Fusion Engineering and Design 193 (2023) 113668

Contents lists available at ScienceDirect

Fusion Engineering and Design

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

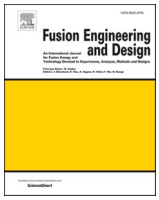

CNN disruption predictor at JET: Early versus late data fusion approach

E. Aymerich [*], G. Sias, F. Pisano, B. Cannas, A. Fanni, the-JET-Contributors [1]

Department of Electrical and Electronic Engineering, University of Cagliari, Cagliari, Italy

A B S T R A C T

This work focuses on the development of a data driven model, based on Convolutional Neural Networks (CNNs), for the real-time detection of disruptive events at JET. The predictor exploits the ability of CNNs in learning relevant spatiotemporal information straight from 1D plasma profiles, avoiding hand-engineered feature extraction procedures. In this paper, the radiation profiles from both the bolometer horizontal and vertical cameras have been considered amongst the predictor inputs, with the aim of discriminating between the core radiation due to impurity accumulations and the outboard radiation phenomena. Moreover, an innovative predictor architecture is proposed, where two separate CNNs are trained to focus on events with different timescales, that is, the destabilization of radiation, electron density and temperature profiles, and the mode-locking and current profile variations. The outputs of the two CNNs are combined with a logic OR function to provide the disruption alarm trigger. The advantages of this data fusion approach impact on the predictor performance, with a very limited number of false alarms (only 1 in the considered test set), and on the model output interpretability as the two different branches are triggered by different types of events.

1. Introduction

Plasma disruptions pose considerable risks to the operation of high-power nuclear fusion devices. In fact, very stringent requirements apply to next-generation tokamaks in terms of the allowable number of unmitigated disruptions. In addition, the scientific community is currently actively working on the task of disruption avoidance, aiming to determine the mechanism that causes sudden discharge termination. The goal is to enable a recovery action or safe termination of the experiment without Massive Gas Injection (MGI). Data-driven methods are very powerful in fault prediction, and many approaches, such as Fully Connected Neural Networks (FC), Support Vector Machines (SVM), Self-Organizing Maps (SOM) and Generative Topographic Mapping (GTM), Classification and Regression Trees (CART) and Random Forests (RF), have been employed in disruption prediction models at JET [1–9], ASDEX Upgrade [10–12], EAST [13], J-TEXT [14], DIII-D [15] and Alcator C-Mod [13]. Moreover, research aims to link the physical phenomena involved in disruptive processes to the inputs of the data-driven prediction models, especially in case of predictive models designed to enable disruption avoidance [6,16-19]. Recently, deep Convolutional Neural Networks (CNNs) have been adopted in several applications and became the state of the art for image processing and computer vision. In fact, they are able to automatically extract features from images, overcoming the need for hand-engineered feature extraction [20–22]. In [8] the authors proposed a disruption predictor based on a Convolutional Neural Network (CNN) for JET tokamak. The CNN processes an image obtained from a set of 1D profile data and 0D signals and returns two likelihoods: the disruptive one (in red in Fig. 1a) and the non-disruptive one (in green in the same Fig. 1a). The predictor was trained and tested on data spanning several years, showing overall good performances. However, a study of the performances evolution over the different campaigns revealed the predictor ageing, with an accuracy degradation, mainly in the false alarm number (Table 3 in [8]). Indeed, during the 2020 JET high power experiments, researchers observed the appearance of localized radiation in the Low Field Side (LFS) [23]. Since the previously developed predictor was only analysing the information from the bolometer horizontal camera, the CNN could not correctly locate the radiation source in these cases, hence triggering a false alarm. Fig. 1 reports a false alarm triggered by the predictor in [8] on pulse #96893. As can be noted, at around 10.5 s, despite a non-disruptive behaviour shown by the HRTS profiles and the 0D signals, high radiation seen from the central lines of sight of the bolometer horizontal camera (Fig. 1d) triggers an alarm, highlighted with a black dashed line in Figs. 1a-f. This behaviour, observed in 4 other pulses, motivated the development of a new version of the predictor, also including the information from the vertical bolometer camera amongst the set of images.

This paper is structured as follows: Section 2 reports the Database used for this study; Section 3 explains the predictor architecture; Section 4 reports the training of the predictor model; Section 5 details the results. Finally, in Section 6 the conclusions of the study are discussed.

* Corresponding author.
E-mail address: enrico.aymerich@unica.it (E. Aymerich).

