Use of MHD Activity for Disruption Prediction in Tokamaks

Tiago Agostinho Martins
tiago.agostinho.martins@tecnico.ulisboa.pt

Instituto Superior Técnico, Lisboa, Portugal

October 2021

Abstract

Disruptions in tokamaks are one of the biggest challenges to the viability of nuclear fusion. One of the main causes of disruptions is the occurrence of locked modes, which are non-rotating unstable magnetohydrodynamic (MHD) modes. In this work, we intend to use deep learning tools to predict these types of disruptions and to understand other possible triggering mechanisms. We developed a model based on a convolutional neural network (CNN) that receives MHD activity in the form of a spectrogram, and predicts whether it is going to disrupt due to mode-locking. The model can reasonably distinguish the classification classes, achieving an accuracy of almost 80%, although it cannot be compared directly with other state-of-the-art predictors. Additionally, an interpretable machine learning tool called class activation mapping (CAM) is used in an attempt to understand which MHD activity is relevant to the model. The results indicate a considerable focus placed by the CAM method in the interruption of MHD activity, in particular internal kink modes, followed by its resurgence before the development of the locked mode. This observation is congruent with the physical explanation given by the spectrogram and with previous studies.

Keywords: Tokamaks; Magnetohydrodynamics; Disruption prediction; Spectrograms; Convolutional neural networks; Interpretable machine learning

1. Introduction

A tokamak disruption is an abrupt event where the confinement of the plasma is practically lost. Due to its destructive potential and imposition of limits on the machine's operational space, it is crucial to develop methods that can detect either the early-stage development of disruptions to prevent them and regain plasma confinement or mitigate its effects on the device.

These events are the direct or indirect consequence of magnetohydrodynamic (MHD) instabilities. One of the most known precursors linked to disruption triggering is the locked mode [1]. The signal of the mode is used for current disruption mitigation systems by defining a critical threshold to it [2]. The defined thresholds, however, may depend on the analyzed dataset and can present some fluctuations [3]. Additionally, there is a risk that a considerable number of discharges are not detected and non-disruptive discharges falsely activate the mitigation systems with the threshold method if they happen very close to the disruption instant [4].

Researchers are putting effort into the use of machine learning and deep learning algorithms [5, 6, 7], such as deep neural networks, which are applied to data coming from the diagnostics available in the reactor to better understand disruptions and predict them. The complexity of the physical phenomena involved in the dynamics of disruptions has not yet, however, allowed a complete theoretical understanding and the establishment of real-time disruption prediction systems which can detect all disruptive discharges. Also, the complexity of the obtained machine learning models may hinder the subsequent analysis of the results obtained, thus interpretability is desired [8].

In this work, we try to predict disruptions due to the locked mode with a deep learning model. Moreover, we retrieve the most important features obtained by the model for a given output to look for physical insight.

2. MHD Stability

The MHD model provides a framework that allows to describe the plasma macroscopic behavior in a tokamak. The stability analysis, including the study of possible precursor modes during a disruption, can be described with a displacement perturbation vector ξ,

$$\xi(\mathbf{r}, t) = \xi(\mathbf{r}) e^{i(m\theta - n\phi + wt)} \tag{1}$$

defined in space and time, where θ and ϕ are the poloidal and toroidal angles, respectively, w is the frequency of the perturbation, m and n are the respective poloidal and toroidal mode numbers. Assuming that the perturbed quantity amplitudes are much smaller than the equilibrium state, the perturbation vector is introduced in the linearized MHD equations for stability analysis within the perturbed fluid velocity term, v_1, as seen in equation 2.

$$v_1 = \frac{\mathrm{d}\xi}{\mathrm{d}t} \tag{2}$$

In non-ideal MHD, the frequency w has both real and imaginary parts (w_r and w_i), and the stability of the mode is determined by the sign of w_i. If w_i is lower than 0, then the system is said to be unstable, and stable otherwise. By solving the linearized equations in terms of ξ, one finds the eigenvalue equation,

$$-w^2\rho\xi = F(\xi) \tag{3}$$

where F is called the force operator. It is by isolating F that one finds the stability limits for a given displacement.

