Fusion Engineering and Design 204 (2024) 114472

Contents lists available at ScienceDirect

Fusion Engineering and Design

journal homepage: www.elsevier.com/locate/fusengdes

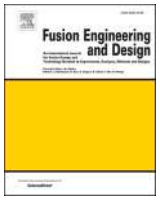

MHD spectrogram contribution to disruption prediction using Convolutional Neural Networks

E. Aymerich, G. Sias *, S. Atzeni, F. Pisano, B. Cannas, A. Fanni, the JET Contributors[1], WPTE Team[2]

Dept. of Electrical and Electronic Engineering, University of Cagliari, Cagliari, Italy

ARTICLE INFO

Keywords:
Deep learning
Disruption prediction
Disruption avoidance
Spectrogram
Mirnov coils

ABSTRACT

The present research focuses on investigating deep neural networks techniques for predicting plasma disruptions in tokamaks. For this purpose, various deep-learning predictive models and several plasma diagnostics will be analyzed, using data gathered during the experimental campaigns conducted at the JET nuclear fusion tokamak between 2011 and 2020. The primary objective of this study is to show that the contribution of the MHD spectrograms increases the disruption predictor performance. The final deep-learning prediction model leverages the capability of Convolutional Neural Networks to directly learn important spatiotemporal information from 1D plasma profiles of temperature, density and power radiation, obtained from High Resolution Thomson Scattering and Bolometer diagnostics, as well as from spectrograms generated by a set of fast magnetic pick-up coils known as Mirnov coils. The Convolutional Neural Network eliminates the need for manual feature extraction methods that characterize the majority of machine learning methods. Using plasma profiles information allows to distinguish between core radiation caused by impurity accumulations and outboard radiation phenomena. Likewise, the decision to incorporate spectrograms from Mirnov coils is based on the diagnostic's ability to measure magnetic fluctuations originating from MHD instabilities, which can lead to disruptions. In order to address phenomena characterized by fast temporal dynamics, the inclusion of the locked mode signal was chosen. This signal is commonly employed at JET to trigger mitigation actions. It is integrated into an alarm scheme that employs both AND/OR logic and optimized thresholds, ensuring its effectiveness. The proposed predictor exhibits significant performance, with only one missed alarm out of 92 disrupted discharges and three false alarms out of 131 regularly terminated discharges in the test set.

1. Introduction

Disruptions pose a significant challenge in the realm of tokamak-based fusion reactors [1,2,3]. The disruptions involve a sudden and extensive loss of plasma stability and confinement within the tokamak device, leading to potential harm to the reactor components and safety concerns [1,2,3].

Extensive efforts are currently underway to develop techniques and strategies for mitigating disruptions and minimizing their impact. These efforts encompass active control methods that involve the injection of gas or impurities to mitigate plasma instabilities. Additionally, improved plasma control algorithms and advanced diagnostic systems are being developed to detect and avoid disruptions in real-time [4].

In recent decades, there has been a comprehensive pursuit to develop models capable of accurately predicting disruptions in tokamaks. This pursuit primarily revolves around utilizing machine learning techniques that leverage big data obtained from various plasma diagnostics during numerous experimental campaigns as shortly referenced in the following. Cannas [5] uses support vector machines (SVM) to predict disruptions and also includes a novelty detection method to assess the reliability of the predictor output. In [6] and [7] SVM is the core of a disruption predictor installed in the JET real-time network and progressively improved. Furthermore, Zheng et al. [8] describe a hybrid neural network structure that significantly enhances the prediction performance of density limit disruptions on the J-TEXT tokamak. Churchill et al. [9] discuss the utilization of deep convolutional neural

* Corresponding author.
E-mail address: giuliana.sias@diee.unica.it (G. Sias).
[1] See the author list of 'Overview of JET results for optimizing ITER operation' by J. Mailloux et al. 2022 Nucl. Fusion 62 042026
[2] See the author list of E. Joffrin et al. Nucl. Fusion, 29th FEC Proceeding (2023)

https://doi.org/10.1016/j.fusengdes.2024.114472
Received 25 January 2024; Received in revised form 18 April 2024; Accepted 22 April 2024
Available online 17 May 2024

networks for multi-scale time-series classification, applied specifically to disruption prediction in the DIII-D tokamak using raw, high temporal resolution, diagnostic data. In 2019 the authors of [10] provide an overview of machine learning algorithms employed for disruption prediction and classification at JET. Also references [11,12,13,14,15] delve into the application of deep learning for disruption prediction in several machines. In [11], the authors use a Long-Short Term Memory (LSTM) network to predict disruptions by inputting plasma radiation profiles. In [12] the authors applied Convolutional Neural Network (CNN) and LSTM to HL-2A tokamak, whereas LSTMs have been employed in both the ADITYA tokamak [14] and EAST [13,15] for the purpose of generating disruption alert outputs.

Some predictive models have also been proposed that extend beyond single-machine approaches, encompassing cross-machine approaches. Rea et al. [16] establish the foundation for comparing data-driven disruption prediction algorithms across DIII-D and JET. Zhu et al. [17] introduce a novel deep-learning disruption prediction algorithm that achieves high accuracy across multiple tokamaks with limited hyperparameter tuning. In [18], the authors employ a hybrid architecture, blending convolutional components with LSTM network, for cross-machine disruption prediction using JET and DIII-D data. The CNN extracts features from electron temperature and density profiles, and after combining output signals, the LSTM predicts the disruption risk, considering factors such as locked mode amplitude.

All these papers have aimed their efforts towards enhancing the performance of data-driven predictive models, from rudimentary single layer neural network-based approaches to advanced and high-performing deep learning methods, while also keeping pace with the advancements in available computing power.

The findings of all these studies suggest that machine learning algorithms hold significant promise as powerful tools for predicting and mitigating disruptions across various tokamaks.

However, incorporating physics-based disruption markers into data-driven algorithms shows promise in establishing a unified framework for disruption prediction and interpretation across different tokamaks. In fact, in a complex field like fusion plasma control, extracting valuable and reliable information from physics quantities, which either represent or have a direct connection to observable variables that can be manipulated to influence the plasma state, can significantly enhance any predictive model. Consequently, the breakthrough in performance occurred when the invaluable insights from physics were utilized to construct the predictive models. In [19], 1D diagnostics have been considered, such as plasma temperature, density, and radiation, to extend the time interval before the disruption, the so-called warning time, in which reliable avoidance actions can be made. While profile-based indicators are just one aspect contributing to plasma performance and stability, they appear to align with the objective of obtaining useful and robust physics-based information to be used as input to the disruption prediction models.

Recent studies have shown that the prediction performance is significantly influenced by both the suitable selection of appropriate diagnostics and careful feature extraction methods [20,21]. The method proposed in [20] introduces a novel approach by employing CNN to process the spatiotemporal information extracted from 1D plasma profiles. This technique involves converting the profiles into images, which are better suited for processing by the CNN. Additionally, it incorporates automatic detection of the pre-disruptive phase of disruptions for selecting the training data.

This research has prompted the exploration of additional diagnostic information that may hold insights into potential precursors of disruptions. By appropriately processing this information, it can be transformed into image representations and subsequently utilized as inputs to CNN models.

Among these diagnostics, magnetic ones are crucial in comprehending plasma physics, ensuring control, and enabling safe operation within a tokamak device. For example, within JET plasmas, a wide range of MHD (Magnetohydrodynamic) instabilities result in magnetic fluctuations and the device is indeed equipped with fast magnetic Mirnov coils for spectral and mode number analysis. In fact, the use of the Mirnov coils [22] allows us to detect the amplitude fluctuations in the magnetic field while the plasma rotates within the torus, providing valuable information about the oscillation modes of the plasma [23]. The frequency of these fluctuations is particularly significant as it helps determine the specific modes of oscillation [24,25]. Certain modes have been identified as harmful to plasma confinement and can potentially lead to severe disruptions [25,26,27].

In the present paper, the effectiveness of MHD analysis for disruption prediction is assessed. Specifically, a model that incorporates the MHD spectrograms, as in [28], together with other 0D and 1D plasma parameters, is proposed. More specifically, the proposal involves using CNNs to extract spatiotemporal features from Mirnov spectrograms as in [29] and appropriately preprocessed images of plasma profiles, including temperature, density, and radiated power. A straightforward alarm system combines the output of the CNN disruption prediction model with the locked mode diagnostic signal, which is commonly employed in plasma mitigation procedures [27,30,31,32].

In order to facilitate the comparison with recent literature, the CNN predictor was trained using data from experimental campaigns conducted at JET from 2011 to 2013. To assess its predictive capabilities, its performance was evaluated on disrupted and regularly terminated discharges spanning a decade of JET experimental campaigns, from 2011 to 2020, as in [20]. This evaluation not only validates the algorithm's robustness, even in the presence of substantial operational condition variations across different campaigns, but also highlights the favorable contribution of the MHD spectrograms.

This paper is organized as follows. Section 2 provides a description of the high-resolution array of Mirnov coils and outlines the steps taken to construct the spectrograms. In Section 3, the procedure employed to transform 1D plasma profiles of temperature, density, and radiation into 2D images is detailed. Section 4 provides a comprehensive explanation of the database utilized for training, optimizing, and validating the proposed deep-learning disruption predictor. In Section 5, the indicators utilized to evaluate the performance of predictive models are introduced. Section 6 presents the fundamental principles of the CNN prediction models, while Section 7 introduces the architecture of the proposed disruption predictor. Section 8 presents a comprehensive analysis of the predictor's output, including a comparison with recent literature using the same test set. Finally, Section 9 offers concluding remarks.