[1] See the author list of Overview of JET results for optimising ITER operation by J. Mailloux et al. to be published in Nuclear Fusion Special issue: Overview and Summary Papers from the 28th Fusion Energy Conference (Nice, France, 10–15 May 2021).

https://doi.org/10.1016/j.fusengdes.2023.113668
Received 31 October 2022; Received in revised form 8 February 2023; Accepted 15 March 2023
Available online 30 March 2023

2. Database

In this work, an upgraded version of the disruption predictor based on the CNN architecture presented in [8] is proposed. For sake of comparison, the database used to train and test the predictor is the same as in [8]. Precisely, the model has been trained using the same 85 disrupted and 70 regularly terminated discharges used in [7,8], selected from the 2011–2013 JET campaigns. Then, the model was tested on 108 disruptive and 149 regularly terminated subsequent pulses from the 2011–2013 (42 disrupted /45 non-disrupted), 2016 (29/41), and 2019–2020 (37/63) JET campaigns. The pulses are analysed in the flat-top phase. For each selected discharge, the flat-top starting time is the first time instant where the plasma is in X-point configuration. The flat-top ending time t_{end}, is the last available time instant in the flat-top for the non-disrupted discharges, while for the disrupted ones, t_{end} is the minimum time between the valve activation time for mitigated discharges, and the disruption time (t_D), for the unmitigated ones. The predictor architecture has been designed to optimize the performance according to the new set of inputs. The 1D plasma profile data, from the High-resolution Thomson Scattering and the bolometer horizontal and vertical cameras were sorted to preserve the spatial ordering of the channels, resampled to 2 ms to maintain a common time basis, and pre-processed to remove outliers and unreliable channels. For the HRTS data, the pre-processing consists in the space interpolation of the channels having a large estimated error [24]. For the bolometer cameras, values higher than 1MW/m^2 were saturated, as well as negative ones which were saturated to 0 [8]. Afterwards, the input images to be fed to the CNN were extracted using a 200 ms sliding window. Then by using the maximum and the minimum values from each diagnostic in the training set, the four different images are normalized and vertically stacked, as shown in Fig. 2.

3. Predictor architecture

In general, the architecture of a CNN contains several layers of blocks, connected in cascade, which filter an input image for feature extraction purposes [20]. Each filtering layer is interconnected with the following one by a nonlinear block (usually a rectified linear unit). Finally, a multi-layer perceptron combines the extracted features to produce the output of the CNN. A dropout layer may be included to improve the generalization of the model. Fig. 3 shows the architecture of the predictor proposed in this work. It consists of two branches, each one being a separate CNN. The top branch, which processes the images of the 1D profiles, has two convolutional units (CU_1, CU_2) followed by a max pooling layer (P_{max}) and an average pooling layer (P_{avg}) respectively. The CU_1 and P_{max} blocks, filter out vertically (spatial dimension) the input image by reducing its size from 154 × 101 to 18 × 101. The other blocks (CU_2, P_{avg}) filter out horizontally (time dimension) the resulting image by reducing the image size to 18 × 20. Each convolutional unit consists of three layers: a convolutional layer (C_k), a batch normalization layer (N_k) and a rectified linear unit (ReLU) activation layer (A_k). The first convolutional layer has a single filter (1-channel kernel) of size 5 × 1, while the second one has one of size 1 × 11. The output of the 2nd convolutional unit is then a 18 × 20 image. The lower branch processes the stacked signals of the internal inductance l_i and the normalized Locked Mode ML_{norm} signals. It consists of a separate Convolutional Unit (CU_3) with 4 filters (4-channel kernel) of size 1 × 5 with dilation 1 × 5 and stride 1 × 1, which process the 0D dimensional data along the horizontal direction. The block is followed by a max pooling layer with size and stride 1 × 5, which also down samples the features along the horizontal direction. The extracted features have a size of 2 × 16 × 4. On both branches, the features are flattened and fed into a Fully Connected (FC) block, which combines them before a SoftMax layer (S). Before the two fully connected layers, a dropout layer with dropout probability of 20% reduces the overfitting on the training set and improves the model generalization.

The SoftMax layer produces the likelihood of the input segment to belong to a disrupted discharge. As an example, Fig. 4a shows the SoftMax outputs for the JET disrupted pulse #96998, where the blue line refers to the disruptive likelihood from the top branch and the magenta line that one from the bottom branch. Finally, for each branch, a final classification layer (CO) simply thresholds the disruptive

