



EUROfusion

WPJET1-PR(18) 21265

DR Ferreira et al.

Applications of Deep Learning to Nuclear Fusion Research

Preprint of Paper to be submitted for publication in
Expert Systems with Applications



This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

This document is intended for publication in the open literature. It is made available on the clear understanding that it may not be further circulated and extracts or references may not be published prior to publication of the original when applicable, or without the consent of the Publications Officer, EUROfusion Programme Management Unit, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK or e-mail Publications.Officer@euro-fusion.org

Enquiries about Copyright and reproduction should be addressed to the Publications Officer, EUROfusion Programme Management Unit, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK or e-mail Publications.Officer@euro-fusion.org

The contents of this preprint and all other EUROfusion Preprints, Reports and Conference Papers are available to view online free at <http://www.euro-fusionscipub.org>. This site has full search facilities and e-mail alert options. In the JET specific papers the diagrams contained within the PDFs on this site are hyperlinked

Applications of Deep Learning to Nuclear Fusion Research*

Diogo R. Ferreira¹, JET Contributors*

EUROfusion Consortium, JET, Culham Science Centre, Abingdon, OX14 3DB, UK

¹Instituto de Plasmas e Fusão Nuclear, Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal

Abstract

Nuclear fusion is the process that powers the sun, and it is one of the best hopes to achieve a virtually unlimited energy source for the future of humanity. However, reproducing sustainable nuclear fusion reactions here on Earth is a tremendous scientific and technical challenge. Special devices – called tokamaks – have been built around the world, with JET (Joint European Torus, in the UK) being the largest tokamak currently in operation. Such devices confine matter and heat it up to extremely high temperatures, creating a plasma where fusion reactions begin to occur. JET has over one hundred diagnostic systems to monitor what happens inside the plasma, and each 30-second experiment (or pulse) generates about 50 GB of data. In this work, we show how convolutional neural networks (CNNs) can be used to reconstruct the 2D plasma profile inside the device based on data coming from those diagnostics. We also discuss how recurrent neural networks (RNNs) can be used to predict plasma disruptions, which are one of the major problems affecting tokamaks today. Training of such networks is done on NVIDIA GPUs.

1. Introduction

The Sun is the star at the center of our solar system, and nuclear fusion is the process that powers the Sun. At its hot and dense core, the Sun fuses hydrogen into helium, and since the result has less mass than the original elements, the excess mass is released as energy.

The idea of replicating this process on Earth (i.e. of using sustainable fusion reactions as a source of energy) dates back to at least the mid-20th century. For practical reasons, one of the most promising reactions is the one involving two isotopes of hydrogen:



where ${}^2_1\text{H}$ (*deuterium*) and ${}^3_1\text{H}$ (*tritium*) fuse into a helium nucleus, releasing a neutron and a relatively large amount of energy, when compared to other possible fusion reactions [1].

However, the engineering a fusion reactor has met some technical difficulties, the most important of which being how to successfully confine matter in a hot, plasma state. In the Sun, such problem is solved by the gravitational field, which holds the plasma together. On Earth, keeping

*This paper is based on a talk presented at NVIDIA's GPU Technology Conference, which took place at the International Congress Center in Munich, Germany, October 9–11, 2018 (GTC Europe 2018).

*See the author list of X. Litaudon et al 2017 Nucl. Fusion 57 102001

the plasma away from any other material (which would eventually cool it down) has turned out to be a surprisingly difficult challenge.

Several different types of fusion reactors have been devised based on the idea of confining the plasma with magnetic fields [2]. One of the most successful designs is the *tokamak* [3], which confines the plasma in a toroidal shape through the use of a toroidal field together with a poloidal field to yield an overall helical field around the torus.

At the macroscopic level, a plasma is approximately neutral (a concept that is referred to as *quasi-neutrality* [4]), but the fact that electrons and nuclei get stripped from each other at very high temperatures allows for the treatment of a plasma as a two-fluid model [5]. The two fluids (of positively and negatively charged particles) can be confined by magnetic fields.

The largest tokamak currently in operation is the Joint European Torus (JET) located near Oxford, in the UK. JET is a D-shaped tokamak with a major radius of 2.96 m and a minor radius in the range 1.25–2.10 m [6]. A larger machine with a similar design – the International Thermonuclear Experimental Reactor (ITER) – is currently being built in Cadarache, France.

