecture: just a multi-layer perceptron with a softmax output layer.

Architecture:

- 4 hidden layers of size 64, 32, 16, 8 respectively.
- ReLU activation functions on each hidden layer
- Dropout of 0.3 after each hidden layer

This serves as a simple baseline model, but is not well-suited for time series analysis and classification.

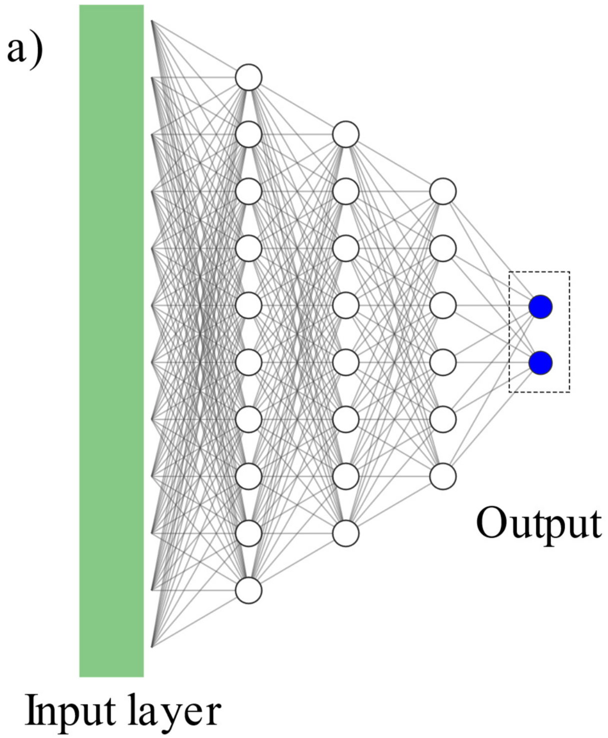

Attention-based LSTM

Adaption of the Long Short-Term Memory network, which itself is an adaption of the Recurrent Neural Network (RNN).

Architecture:

- LSTM cells of size 16 (16 hidden states passed between cells).
- Tanh activations on final LSTM output and attention output
- Combined before a dropout layer of 0.4.

Processes the input vector in chronological order, helping to preserve the temporal structure of the feature vector.

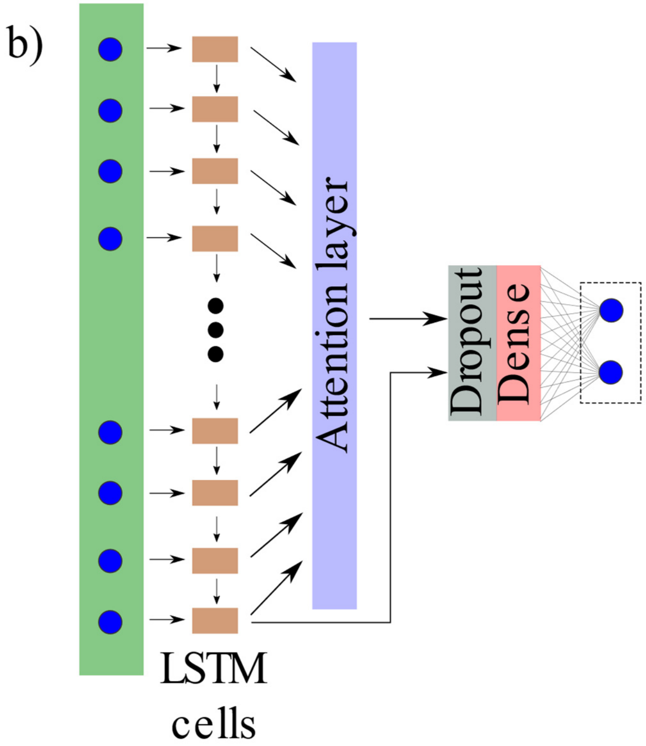

Karim, F., Majumdar, S., Darabi, H. and Harford, S., 2019. Multivariate LSTM-FCNs for time series classification. *Neural Networks*, 116, pp.237-245.