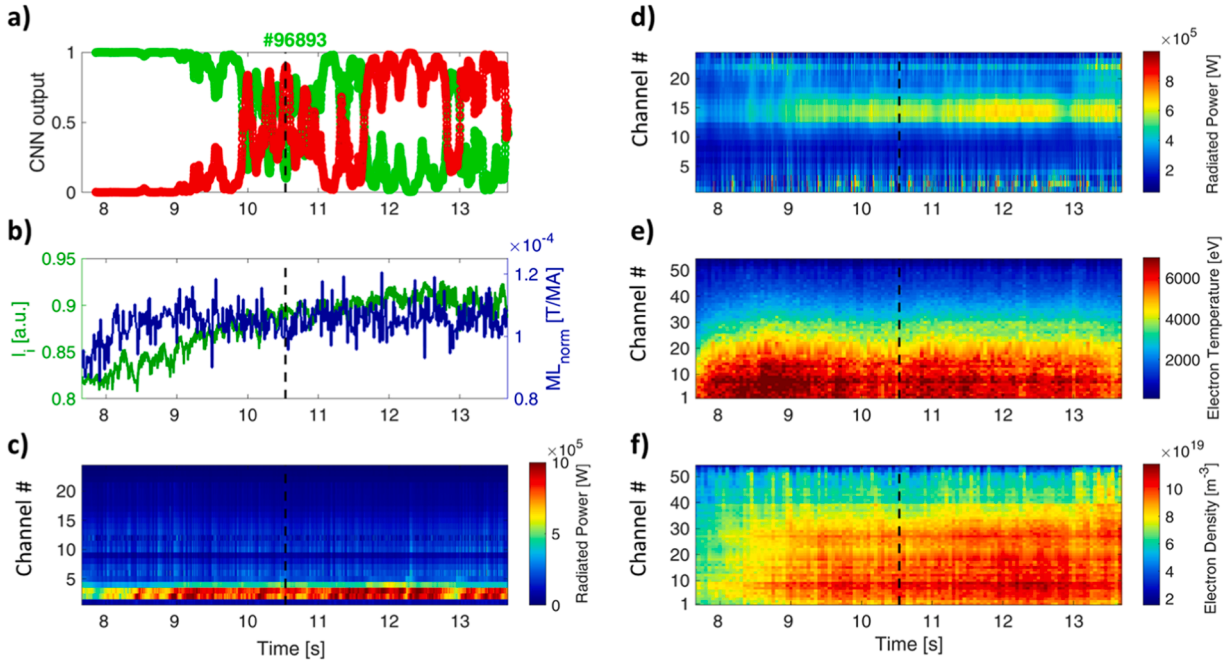

Fig. 1. False alarm triggered by the predictor in [8] on the regularly terminated discharge #96893. (a) CNN likelihood curves, where the green line is the regularly termination likelihood and the red line is the disruption likelihood; (b) internal inductance in green and mode lock normalized by the plasma current in blue; (c) radiated power from the bolometer vertical camera; (d) radiated power from the bolometer horizontal camera; (e) electron temperature from the HRTS; (f) electron density from the HRTS. The dashed black line indicates the CNN alarm time.

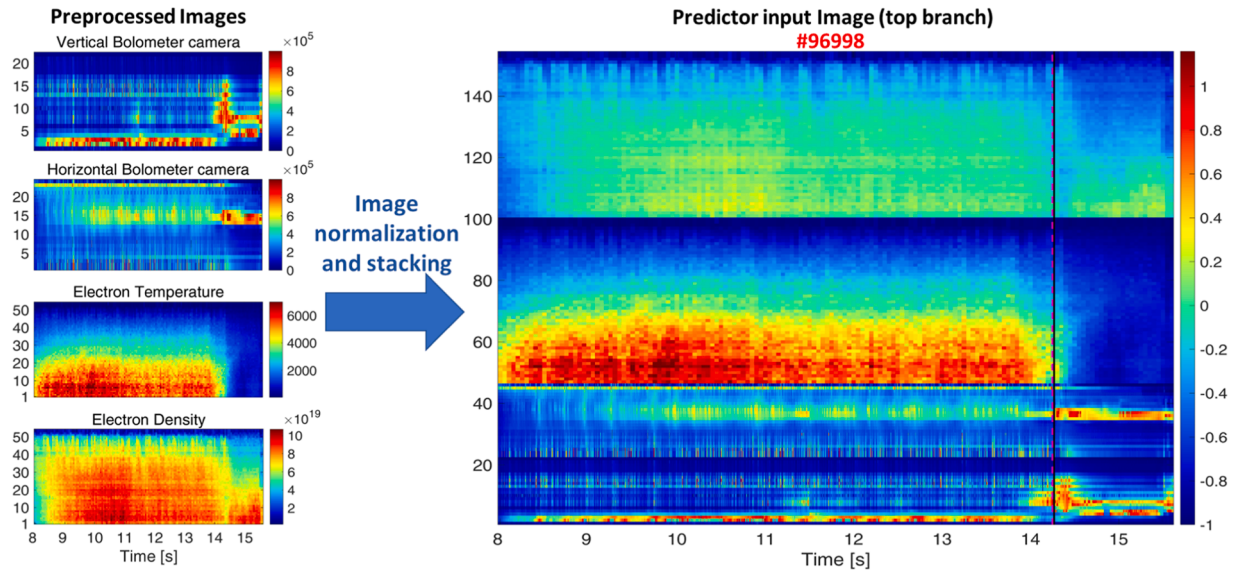

Fig. 2. Sketch of the pre-processing steps applied to pulse #96998 to generate the input images for the predictor left) Pre-processed data are converted into images; Right) Input data for the CNN model, obtained by normalizing the data with the training set ranges and by vertically stacking the different diagnostics. The black solid line indicates the disruption alarm triggered by the proposed predictor whereas the dashed magenta line indicates the $t_{pre-disr}$.