When taking into consideration the resistive MHD theory a specific type of unstable modes driven by current gradients, called tearing modes, can appear. Destabilizing these modes can produce poloidal regions where there is a reconnection of the magnetic field lines. This is typically associated with the formation of magnetic islands [9], where the magnetic field lines break and reconnect, forming island-like structures. A similar instability to the tearing mode can also occur by perturbations in the bootstrap current, facilitating the growth of the magnetic island. These are called neoclassical tearing modes (NTM) [10], and they can contribute to confinement degradation, in particular with $m = 2$, $n = 1$ modes.

During the growth of a tearing mode or a NTM, the width of a magnetic island can be large enough to have considerable interaction with the walls of the reactor. This interaction can cause the deceleration of the mode, eventually stopping its rotation and making it to lock. The deceleration is attributed to applied torques in the vacuum vessel wall [11]. This is a brief description of the locked mode.

3. Experimental Setup

In this work, we use two main data sources from the Joint European Torus (JET) tokamak. Both are part of the magnetic diagnostics of the machine. They consist of a system of coils and probes installed along the vacuum vessel that measure the magnetic field fluxes, allowing important parameters such as the plasma position on the vessel, the rotation frequency, and plasma current, to be retrieved [12]. In order to study the plasma equilibrium, stability, and spectral response, a set of inductive, cylindrical coils are used, designated as Mirnov coils. They consist of N loops of titanium wire that measure the time variation of the local poloidal magnetic field, B_θ.

A single Mirnov coil (H305 coil) was used for this work. It is part of an array of magnetic probes (High Resolution Array Coils) distributed at different toroidal angles and in the same poloidal section. The integrated data acquisition system of the H305 coil provides a maximum sampling rate f_s of 2 MHz. However, a downsampling to 125 kHz was made.

Figure 1 shows the magnetic pick-up coil data from the database discharge 92213. The disruption instant is visible with a well-defined spike. It is also clear that the signal oscillates considerably throughout time, with alternating peaks. Due to this behavior, the use of processing techniques for non-stationary signals are more suitable to capture the frequency component.

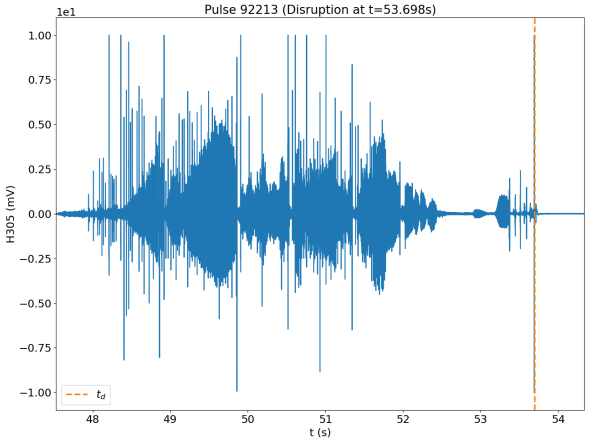

Figure 1: Raw data from the H305 coil. The dashed orange line indicates the disruption instant.

Of equal importance to this work is the measurement and use of the locked mode amplitude. This measurement is possible due to a set of saddle flux loops, which are located in the same poloidal plane at different radial positions, outside the vacuum vessel. According to [13], the amplitude of a $n = 1$ locked mode, B_{LM}, can be measured from the combination of magnetic field fluxes generated by various current sources of the setup. Examples of the locked mode signal from both disruptive and non-disruptive discharges can be seen in figure 2. An increase in the locked mode amplitude at around $t = 53.40$ s in discharge 92213 is noticed. This increase is relevant in threshold-based methods to trigger an alarm.

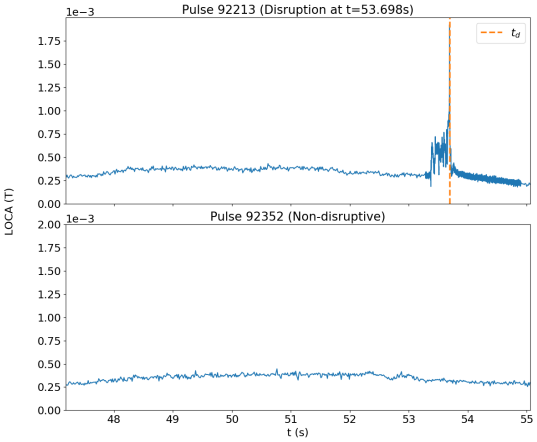

Figure 2: Example of the locked mode signal in a disruptive (top) and non-disruptive (bottom) discharge.