Convolutional Neural Network

This network uses 1D convolutional layers to learn filters that can be used to build up higher level representations of features in the time series.

Architecture:

- 3 1D convolutional layers, each layer followed by a batch normalisation (BN) step
- 1DConv layers have num filters (32, 64, and 32) and filter size (8, 5, and 3) respectively.
- Squeeze and Excite step after BN.
- Global max pooling before dropout of 0.4

CNNs learn shapes and edges that can be used to identify whether a time series “looks” like an anomaly.

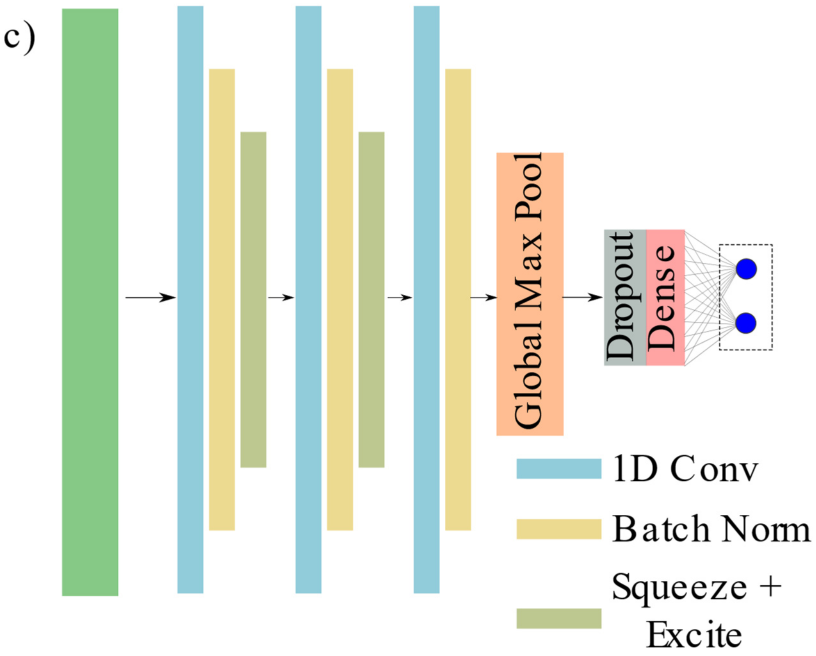

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Karim, F., Majumdar, S., Darabi, H. and Harford, S., 2019. Multivariate LSTM-FCNs for time series classification. *Neural Networks*, 116, pp.237-245.

Results

Accuracy was calculated by evaluating the trained models on the holdout test set to get the percentage of correct classifications.

CNN was the best performer, but feedforward NN not much worse.

On further analysis of misclassifications in the test set, it turned out that the majority of "misclassifications" were just bad labels – i.e. CNN was outperforming the labelling process.

Accuracy on the test set

Network Type	Accuracy (%)
Feedforward NN	97.5
ALSTM	96.0
CNN	98.3

CNN Misclassifications

CNN classification given first,
“ground truth” label given in brackets.