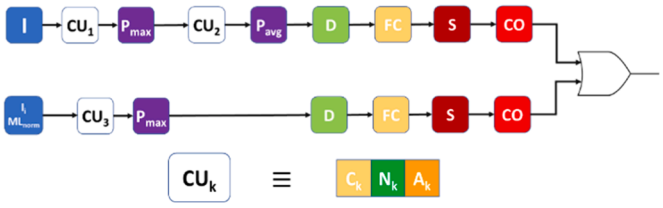

Fig. 3. CNN architecture, where: I is the image input; CU_k is the kth convolutional unit, composed by the cascade of a convolutional layer (C_k), a batch-normalization layer (N_k) and a nonlinear activation layer with ReLU functions (A_k); P_{max} and P_{avg} are the max-pooling and average-pooling layers, respectively; D is a dropout layer; FC is a fully-connected layer; S and CO are the SoftMax and classification output layers, respectively. Finally, an OR logic block activates the predictor whether one of the two branches output is 1.

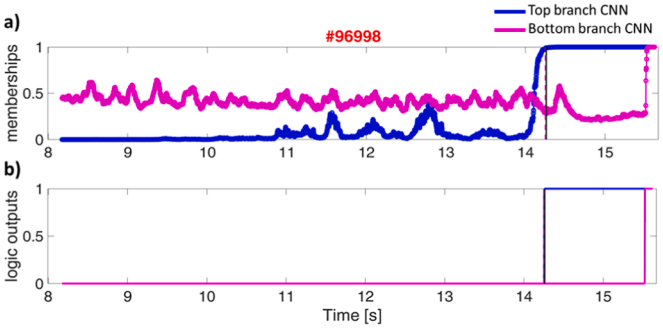

Fig. 4. Disrupted pulse #96998 a) Disruptive membership functions for each predictor branch, where the blue line is the top branch one, and the magenta line is the bottom branch membership; b) Logic output for each branch (blue for the top branch, magenta for the bottom one) for the same pulse. The dashed magenta line indicates the $t_{pre\text{-}disr}$, the solid black line indicates the alarm time.

likelihood to perform the image classification. For each branch, a threshold on the likelihood is optimized by minimising the errors of the entire predictor on the training set, as detailed in Section 4. Fig. 4b shows the branch binary outputs, which are obtained by setting to 1 the membership values greater than or equal to their own optimized threshold, and by setting to 0 the remaining ones. A disruptive behaviour is detected by a branch when its binary output equals 1 (blue curve for top branch and magenta curve for the bottom branch in Fig. 4b). The logic OR function produces the final disruption trigger.

4. Predictor training

Since the CNN is employed in a supervised learning framework, it is necessary to explicitly label the time windows (or time slices) in the training dataset. In this work, the non-disrupted time slices are selected from the non-disrupted discharges, while the disrupted ones are obtained by identifying specific time windows within the disrupted discharges. Since the two CNNs were trained independently from each other, two different criteria have been adopted for defining the disruptive phase. The reason for adopting a different definition is the training of two specialized CNN branches, where each of them is focusing on events with different timings. In particular, the destabilization of the profiles at JET is usually due to the process of impurity accumulation or to the edge cooling [25], revealable by the plasma radiated power and density profiles, and it is exhibited at longer timescales than the insurgence of the locked mode. Hence, the two branches aim to increase the performance of the entire model exploiting the different information carried out by the profiles and the 0D signals. For the 1D profile images, the onset of disrupted phase is defined by the automatically identified pre-disruptive times $t_{pre-disr}$ as in [7,8], whereas, for the 0D signal images, the onset of disrupted phase is defined by the mode locking time (t_{ML}). To this purpose, a threshold has been optimized, resulting in $2{\cdot}10^{-4}$ mT/MA, on the Locked Mode signal normalized by the plasma current. The time interval [t_{ML}, min(t_{ML}+0.3 s, t_{end})] has been labelled as disruptive phase.

Due to the unbalance between the number of non-disrupted and disrupted samples, caused by the different duration of the two pre-disrupted phases, different subsampling strategies for the 200 ms sliding window have been adopted for the training. For the CNN top branch, one image every 24 ms has been sampled from the disrupted discharges in the timespan after $t_{pre-disr}$, whereas one image every 150 ms has been extracted from the non-disrupted discharges. This choice is motivated by the low resolution of the HRTS, which has a 50 ms sampling period. Instead, for the 0D signals every segment of pre-disrupted phase (i.e., one every 2 ms) is considered for the training, whereas one

segment every 200 ms is sampled from the regularly terminated discharges. In the test instead, the sliding window has a stride of 2 ms, so that every sample of all the test discharges (regularly terminated and disrupted) has been classified. The alarm thresholds of the CO layers have been chosen by optimizing the full predictor performances on the training data. In the disruption prediction literature, the following metrics are generally considered when comparing different predictors:

- Successful predictions (SP): pulses that are correctly predicted (hence, an alarm is triggered in disruptive pulses and no alarm is raised in non-disruptive discharges).
- Missed alarms (MAs): pulses which disrupt where the predictor does not trigger an alarm.
- False alarms (FAs): non-disruptive discharges where the model triggers an alarm.