The key instrument to analyze MHD activity is the spectrogram, a time and frequency representation that results from Fourier analysis applied to the data coming from the H305 coil. More specifically, the used spec-

trogram is calculated using a discrete-time implementation of the short-time Fourier transform (STFT). With the STFT, an input signal is divided into N segments of equal length, by multiplying it with a given moving window function. A Fast Fourier transform (FFT) is applied within each segment. The result of the STFT can be represented as a 2D matrix, $S_{[t,\nu]}$, taking into account the global contribution of each segment and providing a local identification of FFT amplitude $s_{t,\nu}$ at each time t and frequency ν indexes, as seen in equation 4,

$$S_{[t,\nu]} = \begin{pmatrix} s_{1,1} & s_{1,2} & \cdots & s_{1,j} \\ s_{2,1} & s_{2,2} & \cdots & s_{2,j} \\ \vdots & \vdots & \ddots & \vdots \\ s_{i,1} & s_{i,2} & \cdots & s_{i,j} \end{pmatrix} \tag{4}$$

where $i, j \in (t, f)$. To compute the STFT, a segment length of 512 points was chosen, the same value of overlap points. The resolution for this segment length is equal to 4.096 ms.

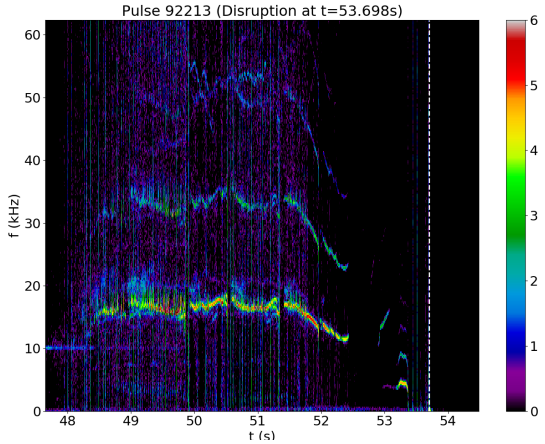

Figure 3: Spectrogram of discharge 92213. The dashed white line indicates the disruption instant.

The spectogram representation of figure 1 can be seen in figure 3. The disruption instants can be well characterized as a sequence of one or more bursts, well defined in time and covering almost the full spectrum of frequencies. Regarding the observation of MHD activity, the dynamics of mode-locking can be visible within this discharge. At $t = 51.5$ s, a $n = 1$ internal kink mode over a $q = 1$ surface starts to decrease in frequency, corresponding to the plasma deceleration. After an interruption in the MHD activity, a $n = 1$, $m = 2$ mode ($q = 2$ surface) appears at a lower frequency after $t = 53.0$ s. Its frequency is reduced and the mode eventually locks at around $t = 53.4$ s.

4. Methodology

The used data was comprised of a total of 486 pulses from JET baseline scenario [14]. According to the prior database definition of disruption time, there are 291 non-disruptive and 195 disruptive pulses. However, as the first step of this work, the study of a direct link between disruptive pulses and the occurrence of mode-locking was made, similarly to other studies, thus leading towards a new definition of a disruptive experiment in the context of this work.

The first goal was to understand how can we split the available database between disruptive and non-disruptive experiments by finding an optimal threshold value to the locked mode amplitude. This initial study was framed as a binary classification problem using machine learning standards, however, in this stage, instead of having a trained classifier to make the predictions for our data, we attribute in advance the classification labels to each experiment. That is, each experiment was labeled according to locked mode amplitude and its maximum value by using the following definitions:

- if the pulse is disruptive and the locked mode amplitude exceeds the threshold before the disruption time, then the pulse is classified as a true positive (TP);
- if the pulse is non-disruptive and the locked mode amplitude does not exceed the threshold at any point in time, then the pulse is classified as a true negative (TN);
- if the pulse is non-disruptive and the locked mode amplitude exceeds the threshold at any point in time, then the pulse is classified as a false positive (FP);
- if the pulse is disruptive and the locked mode amplitude does not exceed the threshold at any point in time or, if it exceeds, it does so on or after the disruption time, then the pulse is classified as a false negative (FN).

Deciding the optimal threshold is the same as obtaining the best combination of previously chosen performance metrics for binary classification.