2. High-resolution Array Coils (Mirnov coils)

JET utilizes multiple diagnostics to monitor the MHD activity of the plasma and to detect possible anomalies [25]. The locked mode signal, which is one of the most commonly used precursors for disruption detection, is acquired using diagnostics like pickup coils, Mirnov coils, or other magnetic field sensors [27,32]. Following an analysis of the locked mode signal and its impact on plasma stability, a decision is made to initiate appropriate mitigation measures. Typically, the response to a disruption involves the application of a threshold to the amplitude of the locked mode. The measurement of the locked mode amplitude on JET is carried out using a set of 2×4 saddle flux loops. The locked mode amplitude, which represents the dominant odd $n = 1$ mode, is determined by analyzing the signals from these flux loops through various processes [33]. The resulting signal is expressed in Tesla units. The warning time provided by thresholding this signal is frequently insufficient for avoiding the disruption and, when a locked mode occurs and becomes detrimental to plasma stability, only mitigation action can be performed.

By measuring fluctuations in the magnetic field, researchers can obtain valuable information about plasma dynamics and explore the influence of various plasma parameters and experimental conditions.

Magnetic fluctuations coming from MHD instabilities can be analyzed also using the data from a set of fast magnetic pickup coils, the so called Mirnov coils, designed for high frequency MHD activity study (up to 500 kHz) [34,35,36]. They consist of a set of 18 coils, distributed at different poloidal and toroidal location. Among them, 8 High-Resolution Array Coils (HRACs) are basically used for high toroidal (n) and poloidal (m) mode analysis, they are located outboard of the plasma, distributed toroidally and poloidally at Octant 3 [34], as shown in Fig. 1.

Among the HRACs, the coil H302 has been considered here, having a sampling rate of 0.5 μs. When the signal from H302 was unavailable, coil H305 was considered due to its closeness to H302 and having the same sampling rate. In order to restrict the frequency analysis to the kilohertz range, the signal from the Mirnov coil has been resampled at a rate of 8μs. The spectrogram for each pulse was generated using a Short-Time Fourier Transform (STFT) [37], employing a time window of 2.048ms with a 48μs overlap. The potential of the information derived from the analysis of MHD activity is depicted in Fig. 2. The upper plot displays the plasma current (red line) and the locked mode signal (blue line) of the JET pulse #96471, the middle plot shows the time trace of the H302 coil signal, while the bottom plot presents the corresponding spectrogram obtained using the previously described method.

In the bottom plot, the spectrogram of H302 signal for the JET pulse #96471 is depicted, which exhibits a disruption at time $t = 12.4\ s$. It can be observed that at time $t = 11.6$, the plasma begins to decelerate, displaying progressively lower frequencies. This phenomenon coincides with the moment when the external heating systems, responsible for maintaining the plasma at extremely high temperatures, were turned off. The shutdown of these systems resulted in a loss of momentum within the plasma, causing it to rotate at an increasingly slower pace, leading to a sudden cessation of oscillation modes at time $t = 12.1\ s$ and subsequent disruption. Hence, before locking, the H302 signal increases in amplitude, and a decrease in frequency can be easily observed by the spectrogram until the oscillation ends, resulting in the definitive locking of the mode. The Mirnov coil measures embed both the information about the rotation frequency and the amplitude of the MHD modes present in the plasma, thereby detecting the formation of islands during the discharge. Comparing the plots in Fig. 2, it is noticeable that at time $t = 12.1s$, there is an increase also in the amplitude of the locked mode signal, which corresponds to the moment when the rotation of the plasma, analyzed in the spectrogram of the same Fig. 2, ceased. However, the warning time obtainable by thresholding the amplitude of the locked mode signal is limited to 300ms, while the spectrogram allows to detect the slowing down of the mode before its locking. Hence, the spectrogram serves as an informative input for predictive models, offering valuable insights. Additionally, its representation as a 2D image makes it particularly compatible with CNNs, which excel at extracting meaningful details from images.

3. Plasma profiles

Although it has been proven that there is a strong correlation between locked modes and plasma disruptions, not all disruptions are caused by intense MHD activity or the locked mode itself [3]. Hence, to construct a disruption predictor, incorporating additional diagnostics associated with plasma parameters becomes imperative. Among the most commonly used diagnostics in the literature [11,20], temperature, density, and radiation profiles contain valuable information for early identification of events that could lead to plasma disruptions.

At JET, High-Resolution Thompson Scattering (HRTS) diagnostics have been used to obtain profiles of electron temperature (Te) and electron density (ne) by up to 63 lines of sight along the outer radius of the plasma (2.9–3.9 m) at 20 Hz. Data regarding the profile of radiated power is obtained from bolometer cameras. In JET, there are two cameras, one vertical and one horizontal. These cameras, through lines of sight, cover the entire analysed plasma section, allowing for the calculation of electromagnetic radiation at each point. In the present work,

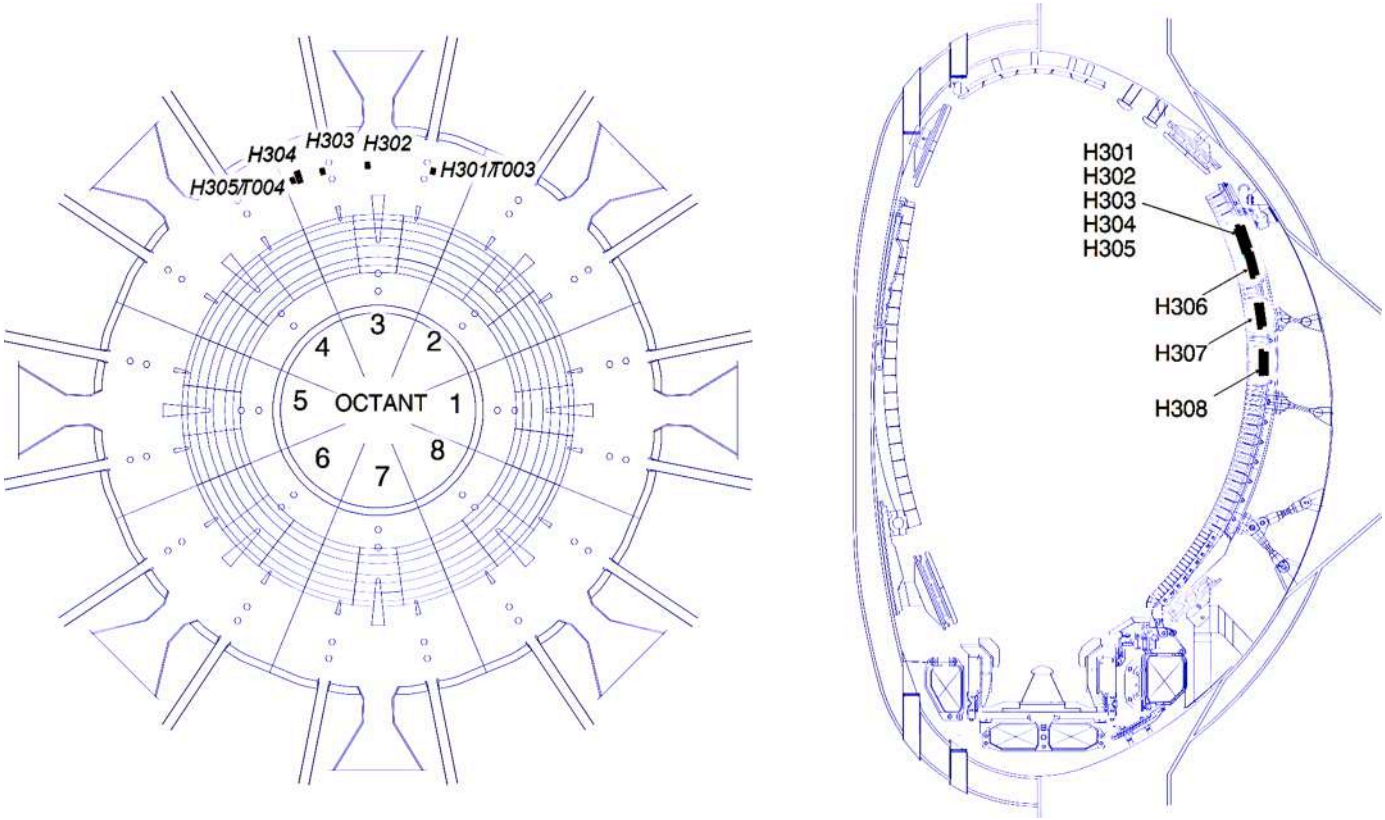

Fig. 1. Toroidal and poloidal location of the High-Resolution Array Coils at JET, named H30X, where X range from 1 to 8 [35].

