ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB260

DETECTION AND CLASSIFICATION OF COLLECTIVE BEAM BEHAVIOUR IN THE LHC

L. Coyle*[1], F. Blanc, T. Pieloni[1], M. Schenk[1], EPFL, Lausanne, Switzerland
X. Buffat, M. S. Camillocci, J. Wenninger, CERN, Geneva, Switzerland
E. Krymova, G. Obozinski, SDSC, ETH Zürich and EPFL, Lausanne, Switzerland
[1]also at CERN, Geneva, Switzerland

Abstract

Collective instabilities can lead to a severe deterioration of beam quality, in terms of reduced beam intensity and increased beam emittance, and consequently a reduction of the collider's luminosity. It is therefore crucial for the operation of the CERN's Large Hadron Collider (LHC) to understand the conditions in which they appear in order to find appropriate mitigation measures. Using bunch-by-bunch and turn-by-turn beam amplitude data, courtesy of the transverse damper's observation box (ObsBox), a novel machine learning based approach is developed to both detect and classify these instabilities. By training an autoencoder neural network on the ObsBox amplitude data and using the model's reconstruction error, instabilities and other phenomena are separated from nominal beam behaviour. Additionally, the latent space encoding of this autoencoder offers a unique image like representation of the beam amplitude signal. Leveraging this latent space representation allows us to cluster the various types of anomalous signals.

INTRODUCTION

When operating the LHC, understanding and mitigating instabilities is essential for efficient and safe operation of the machine. As such, the turn-by-turn and bunch-by-bunch transverse beam position data, essential for the LHC's feedback system, is made available through the transverse feedback system Observation Box (ObsBox) [1].

The ObsBox maintains rolling buffers of the transverse bunch position signals and saves the buffers to disk when receiving manual or automatic triggers. An automatic triggering system was developed to automatically detect instabilities and save the corresponding buffer [2]. This instability detection system, being very sensitive, produces a very large amount of false positives. In practice, the vast majority of saved data does not contain any instabilities. The analysis of this instability data is hindered by the large amount of false triggers, and any analysis is done on a per case level.

This paper describes a novel method of detecting instabilities in the transverse beam position data stream and provides a framework with which to automatically cluster similar anomalous behaviour.

METHODS

The method described in this paper makes use of the large amount of accumulated false triggers to train an anomaly detection model. This model, once trained, can be applied on the transverse beam position data stream to filter out anomalous from nominal beam behaviour. As such, the anomaly detection model is designed with online operation in mind.

The clustering of these anomalous signals is then performed by a clustering model. The clustering relies on extracting features from the input time series, these features vectors are then used to cluster the anomalies.

A schematic diagram of the models is shown in Fig. 1.

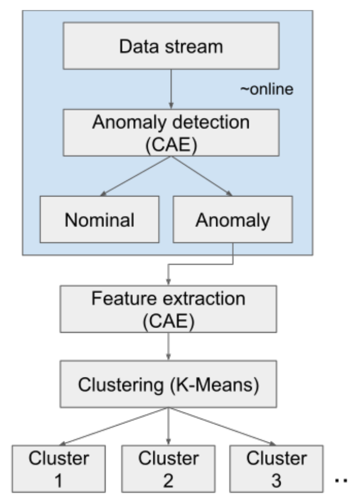

Figure 1: Model overview.

Autoencoders

Both the anomaly detection and feature extraction models use Convolutional Autoencoders (CAE).

Autoencoder [3] models are trained to reconstruct the input data despite having to pass through a bottleneck. By doing this they learn to compress and decompress the input to and from a latent space representation. If the latent space is successfully learnt, the bad quality of the reconstruction is a sign of an anomalous input sample. The larger the reconstruction error the more anomalous a given data sample is.