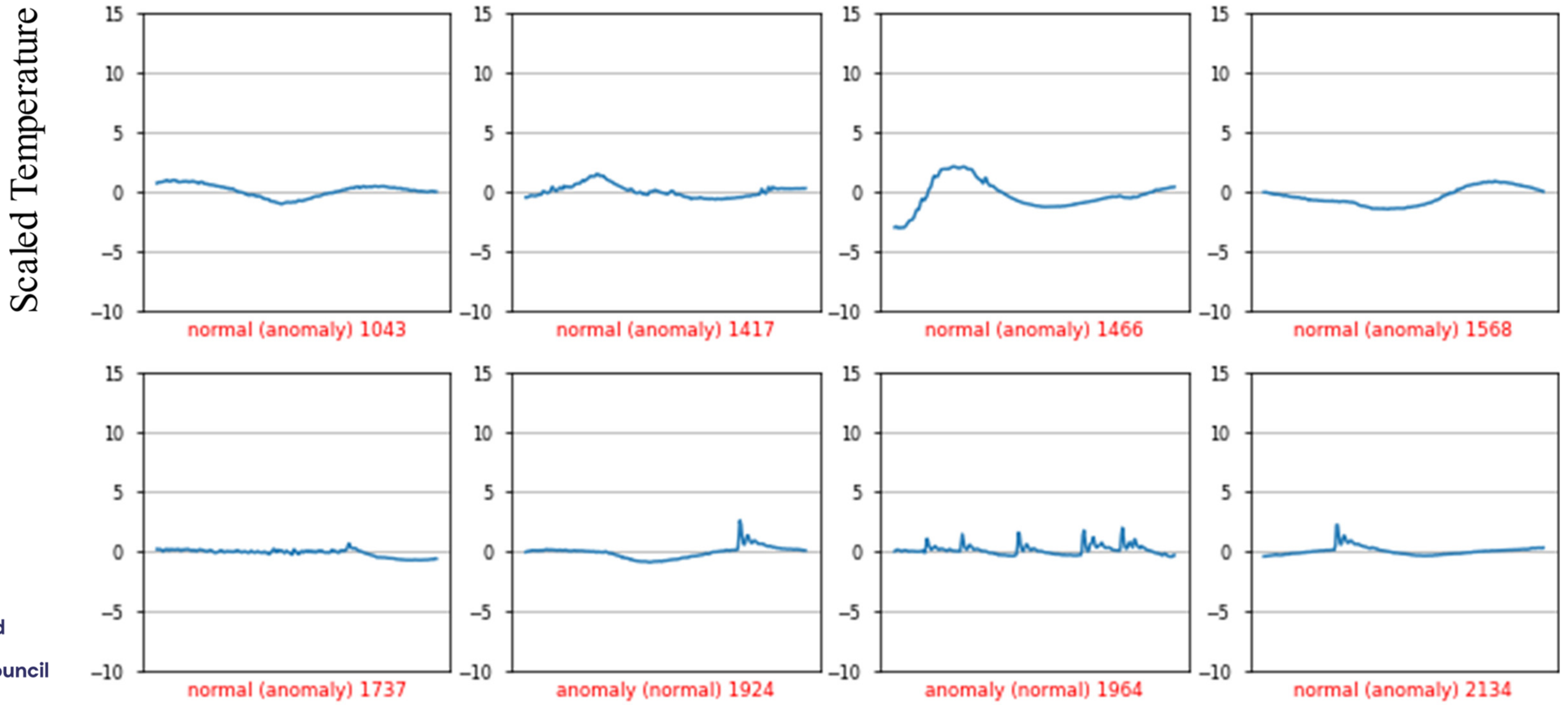

Explainability

Feedforward neural networks are traditionally seen as black box models.

Benefits of the ALSTM and CNN architecture is there are convenient methods for analysing why the network has made the classification it does.

Useful for operators in more complex problems where it may not be immediately obvious why a network has made the classification it does.