Considering these metrics, the single branch thresholds have been selected by minimizing the sum of the full predictor MAs and FAs, and then the distance between the alarm times and the $t_{pre-disr}$ on the training discharges. In fact, firstly a scan of the different thresholds identifies the combinations where the sum of the FAs and MAs is minimized. In this subset, the thresholds which minimize the mean distance between alarm times and $t_{pre-disr}$ are selected. The optimized thresholds result in 0.99 for the top branch and 0.925 for the bottom one.

5. Results and discussion

The results of the predictor are reported and compared with [8] in Table 1. The new model performs better both in the training and in the test sets. In particular, the predictor allows to greatly reduce the number of false alarms in the test set (from 14 to 1).

Another important metric for evaluating disruption predictors designed for avoidance and/or mitigation purposes is the warning time distribution, used to statistically evaluate the available time from the predictor alarm before the t_{end}. An early warning time could allow the adoption of automatic procedures to try to recover the disruptive plasma state or to safely terminate the experiment, while with a short warning time the disruption is generally mitigated. However, to allow the adoption of disruption avoidance strategies, the model should also provide information on the type of instability which is destabilizing the discharge. Fig. 5 reports the warning times of the top branch (blue line), bottom branch (green line), and full predictor (black line) in the test dataset. If both branches are triggered in the same discharge, only the first alarm is plotted. In fact, by reading in correspondence to the time to disruption equal to 10^{-3}, in Fig. 5, it is possible to notice that around 54.62% and 43.51% of the disruptions are detected by the top and bottom branches, respectively. Note that, the top branch CNN, which processes the 1D profile data, can provide larger warning times than the bottom one, which instead detects the mode-locking phase. The separation of the two different mechanisms makes the predictor alarm more interpretable, in view of the development of avoidance schemes. Finally, the vertical red dashed line highlights that disruptions should be identified at least 10 ms in advance to adopt mitigation actions at JET. Detections with a warning time shorter than 10 ms are late or tardy alarms. The predictor can detect different disruptive patterns. For instance, in pulse #96998, as visible from the Fig. 2 and Fig. 4, the top branch detects a change of the plasma profiles, and it allows the predictor to trigger an alarm at 14.25 s in correspondence of the $t_{pre-disr}$, identified by a dashed yellow line.

Table 1
Predictor performance.

Dataset	SP	MA	FA
Train	98.71%	0%	2.85%
Train [8]	98.00%	0%	4.28%
Test	98.83%	1.87%	0.67%
Test [8]	93%	3.7%	9.4%

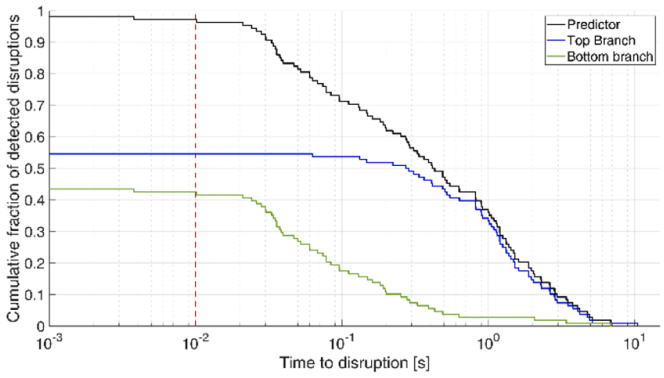

Fig. 5. CNN model warning time distributions in the test set for the top branch (blue line), the bottom one (green line) and full predictor (black line). Only the first alarm is reported. The vertical red dashed line allows to identify tardive detections.

In fact, the electron temperature (T_e) flattens and the electron density (n_e) peaks. This phenomenon is synchronous with strong radiation from the central channels of the horizontal and vertical bolometer. On the other hand, the membership function of the bottom branch in Fig. 4 (magenta line) rises later, at around 15.7 s, in correspondence to the Locked mode onset. Hence, the top-branch is trained to detect destabilizations in the 1D profiles distributions, while the bottom branch on detecting the onset of a locked-mode and a late disruption pattern. Fig. 6, instead, shows the disruption detection in pulse #92426. The membership function of bottom branch (blue line, Fig. 6a) triggers the alarm at around 13.80 s because of the li and ML_{norm} growth (Fig. 6b), while the membership function of the top branch slowly grows following the late change of n_e and T_e profiles.

Fig. 7 shows the regularly terminated pulse #96,893, which was detected as disruptive in [8]. In this case, the predictor does not trigger an alarm, because the high radiation pattern at chords #13–16 of the horizontal bolometer is not coincident with a high radiation from the central lines of sight of the vertical bolometer camera.