In this case, as the data are time series, successive 1D Convolutional layers are used in the AE to learn filters which convolve over the time dimension of the input. These convolutional filters detect and encode recurring patterns in the time series. In this case, the CAE used are tuned to obtain a 4×4 latent space which empirically was found to provide

* loic.coyle@epfl.ch, loic.thomas.coyle@cern.ch

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB260

a sufficiently low dimensional space for the clustering algorithm while still allowing for good signal reconstruction through the anomaly detection CAE.

The anomaly detection CAE and feature extraction CAE, while being identical in structure, serve two distinct purposes. The anomaly detection CAE uses the reconstruction error to detect anomalous samples, whereas the feature extraction CAE is trained from anomalous data so that its latent space encoding provides unsupervised feature extraction for the clustering algorithm. The latent space encoding of the anomaly detection CAE is not used for feature extraction as we are interested in clustering the anomalous signals only, which, by definition, are badly reconstructed, thereby indicating that their latent space encodings are not representative of the input time series.

Clustering

Once the feature are extracted, any vector based clustering algorithm can be used. In this paper K-Means [4] is showcased. In order to determine the optimal number of clusters for the K-Means algorithm, the elbow criteria [5] is employed.

DATA PREPROCESSING

The data used to train the model, is the ObsBox's instability buffer which is 65536 turns long. Specifically the beam 1 horizontal plane data is used. This data covers all beam modes and operation schemes.

In order to facilitate the training of the model, some preprocessing is applied to the data. However, it is important to keep in mind that this preprocessing should be minimal so as to facilitate online integration. Namely, the rolling mean and standard deviation are computed from the raw transverse beam position buffer. The 65536 turn buffer is then split into smaller 2048-turn chunks. Each chunk is individually normalized between 0 and 1.

TRAINING

Anomaly Detection

The CAE anomaly detection model is trained for 50 epochs on a random subset of the entire data. A validation dataset is used to monitor overfitting. As this is an unsupervised problem, no labelled test data set is available. Overall, the training dataset contains roughly 4 million 2048-turn samples.

As shown in Fig. 2 the training has converged and, given that loss functions values are very close for both training and validation datasets, there is no obvious overfitting.

Once trained, we are able to select the anomalous signals by selecting data with a high reconstruction error when passed through the CAE model. It is clear that the anomalies contain more than just collective instabilities, they contain anomalous behaviour such as injection oscillations, beam dumps, logging errors, etc. By tuning the reconstruction error threshold we are able to control the sensitivity of the anomaly detection model.

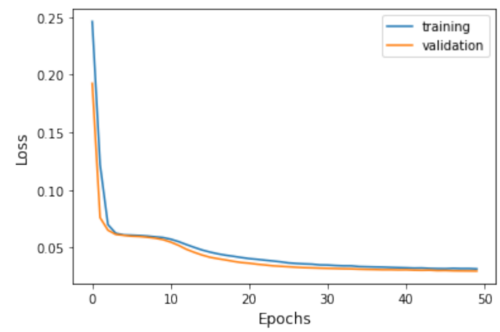

Figure 2: Training loss.

Notice that two similar transverse beam position signals, despite being very similar have drastically different latent space encodings as the anomaly detection CAE's encoding is accurate for the nominal signals and not the anomalous signals, as shown in Fig. 3.

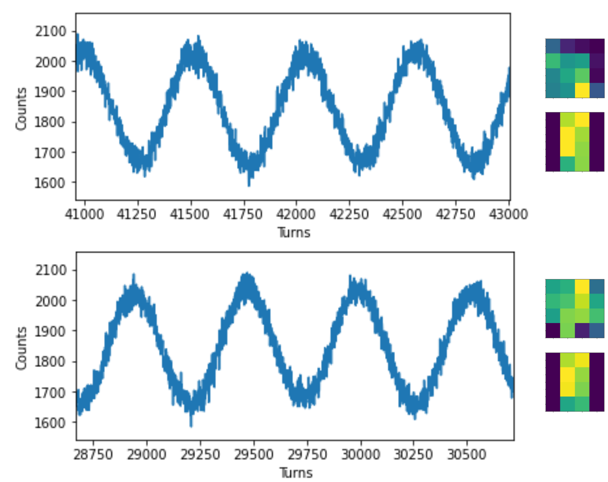

Figure 3: Two similar anomalous time series, from fill 6364, and their latent space encodings. (Left) raw transverse beam position time series. For each time series, the encoding of the anomaly detection CAE (top 4×4 patch) and of the feature extraction CAE (bottom 4×4 patch) are represented.