When the binary classification metrics were applied according to the previous labels to our data, an optimal threshold value of 1.069×10^{-3} T was obtained. The fact that most of these metrics are above 95% (see Figure 5) shows a quite satisfactory discrimination between non-disruptive and disruptive experiments. This value is also in line with previous studies on mitigation systems that use the locked mode amplitude [2]. The obtained performance metrics values for the optimal threshold can be seen in table 1. These first obtained results were used to elaborate the following step, in which a binary

Table 1: Performance results for the optimal B_{LM} threshold.

Metric	Result
Accuracy	0.979
Balanced accuracy	0.977
F1-Score	0.974
Matthews correlation coefficient	0.957

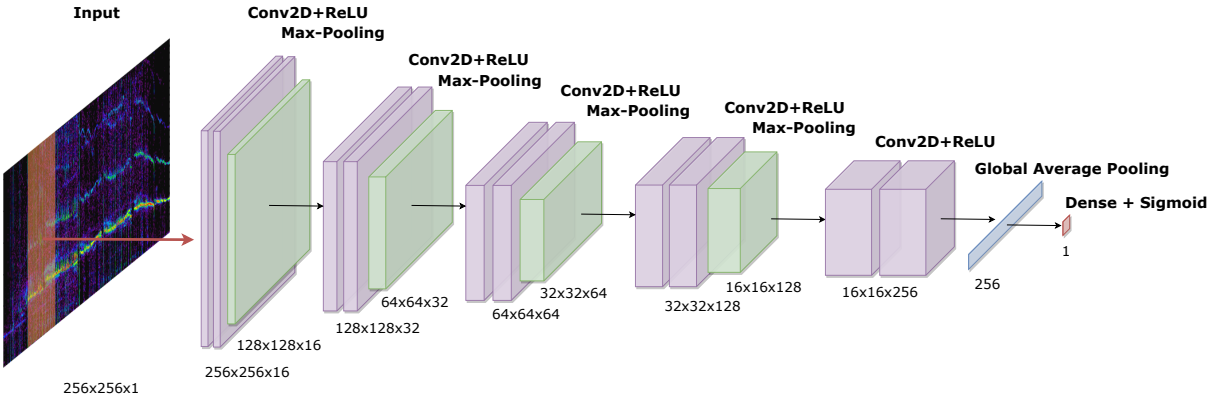

Figure 4: Proposed CNN model.

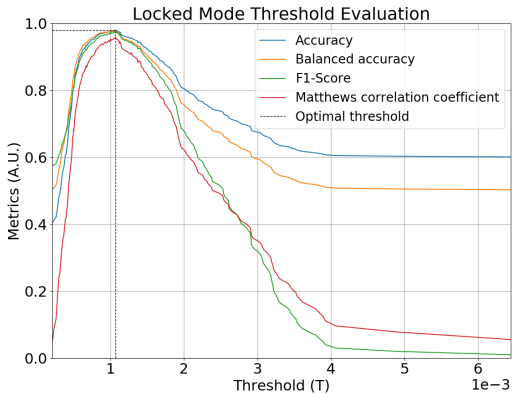

Figure 5: Binary classification metrics with a threshold range on the locked mode amplitude.

classifier based on a deep learning model is developed to receive a sample of an MHD spectrogram and assign as 0 if the locked mode amplitude does not exceed the optimal threshold, and 1 if otherwise.

The proposed deep learning model (Figure 4) is based on convolutional neural networks (CNN). These are a specific class of neural networks, particularly adapted to deal with image classification tasks and other 2D input data [15, 16]. As indicated in the red region, at the beginning of the network, the input consists of sample windows from the spectrogram (previously computed from the Fourier analysis) of 256 to 256 points in both the time and frequency axis. This means that it covers approximately 1.049 seconds in the time axis and the whole frequency axis. In each 2D convolutional layer, the two initial values refer to the width and height of the resultant feature maps, while in each max-pooling layer it refers to each feature map size after downsampling. The third term is the number of applied kernels to the input in the convolutional product and, for instance, the number of obtained feature maps.

Following the input layer, the model is composed of five convolutional blocks, each block with two convolutional layers. A rectified linear unit (ReLU) activation function is used within these layers to positively normalize its inputs and to prevent the full stack of neurons to be activated [17]. A max-pooling layer is placed after the two convolutional layers, except at the last block. Due to the implementation of the interpretation technique, the max-pooling layer in the final block is substituted with a global average pooling layer. To conclude, the resulting vector from the previous layer is fed to a one unit dense layer with a sigmoid activation function to provide a probability value between 0 and 1.