6. Conclusions

In this work, a disruption predictor based on CNNs has been developed using data from 2011–2013 campaigns at JET. The test of the model included more recent JET discharges and high power experiments up to the 2020 experimental campaigns. First, the vertical bolometer camera is added to the set of 1D plasma profile features considered in [8]. Then, two different CNN classifiers, whose thresholds are optimized to achieve the best full predictor performance, are trained to detect different destabilizing events. The automatic detection of the pre-disruptive phase of disruptions is used to train the top branch CNN, while an automatically identified locked mode time is employed for training the bottom branch of the model. The model can correctly identify the local perturbations of the 1D plasma profiles, leading to about, 98.87% of SPs, 0.67% of FAs and 1.87% of MA, considering a test set with 108 disruptive and 149 non-disruptive discharges. Moreover, the proposed approach allows to associate the predictor alarm with the destabilizing mechanism of the discharge. The automatic classification of the different profile instabilities, for instance distinguishing between temperature hollowing and edge cooling [25], would be another step forward towards the implementation of machine learning aided avoidance schemes.

Declaration of Competing Interest

The authors declare that they have no known competing financial

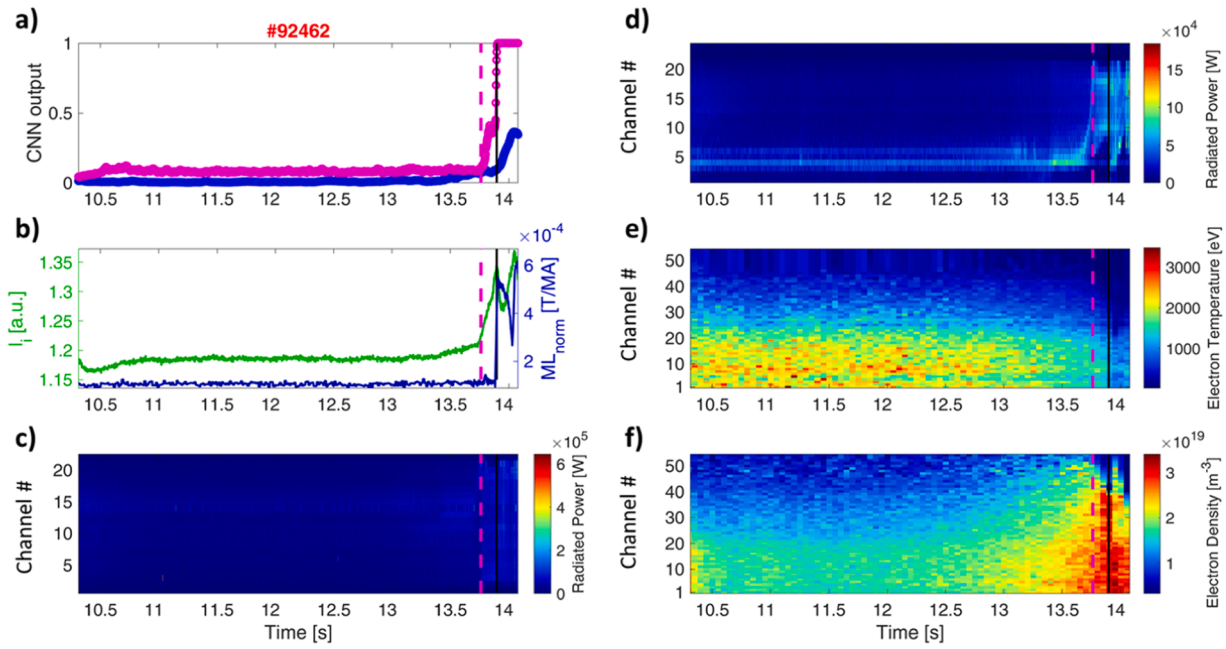

Fig. 6. JET disrupted discharge #92462. (a) CNN membership functions, where the blue line is the top branch membership function, and the magenta line is the bottom branch membership function; (b) internal inductance, in green, and mode lock normalized by the plasma current, in blue; c) radiated power from the bolometer vertical camera; (d) radiated power from the bolometer horizontal camera; (e) electron temperature from the HRTS; (f) electron density from the HRTS. The dashed magenta line indicates the $t_{pre-disr}$, and the solid black line the alarm time.