CNN Saliency Map

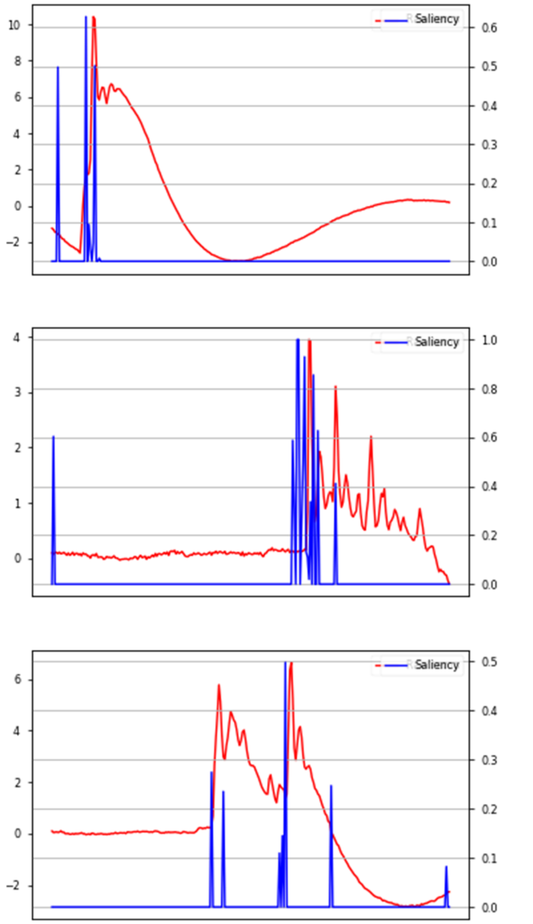

ALSTM Attention activations

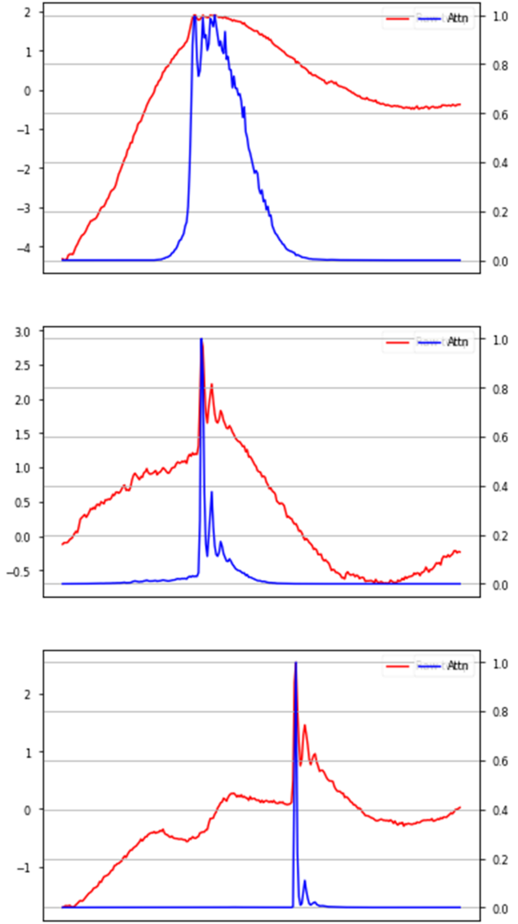

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Deployment

Deployment

Live data is used to create a new bin every 20 seconds, creating a new rolling feature vector.

Models are evaluated and fed back in to the ISIS controls messaging system for display.

Alarms are also generated in the data analysis layer, which are sent back to the Vsystem control system for display in the Main Control Room.

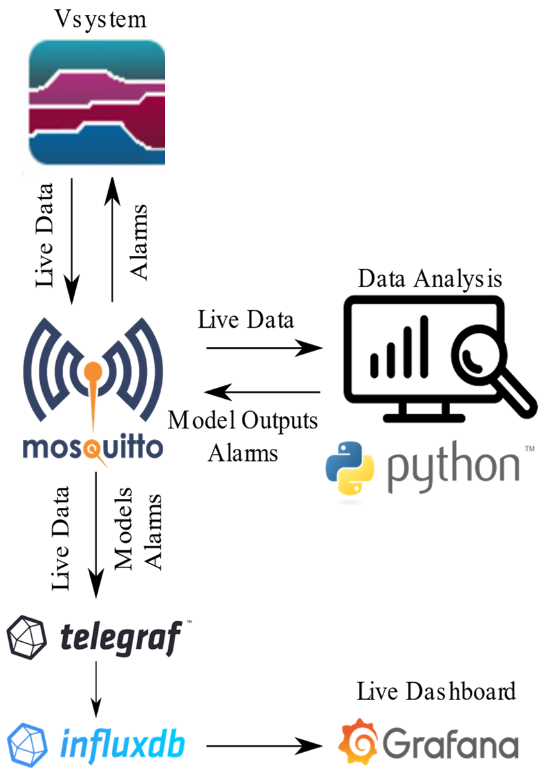

Software stack to apply to models and visualise the live data + model outputs in real time for operators.