The model is trained to discriminate two classes, thus the binary cross-entropy L function is applied. Each parameter is updated by using a variation of Gradient Descent algorithm, called Adaptive Momentum Estimator (Adam) [18].

The model was trained in samples of 485 experiments. The inputs were sampled anywhere on the time axis of the spectrogram. 90% of the total discharges were used for training the model, and 10% for validation. The train was made using a batch size of 437 and a learning rate of 1×10^{-4}. The inputs were sampled anywhere on the time axis of the spectrogram, i.e., at the beginning of the experiment, where there is not any visible evidence of the possible outcome, and sometimes only a very low frequency regime with no MHD activity, or after the disruption instant in a disruptive pulse, if the locked mode has already occurred. This was also done in order to any possible bias in the model towards a desired outcome.

To obtain the physical interpretation from the classifier and relevant MHD activity to the model, the class activation mapping (CAM) technique [19] is used. With CAM, we identify the regions in the spectrogram that contribute the most to the classification of a determined image by the CNN model. To employ CAM, one must change the CNN model according to figure 6.

Both the global average pooling layer and the dense layer are removed and replaced with a single custom layer that computes the weighted sum of 16×16 feature maps produced by the last convolutional layer. In order to overlay the map on the input sample and directly observe the results, it is also upsampled to the same dimensions of the input (256×256).

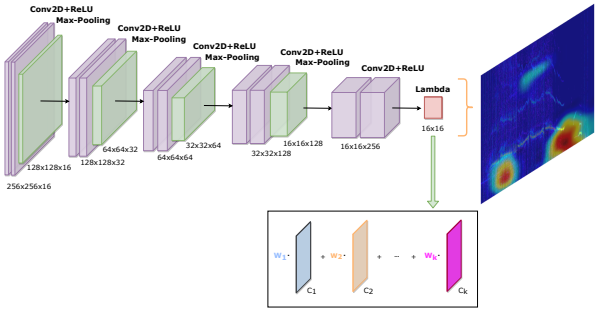

Figure 6: CNN model after the introduction of CAM.

5. Results and Discussion

The results of optimal training, after many trials, are shown in figure 7. At epoch 357, the best value of validation L is achieved (0.38), reaching a validation accuracy of 0.83. After achieving the global minimum, it is clear the development of overfitting.

Figure 7: L function and accuracy metrics obtained during training.

These results could be attributed to the nature of both the input data and the method applied in training. As previously mentioned, no constraints were set when choosing a given sample in training, i.e., they could be selected at any time of the discharge. This means that, at the beginning of the discharge, the chosen sample is identical between non-disruptive and disruptive cases. Thus, if the behavior in frequency is practically the same in two samples, even with different labels, the model can behave very well on training data, but wrongly guess in test data.

With the model fitted in training data, we use discharge 92213 to see the prediction results. Each spectrogram sample window, at a given time index, was used to make a prediction in the next window, from the beginning of the experiment to the last sample. That is, the sample window is positioned always to the past in time from which the prediction is made. The result is in figure 8. At $t = 52.2$ s, the prediction continuously increases to almost 1, and practically stabilizes during MHD activity interruption (from $t = 52.5$ s to $t = 53.0$ s), mode-locking (from $t = 53.2$ s to $t = 53.4$ s) and the disruption instant. In practice, the model was able

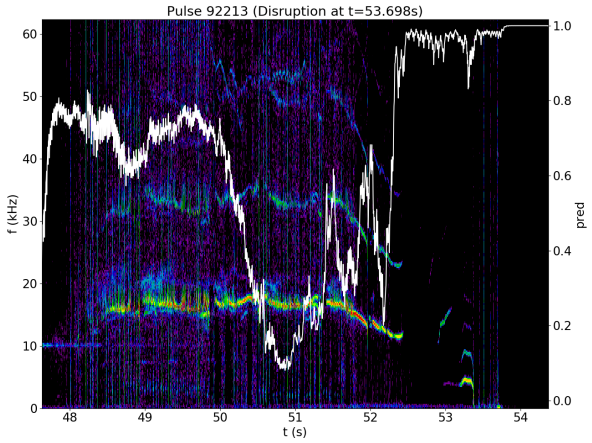

Figure 8: Model prediction (white) for the spectrogram samples of 92213.

to recognize a disruption caused by mode-locking even before the latter shows up, with a time window of approximately 1 s.