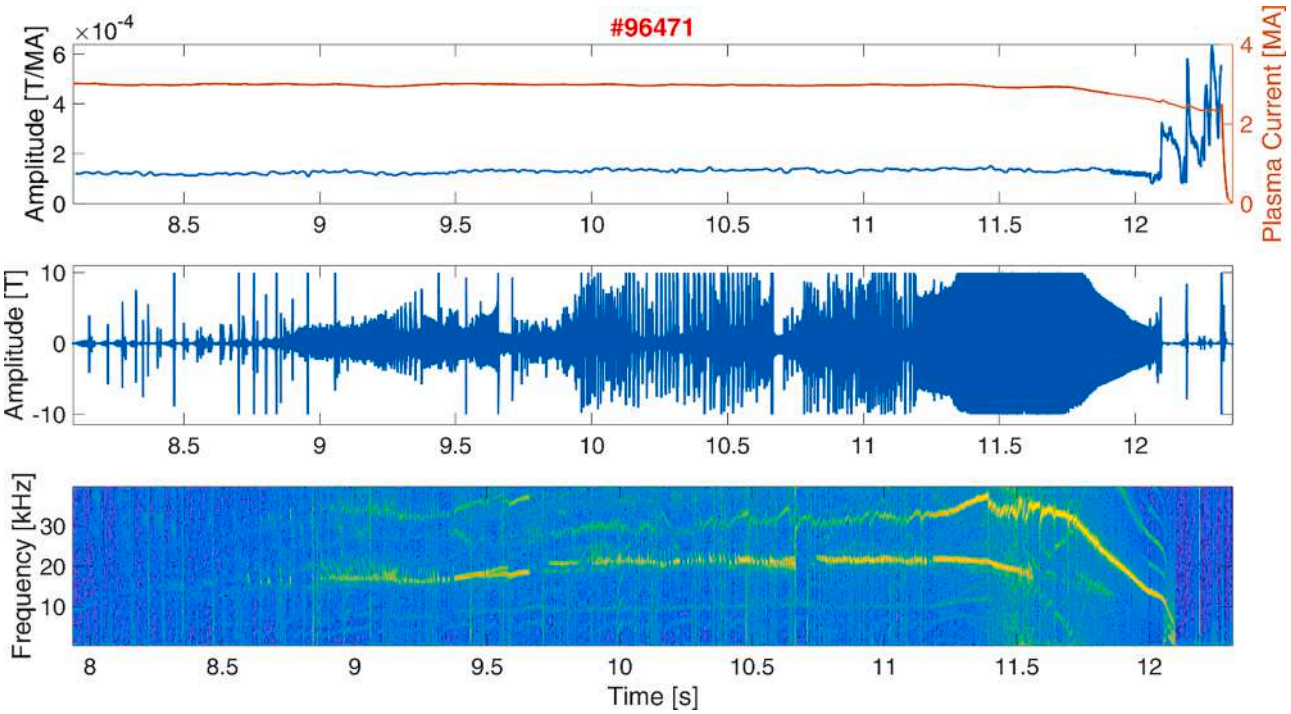

Fig. 2. From the top to the bottom: locked mode amplitude (blue) and Plasma Current (red), signal from the H302 Mirnov coil, and MHD spectrogram of the JET pulse #96471.

for the sake of comparison with literature results [19,20,38] only the lines of sight from the horizontal camera (KB5H diagnostic) are taken into account.

HRTS and bolometer generate one-dimensional (1D) signals that require proper processing to be utilized as inputs for the predictive machine learning model. In fact, even though the time scale of the thermal quench preceding the disruption is in the order of ms, the alteration of the plasma profiles may occur even seconds before the disruption [25], allowing an early detection and the adoption of avoidance actions.

A significant number of machine learning models in the disruptor prediction literature [5,16,19,39] necessitate input in the format of

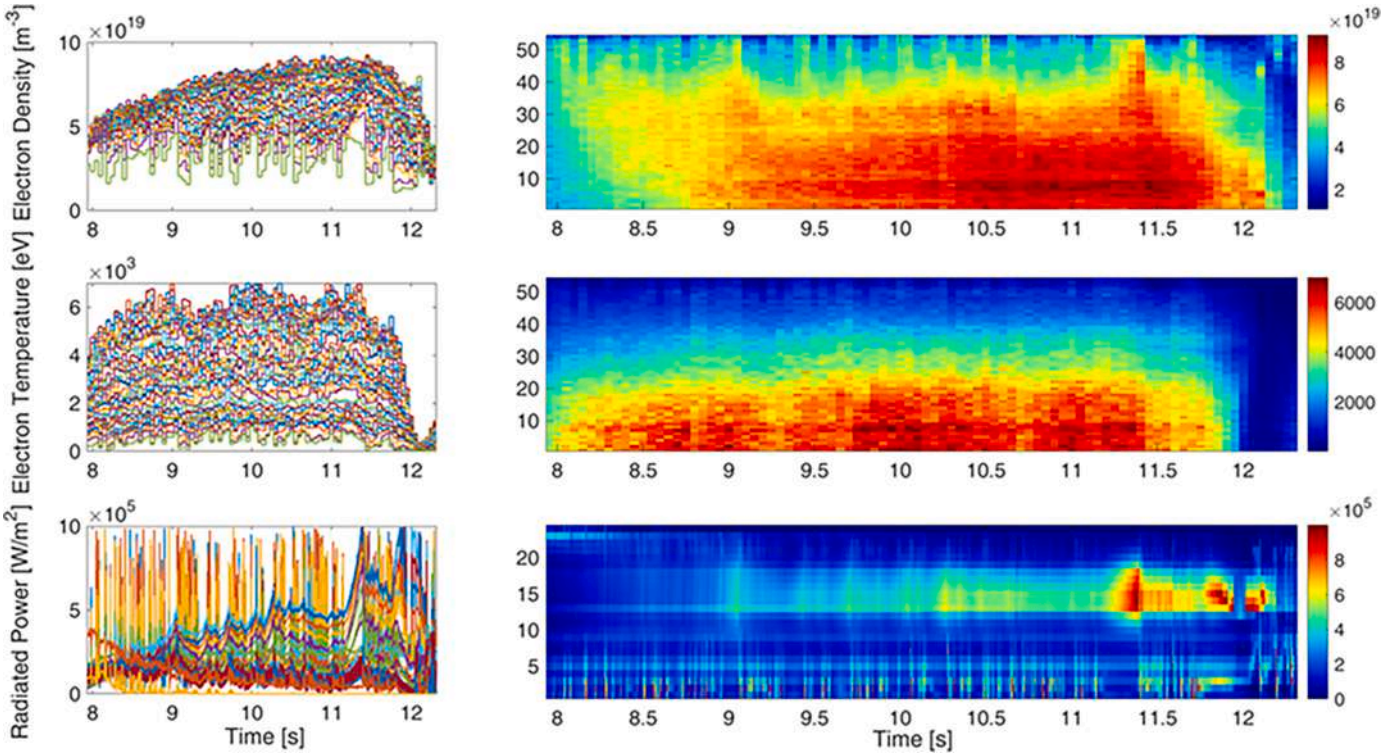

Fig. 3. JET pulse #96741: Left: The Resampled Electron Temperature and Electron Density from HRTS (upper and middle plots), and the Radiated Power from bolometer horizontal camera (bottom plot); Right: the corresponding images.

samples derived from 0D temporal signals. These signals, the so-called peaking factors, have been employed as features to investigate the interplay of their temporal evolution throughout the chain of events leading to disruption [7,16,19,25,38,40,41,42]. As an alternative, for each plasma profile, it is possible to construct a spatiotemporal matrix where each element represents the measurement value at a specific line of sight and at a specific time sample [20]. Fig. 3 illustrates, on the left side, the raw data obtained from the HRTS (top and middle subplots) and the bolometer (bottom subplot), along with the corresponding images, on the right side, specifically related to JET pulse #96471. On the right, the y-axes of the images refer to the different channels of the diagnostics, while the parameter value is expressed by the color scale. It is possible to see how at 11.4 s approximately, the plasma temperature at the core (channels #1÷10) cools down, while the core electron density is high and the density profile peaks, due to the lowering at the edge. The plasma definitely cools down at 12.1 s approximately, when the spectrogram in Fig. 2 confirms the mode locking.

4. Dataset

To train the CNN predictive model, it is necessary to have a dataset of experimental data comprising properly terminated discharges and discharges terminated with a disruption.

In order to assess the effectiveness of MHD analysis for disruption prediction, the same database used in [20] is considered for training and testing the CNN model. This allows for a comparison between the performance of the proposed predictor in this paper and a disruption predictor that does not incorporate MHD activity information. The previously proposed disruption predictor [20] takes inputs such as internal inductance, locked mode amplitude normalized by the plasma current (LM_{norm}), and temporal information from radiated power (Prad), electron temperature (Te), and electron density (ne) profiles. The dataset comprises 193 disruptive and 219 regularly terminated discharges from the JET experimental campaigns conducted between 2011 and 2020. These discharges had a flat-top plasma current exceeding 1.5 MA and a flat-top length longer than 200 ms. Pulses with unavailable Mirnov coil signals are excluded from the dataset, resulting in a smaller set of 167 disrupted pulses and 196 regularly terminated pulses, which are divided into training, validation, and testing sets as detailed in Table 1. For each chosen pulse, the flat-top starting time (t_{SFT}) was determined as the first time instant when the plasma reached the X-point configuration. In the case of disrupted pulses, the flat-top ending time (t_{end}) was defined as the time of valve activation for those terminated by Massive Gas Injection (MGI), or as the time corresponding to the drop off in core temperature and the onset of the plasma current spike (disruption time t_D) for unmitigated disruptions. The dataset does not include disruptions caused by Vertical Displacement Events.

As mentioned before, the input features are obtained from various diagnostics: the Bolometer horizontal camera, which provides 1D profiles of radiated power; the High-Resolution Thompson Scattering (HRTS), which provides electron temperature and density, and the Mirnov coils from which the spectrograms are generated. The database also includes the 0D signal of the locked mode. These diagnostics have different sampling times ranging from 10^{-4} to 10^{-2} *s*. In order to meet the real-time implementation requirement, assuming a common time base with a time interval of 2 *ms*, at each time instant the previous closest sample value of each signal is retained, without additional filtering. Additionally, preprocessing algorithms are applied to eliminate outliers and replace corrupted values in both the 1D and 0D signals, following the approach proposed in [20]. Then, the three images from plasma profiles are vertically stacked, and their ranges are normalized with respect to the signal ranges in the training set, obtaining the final image in Fig. 4. A segmentation of the final image is achieved by applying an overlapping sliding window of 200 *ms*, represented by the dashed black line in Fig. 4. Each sample in the pulse is associated to a time slice of 132×101 pixels.

Table 1
Number of pulses in the training and test sets.

	JET Campaigns	Disruptions	Regulars
Training	C28-C30	57	50
Validation	C28-C30	18	15
Test	C36, C38	92	131

In order to be included in the database, the spectrogram of the Mirnov coil signal underwent a Short-Time Fourier Transform (STFT) with a window size of 2.048*ms* with a 48μs overlap, resulting in a time resolution of 2*ms*. The frequency of the spectrogram is then cut to 40 kHz. The segmentation process generates input images of size 81 × 101 pixels.