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Application – Grafana Dashboard

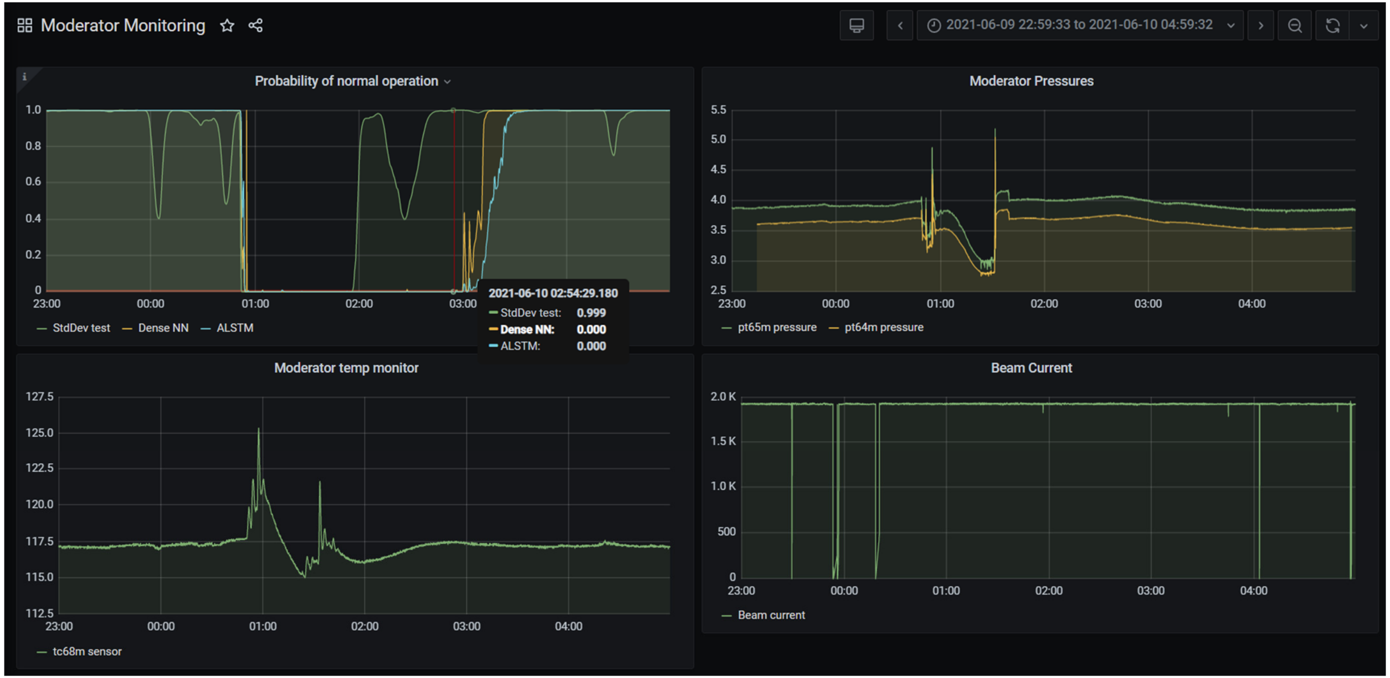

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Future Work

Anomaly Prediction

Work was extended to predicting anomalies – this required better labelling process that found the timestamp of anomalies.

Once timestamp found, feature vector is the period before the anomaly.

Much harder problem, so used several extra signals such as pressures and beam current – becomes a multivariate time series classification problem.