To illustrate the clustering, the 64 most anomalous signals are kept for clustering.

Clustering of Anomalies

The second feature extraction CAE is trained on the anomalous signals. This provides an accurate latent space encoding of these anomalous signals. This feature extraction CAE is identical to the anomaly detection CAE, with the only difference being the training dataset.

The latent space encoding of this second CAE is more consistent, as shown in the bottom encoding images of Fig. 3, similar time series have similar latent space encodings.

Once the latent space encodings of the anomalous signals are obtained, K-Means [4] is applied. Using the elbow method [5] with the distortion criteria, as shown in Fig. 4, the optimal number of clusters is determined to be 6.

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB260

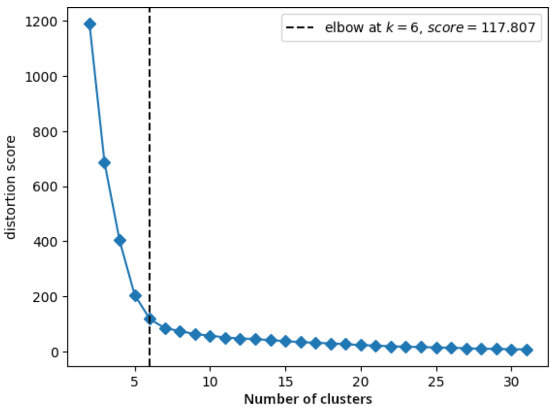

Figure 4: Distortion score elbow for K-Means clustering.

RESULTS

Using K-Means on the extracted features, we obtain 6 clusters; a subset of each cluster is shown in Fig. 5. The complete cluster assignments can be found online at [6].

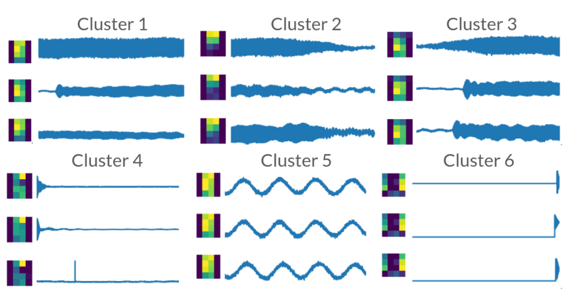

Figure 5: The raw transverse beam position data of a few of the anomalous samples in each cluster, along with their respective latent space representation.

We observe that cluster 6 contains the start of injection oscillations. Cluster 4 mainly contains the end of the injection oscillations. Cluster 5 contains lower frequency orbit oscillations potentially caused by longitudinal oscillations. Cluster 1, 2 and 3 contain interesting growths and drops in oscillation amplitudes some of which are instability candidates. Since there is no ground truth available for the clustering problem, no quantitative performance metric is available to evaluate the quality of the clustering. However, applying the same the clustering model on a subset of the 64 anomalous samples results in the same cluster assignments, indicating good cluster stability.

Additionally, for comparison, we applied an alternative clustering method on the same anomalous signals. This time no feature extraction step is required, instead, Dynamic Time Warping (DTW) [7] is used to compute the similarity between time series and Hierarchical Clustering (HC) [8] to assign the clusters, see [9] for the complete cluster assignments. The cluster cut off threshold is tuned to obtain 6 clusters so as to be comparable with the CAE & K-Means approach.