As the CNN is a supervised algorithm, during the training a label must be assigned to each time windows (or time slice). Thus, the phase where the disruption precursors start to act must be identified for disruptive discharges. In [40], a multivariate statistical procedure was proposed to automatically estimate, for each disrupted pulse, a time instant, called $t_{pre-dis}$, that identifies the onset of disruption precursors, thus enabling the identification of a so-called "precursor phase." Thus, the time slices belonging to the regularly terminated discharges, and those coming from the disruptive ones between t_{SFT} and $t_{pre-dis}$ are labelled as "0"; whereas the time slices coming from the disruptive ones after $t_{pre-dis}$ are labelled as "1". To address the class imbalance issue caused by the varying durations of the two phases, different approaches are used for generating time slices. For regularly terminated discharges and the not disrupted phase of disruptive ones (before $t_{pre-dis}$), time slices are generated and then subsampled to obtain one image every 200 ms. Instead, for the disruptive pulse after $t_{pre-dis}$the subsampling retains one sample every 24 ms. This helps to alleviate the unbalance between the two classes. In order to adhere to real-time constraints, during the test phase, all the consecutive time slices are presented to the CNN.

5. Performance indices

Similar to the majority of existing literature, the evaluation of the disruption predictor's performance is typically assessed in terms of:

- ✓ Successful predictions (SP): Alarms triggered by the predictor in response to disruptive pulses.
- ✓ Missed alarms (MAs): disruptions for which the predictive system does not trigger any alarm.
- ✓ False alarms (FAs): regularly terminated discharges for which the predictive system triggers an alarm.

In the case of disruptive pulses, the effectiveness of the predictor in enabling mitigation and avoidance actions can be assessed by analyzing the cumulative fraction of detected disruptions over time to disruption, i.e., over warning time.

The cumulative fraction curve in Fig. 5 provides the value, per unit, of alarms correctly activated with an anticipation time of at least the corresponding warning time $(\Delta t_{warning})$ defined as the difference between the disruption time (t_D) and the triggering of the predictor alarm (t_{alarm}). The graph also provides an overview of early alarms and their corresponding anticipation times. Additionally, it allows the reading of the fraction of correct predictions (SP), which corresponds to the intersection between the cumulative curve and the minimum time of anticipation required to activate the mitigation system ($t_D - t_{valve} = 10$ *ms* at JET), and the number of tardy alarms and missed alarms (TD+MA) as (1-SP). This graph also serves as a powerful means of comparing different models. Fig. 5 presents the cumulative curves proposed in [43] for the comparison of three different predictive models based on three different Machine Learning methods: Multilayer Perceptron (MLP), Generative Topographic Map (GTM), and Convolutional

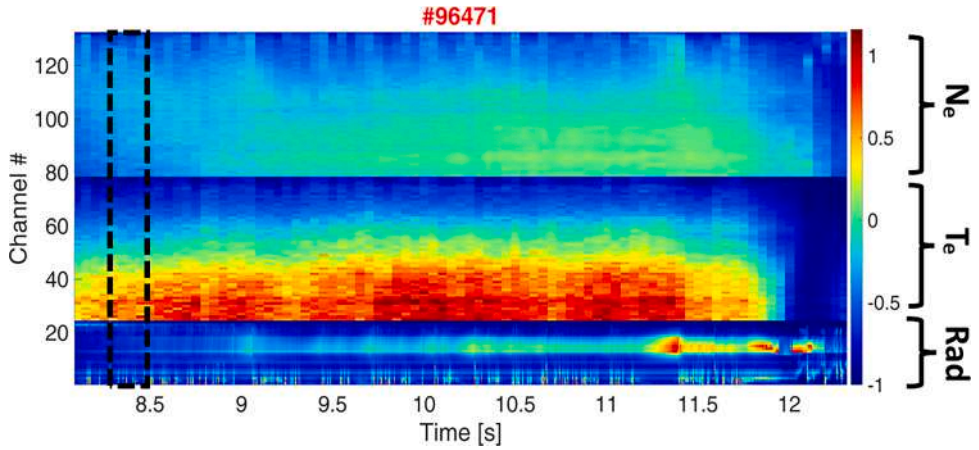

Fig. 4. JET pulse #96741: Image resulting from the processing of the plasma profiles. An overlapping window of 200 ms produces the segmented images fed to the prediction model.

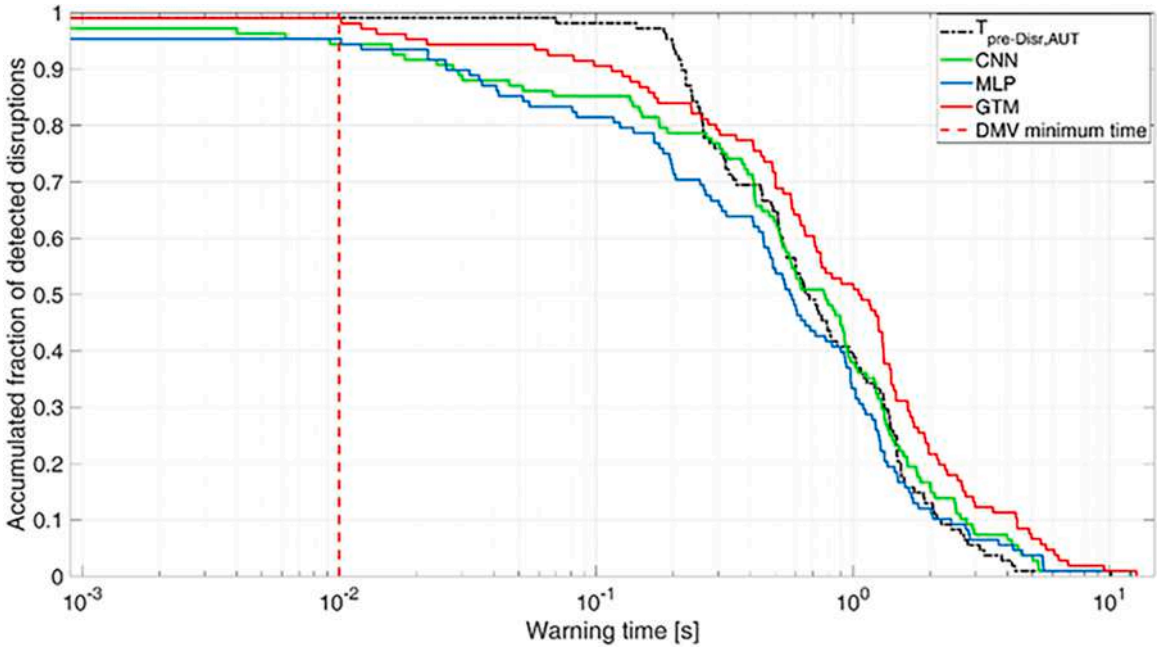

Fig. 5. Cumulative fraction of disruptions detected by the MLP-NN (blue curve), GTM (red curve), and CNN (green curve) as a function of the warning time in the test set in [43]. The dashed vertical red line, corresponding to the minimum time of anticipation required to activate the mitigation system at JET, highlights late alarm (TD) fraction. The dashed black curve refers to the reference time $t_{pre-disr}$ (graph from [43]).

Neural Network (CNN) on the same JET campaigns used in the present proposal (without taking into account the signals from the Mirnov coils).

6. Deep-CNN predictor architecture

A deep-CNN architecture typically consists of a series of filtering blocks (FBs) with multiple layers, which perform significant feature extraction from input images [44,45], see Fig. 6. A fully connected layer (FC) combines the extracted features, and in case of a classification task, it subsequently feeds them into the SoftMax (S) layer. The SoftMax layer's role is to produce two likelihoods, which sum up to one, for the sample to be disrupted or non-disrupted. To mitigate overfitting on the training set and enhance generalization, a dropout layer (D) is commonly inserted before the FC. When several images are used as input, separate branches, formed by cascades of FBs, can extract features independently from each image. These features can then be combined

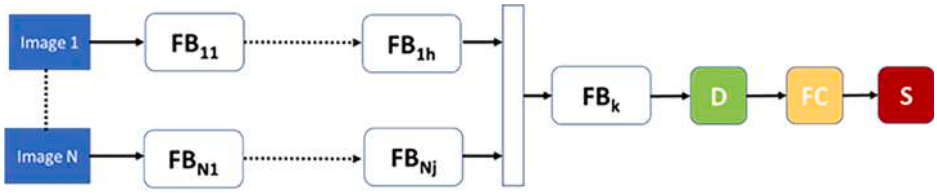

Fig. 6. General architecture of a deep-CNN.

and optionally passed to a new cascade of FBs, as depicted in Fig. 6, or directly forwarded to the fully connected layer.

Each FB generates features from images using a cascade of a Convolutional Unit (CU) followed by a pooling layer (P), as depicted in Fig. 7a. Each CU consists of a series of layers, including a convolutional layer (C), a batch-normalization layer (N), and a nonlinear activation layer with nonlinear activation functions (A), typically using Rectified Linear Units (ReLU), as shown in Fig. 7b. Specifically, the convolutional layer (C) can consist of multiple filtering layers (channel kernels) of the same size, with their weights optimized during the training phase.

Since the aim of this paper is to investigate the potential of the Mirnov coils signals as disruption predictors, the results presented in the following section will be compared to those achieved by a CNN predictor proposed by the same authors in [20], which has the architecture depicted in Fig. 8. The architecture of the reference disruption predictor consists of a cascade of two FBs. An input image, corresponding to a temporal slice extracted from the images obtained by processing the plasma profiles (refer to Fig. 4), is fed to the first filter bank.