One of the relevant tasks when developing binary classifiers is to test them in a new set of data and evaluate its performance metrics to measure its capability to generalize. However, all of the available data was used during training of the model, and thus a test set with new experiments could not be considered in this work. The validation test was used as a test set equivalent to elaborate this task, where 48 discharges were included. Similarly to what was done in the study of the locked mode amplitude, if the probability does not reach a defined threshold value when the discharge has a locked mode disruption, then the discharge is classified as FN, and if it reaches that threshold in a 0 labeled sample is classified as FP. The same performance metrics used in the study of the locked mode threshold were also used, with the addition of the receiver operating characteristic (ROC) curve as seen in figure 9.

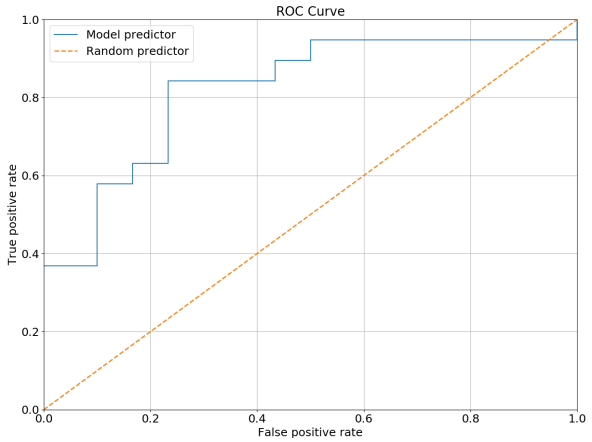

Figure 9: ROC curve of the developed predictor.

At the probability threshold value of 0.883 all the metrics reach their maximum, with an accuracy of about 79.5% (see Table 2). The reason for this considerable high value to be obtained can be attributed to

fluctuations present during the predictions at both non-disruptive and disruptive discharges. The ROC curve indicates that the predictor is reasonably far from a random classifier, which is a good indication, confirmed by the area under curve (AUC) calculation of about 0.82, despite the considerable number of FP discharges.

Table 2: Best performance metrics achieved for a probability threshold of 0.883.

Metric	Result
Accuracy	0.796
Balanced accuracy	0.804
F1-Score	0.762
Matthews correlation coefficient	0.594
AUC	0.819

After the evaluation of the performance metrics, we used the modified model (see Figure 6) to produce the CAM for a given sample. The negative values (as blue color) represent the areas of the spectrogram that contributed the most to the classification of the classifier towards 0, which in the case of a more intense blue indicates the saturation of the sigmoid function to its minimum value. For positive values (as red color), the areas in the spectrogram had a higher influence towards a classification of 1. The intermediate value (as white color) is equivalent to a random guess of the classifier. The resultant map is overlaid on the spectrogram sample to facilitate further interpretations. The application to discharge 92213 can be seen in figure 10.

The right-hand side of the image corresponds to the CAM method when applied to the sample highlighted by the yellow rectangle on the spectrogram at the left side. It can be observed what most influenced the predictor to make a certain classification. In the case of discharge 92213, there are features that converge with the physical interpretation given by the previous analysis of the spectrogram. First, it is interesting to notice the highlight on the interruption of the internal kink mode, specifically at $t = 52.4\,\mathrm{s}$. This highlight by CAM is sufficient to considerably raise the prediction value to approximately 1, about one second before mode-locking development. Secondly, two additional characteristics can be seen by moving the sample window. A negative area, where practically no MHD activity is present (only a residual frequency band), and the positive highlight in a $m = 2$, $n = 1$ mode ($q = 2$ surface), at a lower frequency (approximately 4 kHz), which decreases due to locking afterwords. This kind of behavior is consistent with the work by Sweeney et al. [20], where these modes, namely rotating $m/n = 2/1$ modes, are likely to lock. In Pucella et al. [21], this is also observed and correlated with a growth of a $m/n = 2/1$ islands due to temperature hollowing, in which some impurities start to accumulate in the plasma core.

The CAM method was applied to practically all discharges in the dataset. In 22 discharges, the same pattern described previously is explicitly observed, that is, the interruption of the MHD activity and its resurgence, followed by the locked mode. All these discharges surpass the prediction threshold of the model, thus being classified as disruptive, which is consistent with the prior attributed labels. In some cases, either the prediction value or the positive highlight of CAM is before mode-locking and could provide a sooner warning time than a simple threshold on the locked mode amplitude.