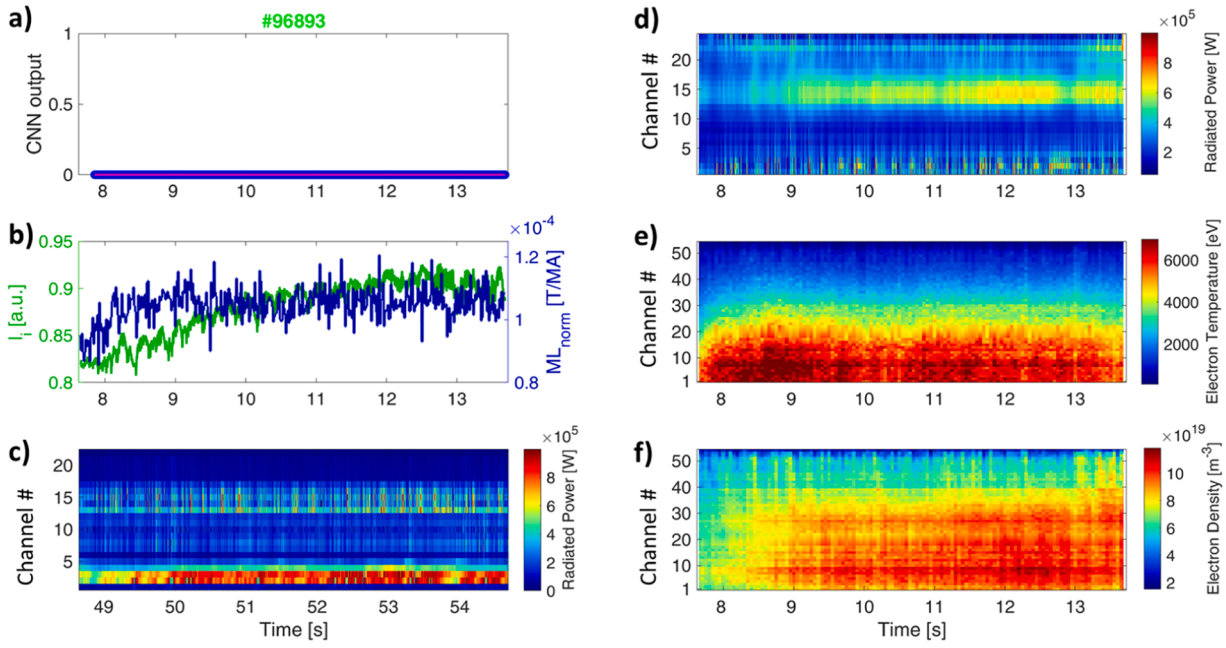

Fig. 7. JET regularly terminated discharge #96893. (a) CNN logic output curves, where the blue line is the top branch logic output, and the magenta line is the bottom branch logic output.; (b) internal inductance, in green, and mode lock normalized by the plasma current, in blue; c) radiated power from the bolometer vertical camera; (d) radiated power from the bolometer horizontal camera; (e) electron temperature from the HRTS; (f) electron density from the HRTS.

interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

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] B. Cannas, A. Fanni, E. Marongiu, P. Sonato, Disruption forecasting at JET using neural networks, Nucl. Fusion. 44 (2003) 68–76, https://doi.org/10.1088/0029-5515/44/1/008.

[2] B. Cannas, A. Fanni, P. Sonato, M.K.Z. and, A prediction tool for real-time application in the disruption protection system at JET, Nucl. Fusion. 47 (2007) 1559–1569, https://doi.org/10.1088/0029-5515/47/11/018.

[3] B. Cannas, R.S. Delogu, A. Fanni, P. Sonato, M.K. Zedda, Support vector machines for disruption prediction and novelty detection at JET, Fusion Eng. Des. 82 (2007) 1124–1130, https://doi.org/10.1016/j.fusengdes.2007.07.004.

[4] S. Dormido-Canto, J. Vega, J.M. Ramírez, A. Murari, R. Moreno, J.M. López, A. P. and, Development of an efficient real-time disruption predictor from scratch on JET and implications for ITER, Nucl. Fusion. 53 (2013), 113001, https://doi.org/10.1088/0029-5515/53/11/113001.

[5] A. Murari, M. Lungaroni, E. Peluso, P. Gaudio, J. Vega, S. Dormido-Canto, M. Baruzzo, M.G. and, Adaptive predictors based on probabilistic SVM for real time disruption mitigation on JET, Nucl. Fusion. 58 (2018), 056002, https://doi.org/10.1088/1741-4326/aaaf9c.

[6] A. Pau, A. Fanni, S. Carcangiu, B. Cannas, G. Sias, A. Murari, F.R. and, A machine learning approach based on generative topographic mapping for disruption prevention and avoidance at JET, Nucl. Fusion. 59 (2019), 106017, https://doi.org/10.1088/1741-4326/ab2ea9.

[7] E. Aymerich, A. Fanni, G. Sias, S. Carcangiu, B. Cannas, A. Murari, A. Pau, the J. contributors, A statistical approach for the automatic identification of the start of the chain of events leading to the disruptions at JET, Nucl. Fusion. 61 (2021), 036013, https://doi.org/10.1088/1741-4326/abcb28.

[8] E. Aymerich, G. Sias, F. Pisano, B. Cannas, S. Carcangiu, C. Sozzi, C. Stuart, P. Carvalho, A. Fanni, 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.

[9] E. Aymerich, B. Cannas, F. Pisano, G. Sias, C. Sozzi, C. Stuart, P. Carvalho, A. Fanni, the JET Contributors, Performance Comparison of Machine Learning Disruption Predictors at JET, Appl. Sci. 13 (2023) 2006, https://doi.org/10.3390/app13032006.

[10] B. Cannas, A. Fanni, G. Pautasso, G. Sias, P. Sonato, An adaptive real-time disruption predictor for ASDEX Upgrade, Nucl. Fusion. 50 (2010), 075004, https://doi.org/10.1088/0029-5515/50/7/075004.