Two 0*D* signals, namely *li* and LM_{norm}, were introduced as inputs to the second convolutional unit. These signals were combined with the output image generated by the max pooling layer (refer to Fig. 8). In the first convolutional unit (CU_1), followed by a max pooling layer (P_{1max}) with a pool size and stride of 8×1, the input image is filtered vertically (along the "spatial" dimension), resulting in a reduction in size from 132 $\times$101 to 16 $\times$ 101. Two segments of size 1×101 from the 0*D* signals are then incorporated. Subsequently, the second convolutional unit (CU_2), followed by an average pooling layer (P_{1avg}) with a pool size of 1 $\times$12 and stride of 1×4, filters the resulting image horizontally (along the "time" dimension), leading to a reduction in image size to 18 $\times$ 20, producing 360 features. In the final stage, a Multi-Layer Perceptron with a Soft Max (S) output layer determines the likelihood of the input segment belonging to specific classes. In this case, the classes are regularly terminated or disrupted discharges. For each input time slice, the S layer provides the disruptive likelihood of the last sample in the considered time slice.

Fig. 9 shows the CNN output for JET disrupted pulse #96471, which represents the disruption likelihood of each sample. It is possible to see how the output rises accordingly to the variation of the temperature profile shown in Fig. 4, and accordingly to the slowing down of the mode seen in Fig. 2. The predictor functions as a classifier, assigning each test sample to a specific class based on its disruptive likelihood value. If the likelihood value overcomes an optimized threshold alarm value, the sample is classified as belonging to the disruptive class; otherwise, it is classified as belonging to the regular (or safe) class. The alarm threshold is optimized using a heuristic procedure that minimizes the number of prediction errors, specifically the combined sum of false alarms and missed alarms, on the training and validation discharges. During the optimization, the threshold values range from a minimum of 0.5 to a maximum of 1.0.

7. Proposed disruption predictor

To evaluate the significance of Mirnov signals in predicting disruptions, an initial study focused on selecting a CNN architecture that utilized only the 2*D* spectrogram images as input. The selected highest performance architecture involves a single-branch CNN (as shown in Fig. 10). Before the fully connected layer, a series of four feature blocks (FBs) are interconnected in a cascade. The initial three feature blocks are comprised of a convolutional unit with a filter size of 3×3 and stride of 1×1, followed by a max pooling layer (P_{max}) with a pool size and stride of $3 \times$ 3. This combination filters the image both vertically and horizontally (across the spatial and temporal dimensions, respectively), resulting in a reduced image size of $3 \times 3 \times 16$ from the original $81 \times 101 \times 1$. The fourth feature block consists of a convolutional unit with a pool size and stride of 3×3 followed by an average pooling layer (P_{avg}) with a pool size and stride of 3×3, further reducing the image size to $1 \times 1 \times 32$. The number of filters of the four convolutional layers starts with 4 of the CU_1 and is doubled at each layer reaching 32 filters for CU_4.

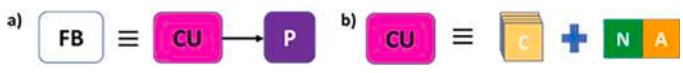

Fig. 7. a) Detail of a filtering block (FB); b) detail of a convolutional unit (CU).

This first investigation confirmed the necessity of integrating spectrograms with data from other diagnostic signals. Indeed, the CNN, exclusively trained on spectrograms, demonstrated a rate of approximately 20 % of missed alarms and a false alarm rate exceeding 50 % on the test set in Table 1. As mentioned earlier, it is important to note that not all disruptions are caused by an MHD activity, and an MHD activity itself could trigger an alarm even in the case of a regularly terminated discharge. Then, the modularity of a CNN has been further examined and a two-branch CNN is employed, as illustrated in Fig. 11. The upper branch has the same architecture of the predictor presented in Fig. 8, but it was trained using training and validation set of this work. On the other hand, the lower branch originates from the last described predictor in Fig. 10. The 352 features resulting from the two branches directly feed the fully connected layer. Note that the convolutional units and the pooling blocks in the two branches are frozen from previous training procedures, only the fully connected layer being retrained. This solution helps in reducing the complexity of the training process, leading to a significant reduction of the computation time despite the larger size of the network. This approach does not heavily impact the accuracy of the network since it reduces the number of parameter updates being performed [20]. The final classification is determined by comparing the alarm, obtained by thresholding the disruption likelihood from the SoftMax block, with the alarm obtained by thresholding LM_{norm}. Both thresholds have been fine-tuned, leading to a threshold value of 0.82 for the *S* block and a threshold value of 0.2 mT/MA for LM_{norm}. A sample is classified as disruptive if either the CNN's disruption likelihood or the LM_{norm} value exceeds its optimized threshold. Otherwise, it is classified as regular, as shown by the alarm scheme in Fig. 12.

8. Results and comparison

Table 2 presents the performance of the reference predictor (illustrated in Fig. 8) along with the proposed predictive model, which integrates the information from temperature, density and radiation plasma profiles with that from the Mirnov coil measurements.

Unfortunately, by the end of 2016, several pick-up coils, essential for both equilibrium reconstruction (slow coils) and MHD analysis (fast Mirnov coils), were lost in JET experiments [34]. Hence, due to the unavailability of the H302 or H305 signals for some shots, the composition of the database has been adjusted for both the training and test sets.

Consequently, for the comparison, the architecture proposed in [20], and shown in Fig. 8, has been tested using the test set reported in Table 1. The newly proposed predictor exhibits enhanced performance in both the training and test sets. Particularly noteworthy is the predictor's remarkable reduction in the number of false alarms triggered in response to regularly terminated pulses in the test set, decreasing from 11 to only 3. Moreover, the performance in terms of disruption prediction results in 1 tardy alarm and 1 missed alarm.

In Fig. 13, the cumulative fraction of disruptions detected by the proposed disruption predictor is presented as a function of the warning time. The red curve corresponds to the training set, the black curve corresponds to the test set, and the green curve represents the disruptions detected solely by thresholding the locked mode signal. In cases where both branches in the alarm scheme in Fig. 12 are triggered in the same discharge, only the first alarm is depicted. It is important to note that the top branch CNN, responsible for processing the plasma profiles and Mirnov coils data, can yield longer warning times compared to the bottom branch, which focuses on detecting the disruption due to mode

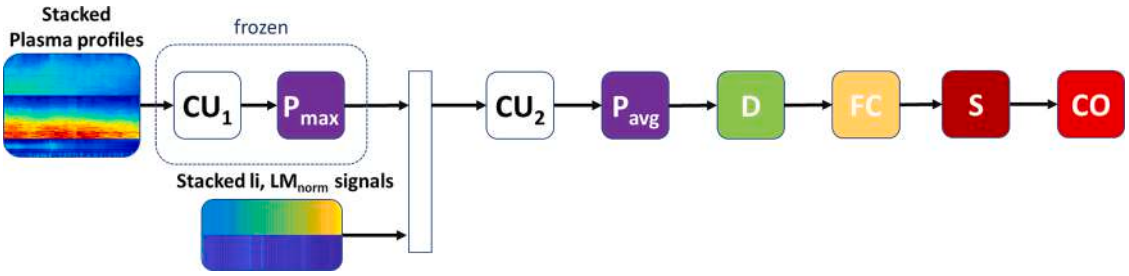

Fig. 8. Architecture of the CNN prediction model used as reference [20], where the input image represents a time slice from the images obtained processing the plasma profiles as described in Section 4 (see Fig. 4), and the lower image is obtained by stacking the 0-D signals of li and LM_{norm}.

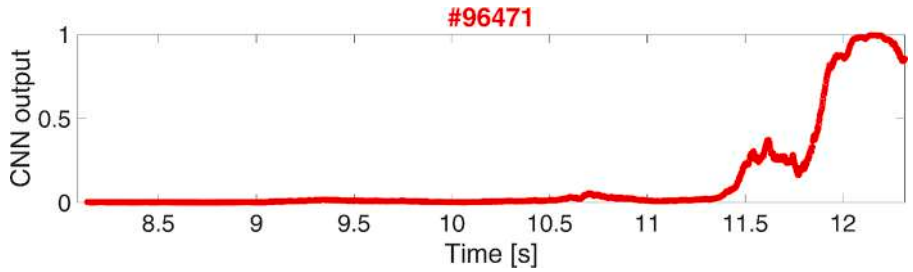

Fig. 9. Disruption likelihood of the JET disrupted pulse #96471.

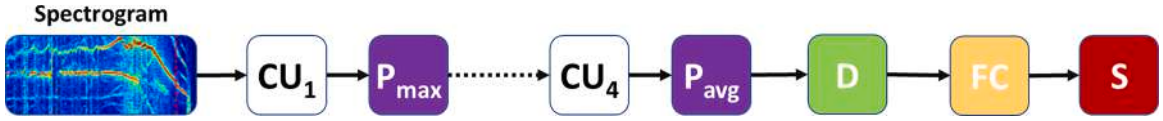

Fig. 10. Architecture of the deep-CNN prediction model that uses only information from Mirnov spectrograms: the input image represents a time slice from the images obtained processing the Mirnov coil signal.

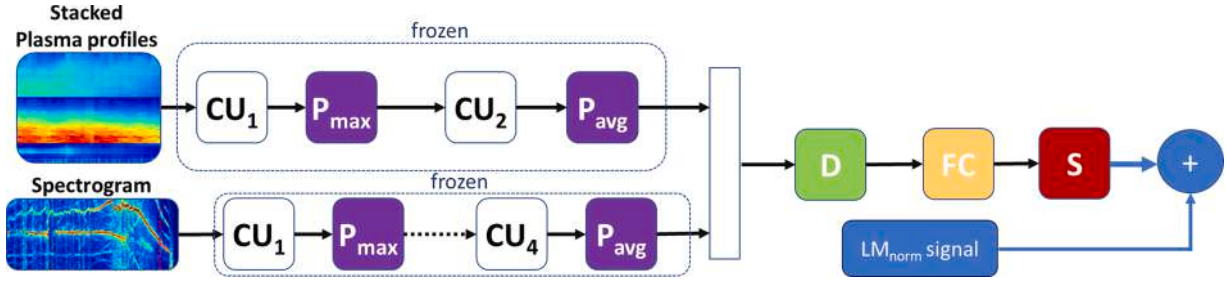

Fig. 11. Architecture of the proposed CNN prediction model, where the input images I and Spectrogram represent a time slice from the images obtained processing the plasma profiles, as described in Section 4 (see Fig. 5), and from the Spectrogram obtained processing the Mirnov coil signal.