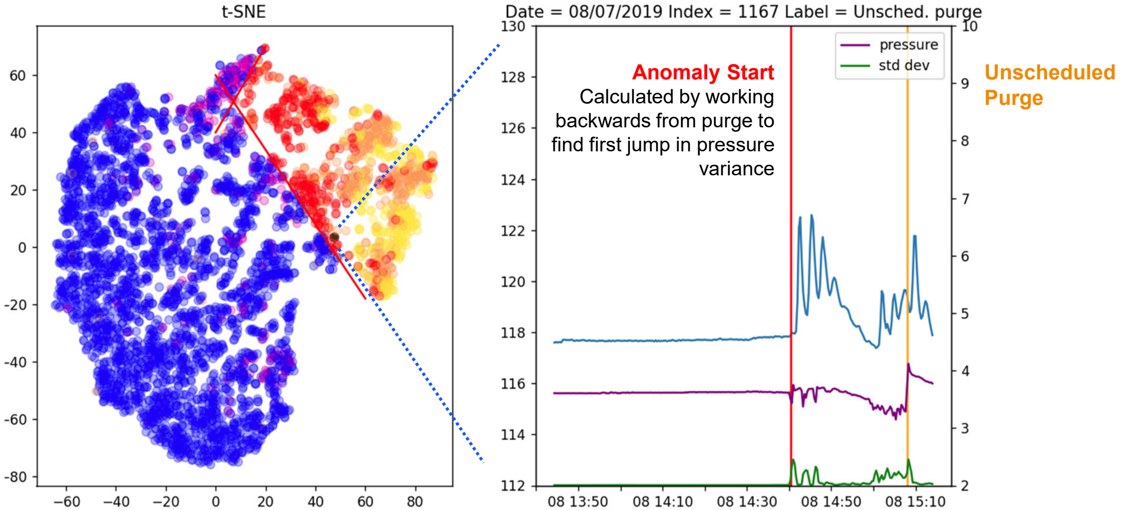

Each point in t-SNE is a single temperature sample.