Visual inspection shows that both methods produce coherent clusters with only a few samples assigned to questionable clusters. The CAE & K-Means approach has the advantage of being much more scalable due to its cheaper computational requirements, especially once trained. Whereas, HC requires the computation of all pairwise similarities which scales as $O(n^2)$ and DTW itself is typically also $O(n^2)$, this quickly becomes unwieldy. Moreover, the first method can easily be applied on new data samples as it is an inductive model, i.e. it can be trained on a dataset and assign clusters on new previously unseen data samples, which is a very appealing feature and is not the case for the DTW & HC method.

Moreover, the latent space encoding used in the CAE & K-means approach could very well prove to be useful for multi-bunch instability clustering. Using this latent space encoding a beam becomes a stack of images on which convolution models could be applied to cluster multi-bunch behaviour.

CONCLUSION

A novel, data driven, instability triggering methodology has been developed based on a CAE model. This model can successfully detect anomalous behaviour in the transverse bunch position signal. The anomalous samples can then be assigned to clusters by using a feature extraction CAE combined with any standard clustering algorithm, in this paper K-Means was showcased.

The quality of the resulting clustering of the anomalous signals' latent representation is comparable to HC with DTW applied directly to the anomalous time series. At the same time the proposed approach is inductive and more scalable.

However some aspects still could be improved upon:

- The anomaly detection model has, as of yet only been trained on a subset of the entire dataset.
- The tuning of the hyper parameters of the various models is challenging in the absence of ground truth.
- Extending the clustering to take into account multi-bunch behaviour.

ACKNOWLEDGEMENTS

This work is partially funded by the SDSC project C18-07.

REFERENCES

[1] L. R. Carver *et al.*, "Usage of the Transverse Damper Observation Box for High Sampling Rate Transverse Position Data in the LHC", in *Proc. 8th Int. Particle Accelerator Conf. (IPAC'17)*, Copenhagen, Denmark, May 2017, pp. 389-392. doi:10.18429/JACoW-IPAC2017-MOPAB113

[2] M. E. Söderén, G. Kotzian, M. Ojeda Sandonís, and D. Valuch, "Online Bunch by Bunch Transverse Instability Detection in LHC", in *Proc. 8th Int. Particle Accelerator Conf. (IPAC'17)*, Copenhagen, Denmark, May 2017, pp. 397-399. doi:10.18429/JACoW-IPAC2017-MOPAB117

[3] M. A. Kramer, "Nonlinear principal component analysis using autoassociative neural networks", *AIChE Journal*, vol. 37, no. 2, pp. 233–243, Feb. 1991. doi:10.1002/aic.690370209

ISBN: 978-3-95450-214-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2021-THPAB260

[4] J. MacQueen, “Some methods for classification and analysis of multivariate observations”, in *Proc. of 5th Berkeley Symposium on Mathematics, Statistics and Probability*, Berkeley, CA, USA, Jun.-Jul. 1967, pp. 281-297.

[5] M. A. Syakur *et al.*, “Integration K-Means Clustering Method and Elbow Method For Identification of The Best Customer Profile Cluster”, *IOP Conference Series: Materials Science and Engineering*, vol. 336, p. 012017, Apr. 2018. doi:10.1088/1757-899x/336/1/012017

[6] CAE_Kmeans_clusters.png, https://cernbox.cern.ch/index.php/s/U44M17XP313xKWY

[7] R. Bellman and R. Kalaba, “On adaptive Control Processes”, *IRE Transactions on Automatic Control*, vol. 4, no. 2, pp. 1-9, Nov. 1959. doi:10.1109/TAC.1959.1104847

[8] D. Müllner, “Modern hierarchical, agglomerative clustering algorithms”, 2011. https://arxiv.org/abs/1109.2378v1

[9] DTW_HC_tree_clusters.png, https://cernbox.cern.ch/index.php/s/EBO3irck0EolZ2z