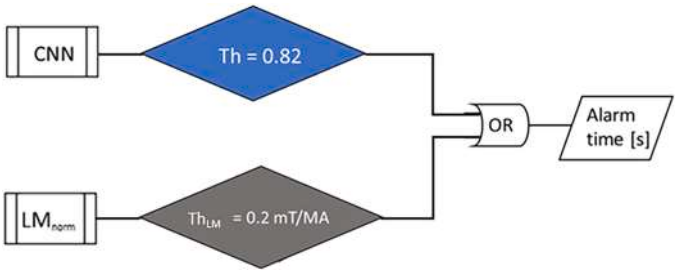

Fig. 12. Alarm scheme characterized by two parallel branches, one fed with the images processed by the CNN and the other fed with the locked mode signal, each with its optimized threshold.

Table 2
Comparison between performance of the reference and the proposed CNN disruption prediction models.

Data set	Reference Predictor			Proposed Predictor		
	TD%	MA%	FA%	TD%	MA%	FA%
Training	2.00	0	4.28	0	0	1.54
Test	0	1.09	8.40	1.09	1.09	2.29

locking.

This separation of the two distinct mechanisms enhances the interpretability of the predictor's alarms, in view of the development of avoidance schemes. Based on the graph in Fig. 13, it can be observed that around 70 % of the disruptions cause the CNN predictor to activate the alarm at least 100 ms prior to the disruption. However, when exclusively relying on the locked mode signal, this percentage drops to

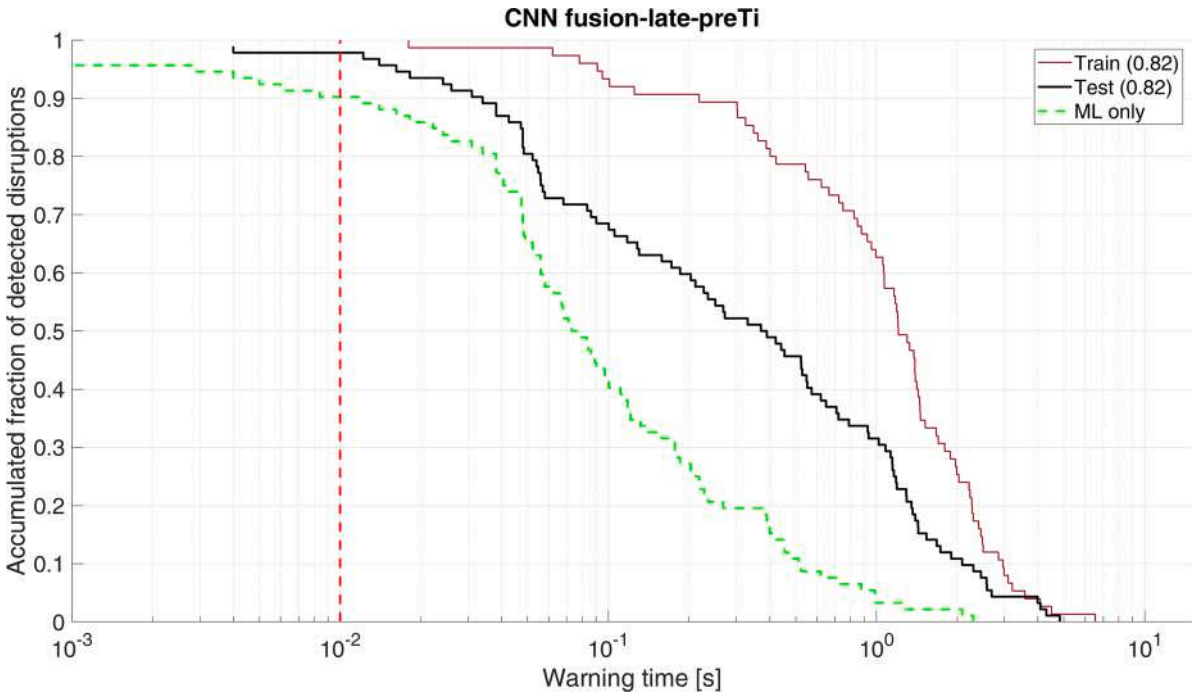

Fig. 13. Cumulative fraction of disruptions detected by the proposed disruption predictor on the training set (red curve), on the test set (black curve), and only by thresholding the locked mode signal in test set (green curve) as a function of the warning time. The dashed vertical red line (corresponding to the minimum time of anticipation required to activate the mitigation system at JET) highlights late alarms.

less than 40 %. In over 30 % of the disruptions, the warning time is a minimum of 1 second, whereas there are only a few alarms activated by the locked mode signal. Moreover, the vertical red dashed line in Fig. 13 emphasizes the disruptions identified at least 10 ms in advance, to allow for timely implementation of mitigation actions at JET. Detections with a warning time shorter than 10 ms are categorized as late or tardy alarms. An example of the working principle of the predictor can be seen in Fig. 14, which refers to the test pulse #95295 (outside the training range). The CNN output in Fig. 14a identifies a clear rise of the disruptive likelihood at about 9.4 s (dashed vertical red line), accordingly with the visible change in the plasma behavior across the input profiles shown (Figs. 14e). Later, when the 200 ms input sliding window

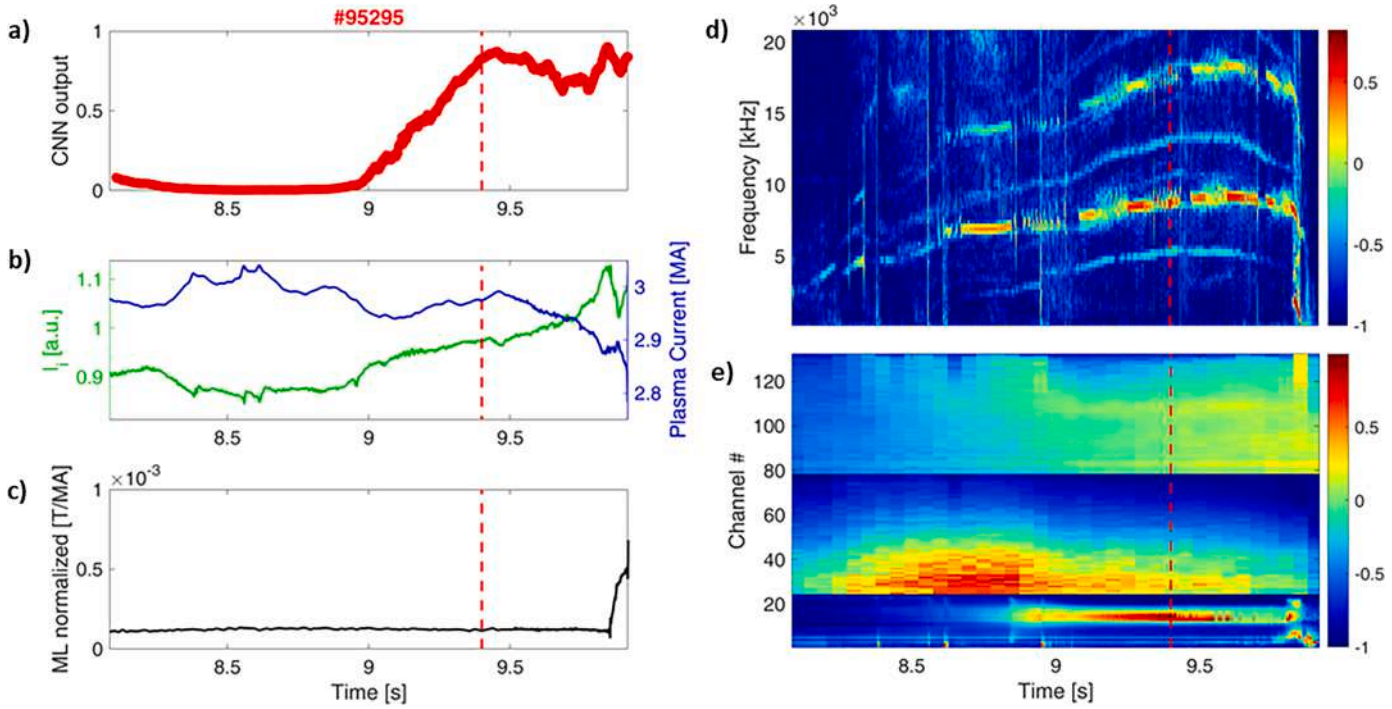

Fig. 14. JET disrupted discharge #95295. a) CNN disruptive likelihood. The dashed vertical red line indicates the CNN alarm time (the first time at which the likelihood overcomes the threshold); b) Internal inductance, in green, and plasma current in blue; c) locked mode signal normalized by the plasma current; d) CNN input spectrogram image; e) CNN input profile image.

is not focused anymore on the profile perturbation, the disruptive likelihood slightly decreases. The likelihood rises finally again in correspondence of the slowing down of the MHD activity at around 9.8 s.