Anomaly Prediction

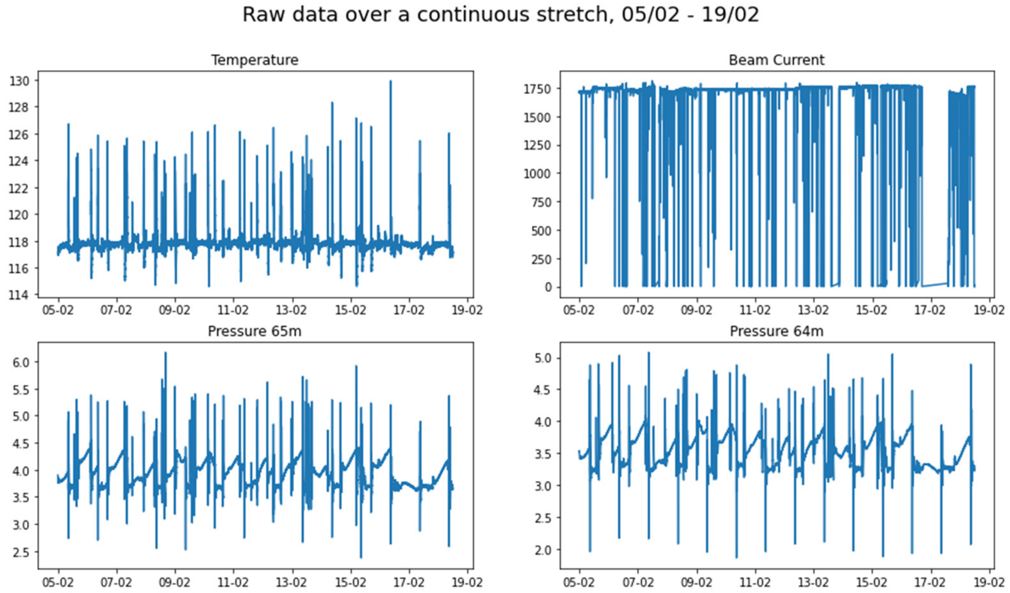

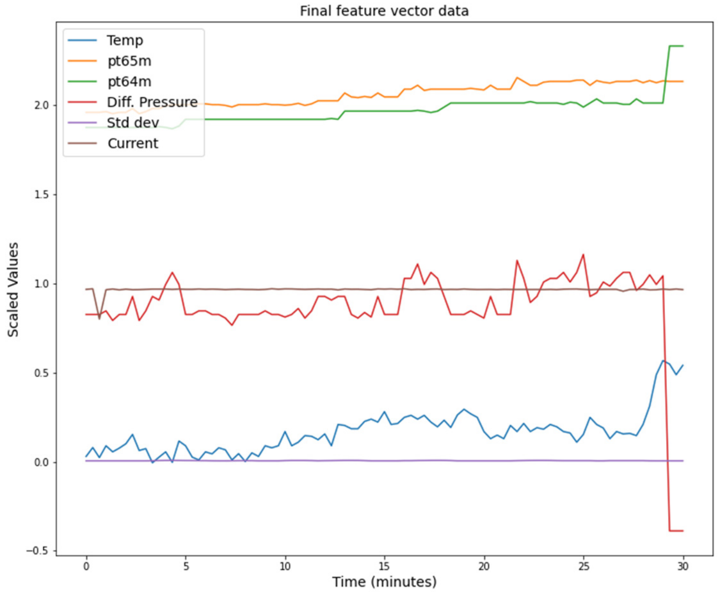

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Anomaly Prediction

Used an ALSTM-FCN architecture – basically just a combined version of the ALSTM and CNN models from earlier.

Results were very poor – anomaly prediction is a much harder task.

Future avenues to explore: more signals + feature engineering, different models like Gradient Boosted Trees, and dealing with the worsened class imbalance through data augmentation.

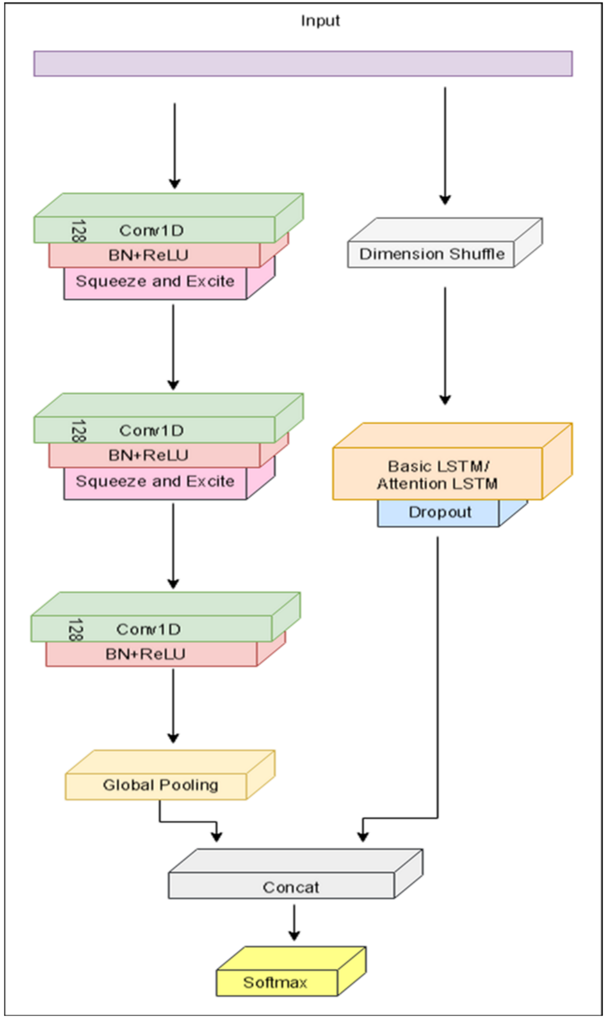

Figure from: Karim, F., Majumdar, S., Darabi, H. and Harford, S., 2019. Multivariate LSTM-FCNs for time series classification. *Neural Networks*, 116, pp.237-245.

Science and Technology Facilities Council

ISIS Neutron and Muon Source

Thank you

isis.stfc.ac.uk @ISISNeutronMuon STFC