[11] B. Cannas, A. Fanni, G. Pautasso, G. Sias, Disruption prediction with adaptive neural networks for ASDEX Upgrade, Fusion Eng. Des. 86 (2011) 1039–1044, https://doi.org/10.1016/j.fusengdes.2011.01.069.

[12] R. Aledda, B. Cannas, A. Fanni, A. Pau, G. Sias, Improvements in disruption prediction at ASDEX Upgrade, Fusion Eng. Des. 96–97 (2015) 698–702, https://doi.org/10.1016/j.fusengdes.2015.03.045.

[13] K.J. Montes, C. Rea, R.S. Granetz, R.A. Tinguely, N. Eidietis, O.M. Meneghini, D. L. Chen, B. Shen, B.J. Xiao, K. Erickson, M.D. Boyer, Machine learning for disruption warnings on Alcator C-Mod, DIII-D, and EAST, Nucl. Fusion. 59 (2019), 096015, https://doi.org/10.1088/1741-4326/ab1df4.

[14] W. Zheng, F.R. Hu, M. Zhang, Z.Y. Chen, X.Q. Zhao, X.L. Wang, P. Shi, X.L. Zhang, X.Q. Zhang, Y.N. Zhou, Y.N. Wei, Y. P, Hybrid neural network for density limit disruption prediction and avoidance on J-TEXT tokamak, Nucl. Fusion. 58 (2018), 056016, https://doi.org/10.1088/1741-4326/aaad17.

[15] C. Rea, R.S. Granetz, K. Montes, R.A. Tinguely, N. Eidietis, J.M. Hanson, B. Sammuli, Disruption prediction investigations using Machine Learning tools on DIII-D and Alcator C-Mod, Plasma Phys. Control. Fusion. 60 (2018), 084004, https://doi.org/10.1088/1361-6587/aac7fe.

[16] A. Pau, A. Fanni, B. Cannas, S. Carcangiu, G. Pisano, G. Sias, P. Sparapani, M. Baruzzo, A. Murari, F. Rimini, M. Tsalas, P.C. de Vries, A First Analysis of JET Plasma Profile-Based Indicators for Disruption Prediction and Avoidance, IEEE Trans. Plasma Sci. 46 (2018) 2691–2698, https://doi.org/10.1109/TPS.2018.2841394.

[17] J. Kates-Harbeck, A. Svyatkovskiy, W. Tang, Predicting disruptive instabilities in controlled fusion plasmas through deep learning, Nature 568 (2019) 526–531, https://doi.org/10.1038/s41586-019-1116-4.

[18] 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 (2020) 912–924, https://doi.org/10.1080/15361055.2020.1798589.

[19] P.C. de Vries, M. Baruzzo, G.M.D. Hogeweij, S. Jachmich, E. Joffrin, P.J. Lomas, G. F. Matthews, A. Murari, I. Nunes, T. Pütterich, C. Reux, J. Vega, JET-EFDA Contributors, The influence of an ITER-like wall on disruptions at JET, Phys. Plasmas. 21 (2014), 056101, https://doi.org/10.1063/1.4872017.

[20] Y. Bengio, Learning Deep Architectures for AI, Found, Trends® Mach. Learn. 2 (2009) 1–127, https://doi.org/10.1561/2200000006.

[21] 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. Plasmas. 27 (2020), 062510, https://doi.org/10.1063/1.5144458.

[22] 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 (2020), 026007, https://doi.org/10.1088/1741-4326/abc664.

[23] J. Garcia, F.J. Casson, C. Challis, D. Frigione, D. Van Eester, L. Garzotti, J. Hobirk, A. Kappatou, E.A. Lerche, J. Mailloux, F. Rimini, Integrated Scenario Development at JET for DT Operation and ITER Risk Mitigation, in: 2021. https://pure.mpg.de/pubman/faces/ViewItemOverviewPage.jsp?itemId=item_3320851 (accessed October 11, 2021).

[24] M.J. Leyland, M.N.A. Beurskens, J.C. Flanagan, L. Frassinetti, K.J. Gibson, M. Kempenaars, M. Maslov, R. Scannell, JET Contributors, Edge profile analysis of Joint European Torus (JET) Thomson scattering data: Quantifying the systematic error due to edge localised mode synchronisation, Rev. Sci. Instrum. 87 (2016), 013507, https://doi.org/10.1063/1.4939855.

[25] G. Pucella, P. Buratti, E. Giovannozzi, E. Alessi, F. Auriemma, D. Brunetti, D. R. Ferreira, M. Baruzzo, D. Frigione, L. Garzotti, E. Joffrin, E. Lerche, P.J. Lomas, S. Nowak, L. Piron, F. Rimini, C. Sozzi, D.V. Eester, J.E.T. Contributors, Onset of tearing modes in plasma termination on JET: the role of temperature hollowing and edge cooling, Nucl. Fusion. 61 (2021), 046020, https://doi.org/10.1088/1741-4326/abe3c7.