Fig. 15 instead reports the CNN inputs and output for the regularly terminated test pulse #95293. As it can be seen, the CNN triggers a FA at 12.2 s nearby a high radiation from the central lines of sight of the bolometer horizontal camera, together with a decrease of both the electron temperature and the peaking of the electron density at the core (see Fig. 15e). The same false alarm was already discussed in [20] for the reference predictor. The other two FAs triggered by the proposed predictor are due to the rise of the normalized locked mode signal over the optimized threshold. Thus, no FAs seem to be triggered solely because of the Mirnov coil measurements.

9. Conclusions

This paper investigates the impact that additional diagnostic information may have in a deep-CNN model for disruption prediction. In particular, the information from the fast pick-up coils has been processed using the STFT transform [37] and has been converted into an image which can be processed as input by a CNN. This method extends the capabilities of the disruption prediction model based on a deep-CNN. First, a CNN model processes the spectrogram information, but this information alone is not sufficient to obtain a reliable disruption predictor. For this reason, the features of this CNN are stacked together with the ones obtained from the profile data of the HRTS and the Bolometer, and the output is processed by a fully connected neural network. This final set of features allowed the predictive model to detect both localized destabilizations of the temperature, density, and radiation and slowing down of the MHD modes, and it also leads to a significantly lower number of false alarms with respect to the previous predictor, for a test set consisting in 92 disruptive and 131 regularly terminated discharges. In particular, a comparison of the current model performance with the predictor proposed in [20] highlighted a drop in the FA rate from 8.4 % to 2.3 %, with a slight increase of TDs from 0 % to 1.1 % (1 case). The proposed predictor is compatible with the real-time use since the CNN takes less than 2 ms to compute a prediction and the preprocessing of the input signals is real-time compliant. Among the preprocessing procedures the bottleneck of the computation time is the FFT, which takes 5 ms when computed by a Matlab code with 32 GB RAM. However, this time could be significantly reduced by using dedicated hardware and by optimizing the algorithm on a compiled programming language. The use of the 1*D* plasma profiles as model input eases the a-posteriori interpretation between the disruption chain of events and the predictors decisions.

Therefore, the work provides a way for the integration of additional features which can detect the general physical mechanisms causing the disruptions and the development of more complex models, demonstrating that a large set of relevant diagnostics can be used to identify the different physical mechanisms which can trigger disruptions. Note that the interpretability of CNN disruption predictors could be increased by applying Explainable AI methods such as the Grad-CAM [29,46], which highlight the regions of the image that explain the model's decision. In this way, it is possible to correlate the model response to the detected phenomenon. Explainable AI approaches will be investigated in future works to understand better the impact of each diagnostic in the prediction model.

One of the challenges of data-driven predictors for disruptions is how to apply them to ITER or future devices, which is still an unresolved problem. Previous studies [17,18,47] explored the cross-tokamak method for disruption prediction using neural networks. The results showed that the neural network approach could be used for different tokamaks, but it required some adjustments of the output threshold and the operating parameter ranges based on the test data. It also suggested that using features that are consistent across different devices would be beneficial for cross-machine disruption prediction. The current predictor in this work is not directly applicable to ITER plasmas, because it uses input parameters that are not dimensionless. However, the authors are already working towards the standardization of the profile data in view of cross-machine application of CNN disruption algorithms.

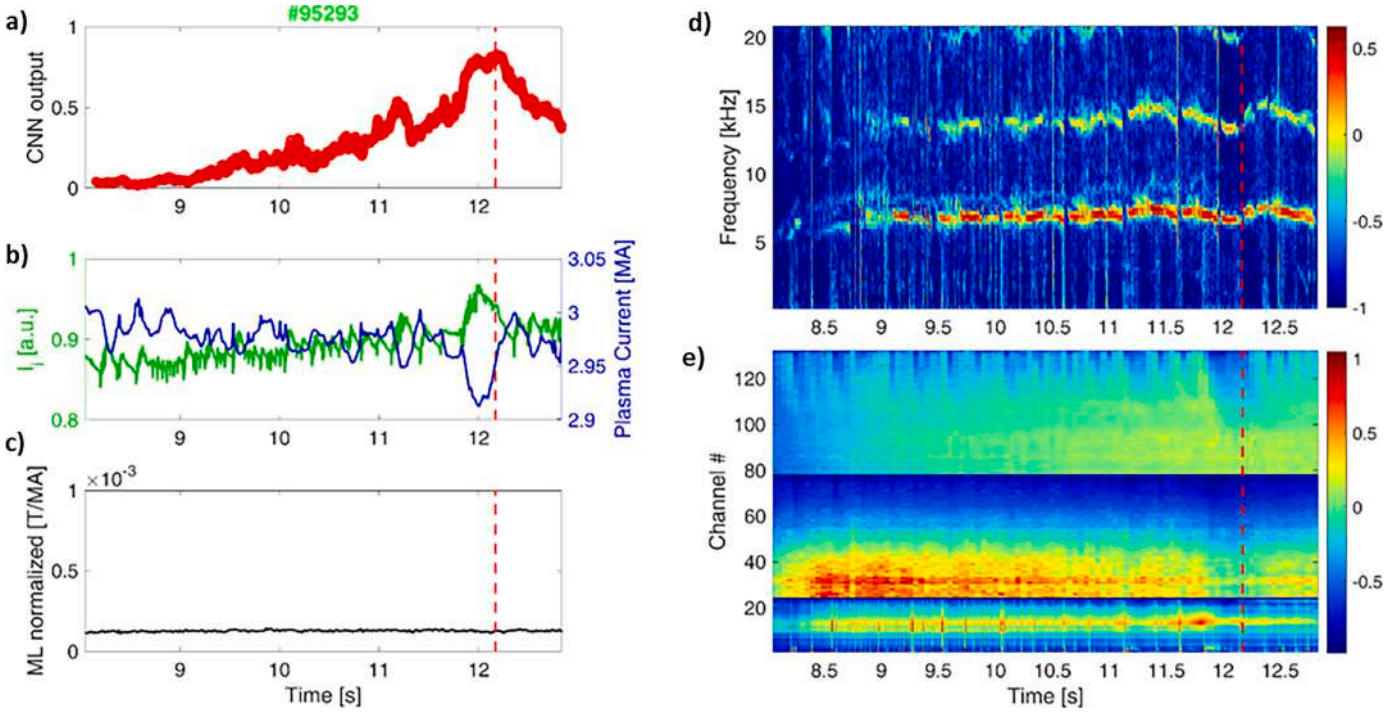

Fig. 15. CNN output on the regularly terminated discharge #95293. a) CNN disruptive likelihood, the dashed vertical red line indicates the CNN alarm time (the first time at which the likelihood overcomes the threshold); b) Internal inductance, in green, and plasma current in blue; c) locked mode signal normalized by the plasma current; d) CNN input spectrogram image; e) CNN input profile image.

Finally, the Mirnov coils processing adopted in this work is quite standard and could be adopted to other coils with the same sampling frequency. The reliability of this approach will be investigated in future works considering JET and ASDEX Upgrade data.

CRediT authorship contribution statement

E. Aymerich: Writing – original draft, Software, Formal analysis, Data curation. **G. Sias:** Writing – original draft, Methodology, Formal analysis, Data curation, Conceptualization. **S. Atzeni:** Software. **F. Pisano:** Software, Methodology. **B. Cannas:** Methodology, Conceptualization. **A. Fanni:** Writing – original draft, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data.

Acknowledgments

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (grant agreement no 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