However, these conclusions need to be carefully constrained. In some of the identified discharges, the CAM method also gives a particular focus to the described patterns, but it may not increase considerably the probability value in those windows, as previous MHD activity could enhance it. Additionally, the presence of many FP discharges when the calculations of the performance metrics were made can indicate that the instability of the probability values throughout each discharge makes the model wrongly classify some of the non-disruptive discharges as disruptive. To conclude, some discharges could be left out of the CAM behavior identification due to the short duration of the locked mode.

6. Conclusions

In this work, we have shown a deep learning method to be applied directly on spectrogram analysis of MHD activity to predict the occurrence of disruptions with mode-locking. Although it cannot be directly compared with other predictors, due to the lower accuracy value and the fact that no test set was used to measure the performance metrics, it can reasonably discriminate the classes of interest, having into account that only the data coming from a magnetic coil was used as an input. Additionally, the CAM method was implemented to allow the addition of physical understanding of the results and retrieve relevant MHD activity in the perspective of the model.

Based on the application of CAM mapping to various discharges, the most common feature to be compatible with available physical insight is the interruption of MHD activity, followed by its resurgence. In some discharges, the model can provide an alarm before mode-locking itself, and it can even surpass current predictors in the time window for prediction.

To sum up, this work delivers a new technique to analyze MHD activity in time and frequency representations, as well as provides some insight when data-driven models are applied in nuclear fusion data. Furthermore, it contributes to the new paradigm in the disruption prediction field, where besides the capabilities of the predictors to correctly classify disruptive discharges, it allows to increase the knowledge in the mechanisms involved, as well as to validate the results from the applied deep learning models in fusion data. The introduction of this methodology and other interpretable deep learning frameworks can be an important analysis tool for tokamak physicists, allowing a broader perspective of data-

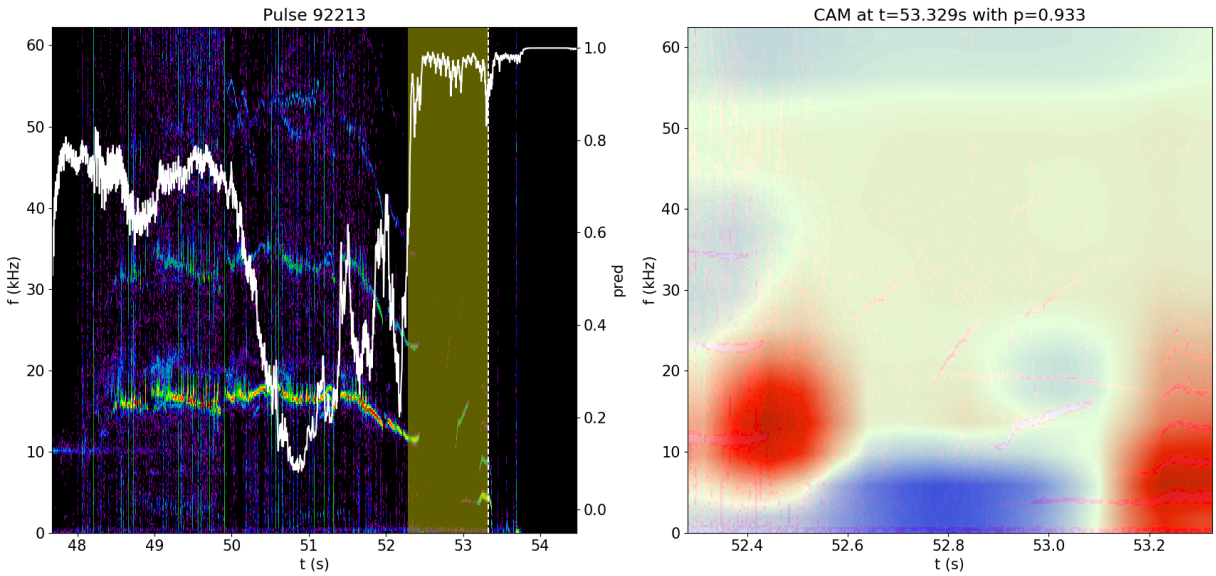

Figure 10: Result of CAM for discharge 92213 at $t = 53.32$ s.

driven models and providing the learned features from the models to extend the analysis of MHD activity and other machine diagnostics.