References

[1] ITER Physics Expert Group on Disruptions Plasma Control, MHD, and I. P. B., Chapter 3: MHD stability, operational limits and disruptions, in: Nucl. Fusion, 39, 1999, p. 2251, https://doi.org/10.1088/0029-5515/39/12/303.
[2] N.W. Eidietis, et al., The ITPA disruption database, Nucl. Fusion 55 (6) (2015) 063030, https://doi.org/10.1088/0029-5515/55/6/063030.
[3] P.C. de Vries, et al., Survey of disruption causes at JET, Nucl. Fusion 51 (5) (2011) 053018, https://doi.org/10.1088/0029-5515/51/5/053018.
[4] S.A. Sabbagh, et al., Disruption event characterization and forecasting in tokamaks, Phys. Plasmas 30 (3) (2023) 032506, https://doi.org/10.1063/5.0133825.
[5] B. Cannas, A. Fanni, P. Sonato, M.K. Z, A prediction tool for real-time application in the disruption protection system at JET, Nucl. Fusion 47 (11) (2007) 1559–1569, https://doi.org/10.1088/0029-5515/47/11/018.
[6] J. Vega, et al., Results of the JET real-time disruption predictor in the ITER-like wall campaigns, Fusion Eng. Des. 88 (6) (2013) 1228–1231, https://doi.org/10.1016/j.fusengdes.2013.03.003.
[7] G.A. Rattá, J. Vega, A. Murari, D. Gadariya, J. E. T. Contributors, PHAD: a phase-oriented disruption prediction strategy for avoidance, prevention, and mitigation in JET, Nucl. Fusion 61 (11) (2021) 116055, https://doi.org/10.1088/1741-4326/ac2637.
[8] W. Zheng, et al., Hybrid neural network for density limit disruption prediction and avoidance on J-TEXT tokamak, Nucl. Fusion 58 (5) (2018) 056016, https://doi.org/10.1088/1741-4326/aaad17.
[9] R.M. Churchill, B. Tobias, Y. Zhu, Deep convolutional neural networks for multi-scale time-series classification and application to tokamak disruption prediction using raw, high temporal resolution diagnostic data, Phys. Plasma. 27 (6) (2020) 062510, https://doi.org/10.1063/1.5144458.
[10] G. Sias, et al., Disruption Prediction Approaches Using Machine Learning Tools in Tokamaks, in: 2019 PhotonIcs & Electromagnetics Research Symposium - Spring (PIERS-Spring), June, 2019, pp. 2880–2890, https://doi.org/10.1109/PIERS-Spring46901.2019.9017280.
[11] D.R. Ferreira, P.J. Carvalho, H. Fernandes, Deep learning for plasma tomography and disruption prediction from bolometer data, IEEE Trans. Plasma Sci. 48 (1) (2020) 36–45, https://doi.org/10.1109/TPS.2019.2947304.
[12] Z. Yang, F. Xia, X. Song, Z. Gao, S. Wang, Y. Dong, In-depth research on the interpretable disruption predictor in HL-2A, Nucl. Fusion 61 (12) (2021) 126042, https://doi.org/10.1088/1741-4326/ac31d8.
[13] B.H. Guo, et al., Disruption prediction using a full convolutional neural network on EAST, Plasma Phys. Control. Fusion 63 (2) (2020) 025008, https://doi.org/10.1088/1361-6587/abcbab.
[14] A. Agarwal, et al., Deep sequence to sequence learning-based prediction of major disruptions in ADITYA tokamak, Plasma Phys. Control. Fusion 63 (11) (2021) 115004, https://doi.org/10.1088/1361-6587/ac234c.
[15] B.H. Guo, et al., Disruption prediction on EAST with different wall conditions based on a multi-scale deep hybrid neural network, Nucl. Fusion 63 (9) (2023) 094001, https://doi.org/10.1088/1741-4326/ace2d4.
[16] C. Rea, K.J. Montes, A. Pau, R.S. Granetz, O. Sauter, Progress toward interpretable machine learning–based disruption predictors across tokamaks, Fusion Sci. Technol. 76 (8) (2020) 912–924, https://doi.org/10.1080/15361055.2020.1798589.
[17] J.X. Zhu, C. Rea, K. Montes, R.S. Granetz, R. Sweeney, R.A. Tinguely, Hybrid deep-learning architecture for general disruption prediction across multiple tokamaks, Nucl. Fusion 61 (2) (2020) 026007, https://doi.org/10.1088/1741-4326/abc664.
[18] J. Kates-Harbeck, A. Svyatkovskiy, W. Tang, Predicting disruptive instabilities in controlled fusion plasmas through deep learning, Nature 568 (7753) (2019) 7753, https://doi.org/10.1038/s41586-019-1116-4. Art. no.
[19] A. Pau, et al., A machine learning approach based on generative topographic mapping for disruption prevention and avoidance at JET, Nucl. Fusion 59 (10) (2019) 106017, https://doi.org/10.1088/1741-4326/ab2ea9.
[20] E. Aymerich, et al., Disruption prediction at JET through Deep Convolutional Neural Networks using spatiotemporal information from plasma profiles, Nucl. Fusion 62 (2022) 066005, https://doi.org/10.1088/1741-4326/ac525e.
[21] E. Aymerich, G. Sias, F. Pisano, B. Cannas, A. Fanni, the-JET-Contributors, CNN disruption predictor at JET: early versus late data fusion approach, Fusion Eng. Des. 193 (2023) 113668, https://doi.org/10.1016/j.fusengdes.2023.113668.
[22] Active control of magneto-hydrodynamic instabilities in hot plasmas, vol. 83, in: V. Igochine (Ed.), Springer Series on Atomic, Optical, and Plasma Physics, vol. 83, Springer, Berlin, Heidelberg, 2015, https://doi.org/10.1007/978-3-662-44222-7.
[23] J.S. Kim, D.H. Edgell, J.M. Greene, E.J. Strait, M.S. Chance, MHD mode identification of tokamak plasmas from Mirnov signals, Plasma Phys. Control. Fusion 41 (11) (1999) 1399, https://doi.org/10.1088/0741-3335/41/11/307.
[24] M. Baruzzo, et al., Neoclassical tearing mode (NTM) magnetic spectrum and magnetic coupling in JET tokamak, Plasma Phys. Control. Fusion 52 (7) (2010) 075001, https://doi.org/10.1088/0741-3335/52/7/075001.
[25] G. Pucella, et al., Onset of tearing modes in plasma termination on JET: the role of temperature hollowing and edge cooling, Nucl. Fusion 61 (4) (2021) 046020, https://doi.org/10.1088/1741-4326/abe3c7.
[26] A. Bondeson, R.D. Parker, M. Hugon, P. Smeulders, MHD modelling of density limit disruptions in tokamaks, Nucl. Fusion 31 (9) (1991) 1695, https://doi.org/10.1088/0029-5515/31/9/008.
[27] P.C. de Vries, et al., Scaling of the MHD perturbation amplitude required to trigger a disruption and predictions for ITER, Nucl. Fusion 56 (2) (2015) 026007, https://doi.org/10.1088/0029-5515/56/2/026007.
[28] B. Cannas, et al., Wavelet analysis of Mirnov coils signals for disruption prediction at JET, in: Mánes: 45th EPS Conference on Plasma Physics, 2018 [Online]. Available: http://ocs.ciemat.es/EPS2018ABS/pdf/P4.1017.pdf.
[29] D.R. Ferreira, T.A. Martins, P. Rodrigues, J. E. T. Contributors, Explainable deep learning for the analysis of MHD spectrograms in nuclear fusion, Mach. Learn. Sci. Technol. 3 (1) (2021) 015015, https://doi.org/10.1088/2632-2153/ac44aa.
[30] C. Reux, et al., Use of the disruption mitigation valve in closed loop for routine protection at JET, Fusion Eng. Des. 88 (6) (2013) 1101–1104, https://doi.org/10.1016/j.fusengdes.2012.12.026.
[31] R.H. Tong, et al., The impact of an m/n = 2/1 locked mode on the disruption process during a massive gas injection shutdown on J-TEXT, Nucl. Fusion 59 (10) (2019) 106027, https://doi.org/10.1088/1741-4326/ab32a3.
[32] S.N. Gerasimov, et al., Locked mode and disruption in JET-ILW, in: 46th European Physical Society Conference on Plasma Physics (EPS), Milan, 8-12 July 2019, 2023. Accessed: Jan. 18[Online]. Available: https://scientific-publications.ukaea.uk/papers/locked-mode-and-disruption-in-jet-ilw/.
[33] G. Sias, B. Cannas, A. Fanni, A. Murari, A. Pau, A locked mode indicator for disruption prediction on JET and ASDEX upgrade, Fusion Eng. Des. 138 (2019) 254–266, https://doi.org/10.1016/j.fusengdes.2018.11.021.
[34] G. Artaserse, M. Baruzzo, R.B. Henriques, S. Gerasimov, N. Lam, M. Tsalas, Refurbishment of JET magnetic diagnostics, Fusion Eng. Des. 146 (2019) 2781–2785, https://doi.org/10.1016/j.fusengdes.2019.05.032.
[35] Gerasimov, S., 'JET magnetic diagnostic reference page https://users.euro-fusion.org/pages/mags/mhd/figures-mhd/hr-coil-array.pdf'.
[36] M. Baruzzo, et al., Fault analysis and improved design of JET in-vessel Mirnov coils, Fusion Eng. Des. 150 (2020) 110863, https://doi.org/10.1016/j.fusengdes.2019.02.123.
[37] J.P.S. Bizarro, A.C.A. Figueiredo, Time–frequency analysis of fusion plasma signals beyond the short-time Fourier transform paradigm: an overview, Fusion Eng. Des. 83 (2) (2008) 350–353, https://doi.org/10.1016/j.fusengdes.2007.12.006.
[38] A. Pau, et al., A First Analysis of JET Plasma Profile-Based Indicators for Disruption Prediction and Avoidance, IEEE Trans. Plasma Sci. 46 (7) (2018) 2691–2698, https://doi.org/10.1109/TPS.2018.2841394.
[39] R. Aledda, B. Cannas, A. Fanni, G. Sias, G. Pautasso, Multivariate statistical models for disruption prediction at ASDEX Upgrade, Fusion Eng. Des. 88 (6) (2013) 1297–1301, https://doi.org/10.1016/j.fusengdes.2013.01.103.
[40] E. Aymerich, et al., A statistical approach for the automatic identification of the start of the chain of events leading to the disruptions at JET, Nucl. Fusion 61 (3) (2021) 036013, https://doi.org/10.1088/1741-4326/abcb28.

[41] A. Murari, R. Rossi, M. Lungaroni, M. Baruzzo, M. Gelfusa, Stacking of predictors for the automatic classification of disruption types to optimize the control logic, Nucl. Fusion 61 (3) (2021) 036027, https://doi.org/10.1088/1741-4326/abc9f3.
[42] R. Rossi, M. Gelfusa, J. Flanagan, A. Murari, J. E. T. Contributors, Development of robust indicators for the identification of electron temperature profile anomalies and application to JET, Plasma Phys. Control. Fusion 64 (4) (2022) 045002, https://doi.org/10.1088/1361-6587/ac4d3b.
[43] E. Aymerich, et al., Performance Comparison of Machine Learning Disruption Predictors at JET, Appl. Sci. 13 (3) (2023) 3, https://doi.org/10.3390/app13032006. Art. no.
[44] S. Pouyanfar, et al., A Survey on Deep Learning: algorithms, Techniques, and Applications, ACM Comput. Surv. 51 (5) (2018), https://doi.org/10.1145/3234150, p. 92:1-92:36.
[45] M.Z. Alom, et al., A State-of-the-Art Survey on Deep Learning Theory and Architectures, Electronics (Basel) 8 (3) (2019) 3, https://doi.org/10.3390/electronics8030292. Art. no.
[46] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-CAM: visual explanations from deep networks via gradient-based localization, in: 2017 IEEE International Conference on Computer Vision (ICCV), October, 2017, pp. 618–626, https://doi.org/10.1109/ICCV.2017.74.
[47] C.G. Windsor, et al., A cross-tokamak neural network disruption predictor for the JET and ASDEX Upgrade tokamaks, Nucl. Fusion 45 (5) (2005) 337–350, https://doi.org/10.1088/0029-5515/45/5/004.