As future work, a new set of possible additional tasks arise after the first approach with these methodologies. The parameter and hyperparameter spaces were not completely explored, and a more extensive tuning, with proper computational resources, could lead to improvements in the model's accuracy. One could also establish a new type of labeling and classification constraints, since using binary labels can be a limitation specifically when samples are at the beginning of the discharge. Finally, it would be interesting to examine other behaviors retrieved by CAM in our data, or add other interpretability techniques, which could bring additional insight.

7. Acknowledgments

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement No 633053. IPFN (Instituto de Plasmas e Fusão Nuclear) received financial support from FCT (Fundação para a Ciência e a Tecnologia) through projects UIDB/50010/2020 and UIDP/50010/2020. The author is thankful for the granted use of computational resources at CCFE (Culham Centre for Fusion Energy) in the UK, as well as to NVIDIA Corporation for the donation of hardware that was used to this project.

References

[1] P. C. De Vries et al., "Survey of disruption causes at JET," *Nucl. Fusion*, vol. 51, p. 053018, 2011.

[2] C. Reaux et al., "Use of the disruption mitigation valve in closed loop for routine protection at JET," *Fusion Engineering and Design*, vol. 88, p. 1101, 2013.

[3] P. C. De Vries et al., "Scaling of the MHD perturbation amplitude required to trigger a disruption and predictions for ITER," *Nucl. Fusion*, vol. 56, p. 026007, 2015.

[4] J. Vega et al., "Disruption precursor detection: Combining the time and frequency domains," in *2015 IEEE 26th Symposium on Fusion Engineering*, pp. 1–8, 2015.

[5] J. Kates-Harbeck, A. Svyatkovskiy, and W. Tang, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning," *Nature*, vol. 568, pp. 526–531, 2019.

[6] B. Cannas et al., "Support vector machines for disruption prediction and novelty detection at JET," *Fusion Engineering and Design*, vol. 82, pp. 1124–1130, 2007.

[7] J. M. Lopez et al., "Implementation of the disruption predictor APODIS in JET's real-time network using the MARTe framework," *IEEE Transactions on Nuclear Science*, vol. 61, pp. 741–744, 2014.

[8] R. M. Churchill et al., "Deep convolutional neural networks for multi- scale time-series classification and application to tokamak disruption prediction using raw, high temporal resolution diagnostic data," *Physics of Plasmas*, vol. 27, p. 062510, 2020.

[9] D. Batchelor et al., "Simulation of wave interactions with MHD," *Journal of Physics: Conference Series*, vol. 125, 2008.

[10] C. C. Hegna, "The physics of neoclassical magnetohydrodynamic tearing modes," *Physics of Plasmas*, vol. 5, no. 5, pp. 1767–1774, 1998.

[11] R. Fitzpatrick, "Interaction of tearing modes with external structures in cylindrical geometry (plasma)," *Nucl. Fusion*, vol. 33, pp. 1049–1084, 1993.

[12] I. H. Hutchinson, *Principles of Plasma Diagnostics*. Cambridge University Press, 2 ed., 2002.

[13] G. Sias et. al., "A locked mode indicator for disruption prediction on JET and ASDEX upgrade," *Fusion Engineering and Design*, vol. 138, pp. 254–266, 2019.

[14] L. Garzotti et al., "Scenario development for D–T operation at JET," *Nucl. Fusion*, vol. 59, p. 076037, 2019.

[15] M. Browne and S. Ghidary, "Convolutional neural networks for image processing: An application in robot vision," in *Australian Conference on Artificial Intelligence*, pp. 641–652, 2003.

[16] M. Dörfler, R. Bammer, and T. Grill, "Inside the spectrogram: Convolutional neural networks in audio processing," in *2017 International Conference on Sampling Theory and Applications*, pp. 152–155, 2017.

[17] H. Ide and T. Kurita, "Improvement of learning for CNN with ReLU activation by sparse regularization," in *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 2684–2691, 2017.

[18] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," 2015. arXiv:1412.6980.

[19] B. Zhou et al., "Learning Deep Features for Discriminative Localization," 2015. arXiv:1512.04150.

[20] R. Sweeney et al., "Statistical analysis of m/n = 2/1 locked and quasi-stationary modes with rotating precursors at DIII-D," *Nucl. Fusion*, vol. 57, p. 016019, 2017.

[21] G. Pucella et al., "Onset of tearing modes in plasma termination on JET: the role of temperature hollowing and edge cooling," *Nucl. Fusion*, vol. 61, p. 046020, 2021.