JET has a vast assortment of diagnostic systems, including magnetic coils to measure plasma current and instabilities, interferometers and reflectometers to measure plasma density, Thomson scattering to determine the electron temperature, spectroscopy to measure ion temperature, and X-ray cameras to measure electromagnetic radiation, among many others [7].

One of the diagnostic systems that we will be focusing on – in fact, the only diagnostic that we will be using in this work – is the bolometer system available at JET [8]. This bolometer system measures the plasma radiation on a cross-section of the fusion device.

The system is illustrated in Figure 1. It comprises a horizontal camera and a vertical camera with multiple lines of sight. Some of these lines of sight are directed towards a region of special interest called the *divertor* [9], where a lot of activity takes place due to plasma exhaust.

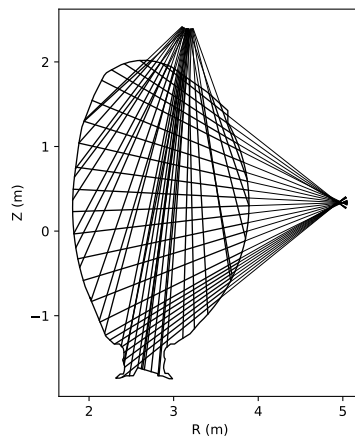


Figure 1: Lines of sight for the vertical and horizontal cameras of the bolometer system at JET.

The diagnostic is called a *bolometer system* because each camera comprises an array of sensors called *bolometers*. Each bolometer measures the radiation along a particular line of sight (technically, it provides a measure of the line-integrated radiation). In essence, a bolometer is a thin metal foil attached to a temperature-sensitive resistor [10]. As the metal foil absorbs radiation, its temperature changes and there is a change in resistance. This provides an effective way

to measure electromagnetic radiation in the range from ultraviolet (UV) to soft X-rays.

Using both the horizontal and the vertical camera, it is possible to reconstruct the 2D plasma radiation profile through a process called *tomography* [11]. The underlying principle is the same as Computed Tomography (CT) in medical applications [12], but the reconstruction technique is different due to the relatively scarce number of lines of sight available.

Figure 2 shows a sample tomographic reconstruction computed at JET. The reconstruction method involves an iterative constrained optimization approach that minimizes the error with respect to the observed measurements, while requiring the solution to be non-negative [13].

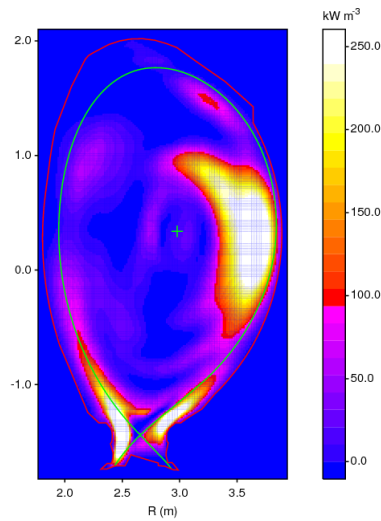


Figure 2: Sample tomographic reconstruction for pulse 89065 at $t = 47.8$ s.

This iterative reconstruction method takes a significant amount of computation time. The total run-time depends on the actual data but, with the code available at JET, it can take up to one hour to produce a single reconstruction.

To appreciate the computational effort involved, one should consider the following:

- The bolometer system has a sampling rate of 5 kHz.
- For the purpose of noise reduction, a window average of 5 ms is usually applied, which corresponds to 25 samples.
- Subsampling by a factor of 25 yields an effective sampling rate of 200 Hz.
- So, in principle, it should be possible to have 200 reconstructions per second.
- For a typical 30-second pulse at JET, this would mean a total of 6000 reconstructions.
- At one hour per reconstruction, this could take up to 250 days of computation time.

Clearly, another way to compute the reconstructions for an entire pulse should be found. The following section describes how deep learning was applied to train a convolutional neural network (CNN) that is able to compute 3000 reconstructions per second on an NVIDIA GPU [14].

2. Deep Learning for Plasma Tomography

Deep learning [15] has had a tremendous impact in several fields, such as image processing and natural language processing. In particular, convolutional neural networks (CNNs) have been very successful at classifying input images into a set of output classes. This has been demonstrated in the recognition of hand-written digits [16], in the classification of Web images [17], and in the annotation of online videos [18], to cite only a few examples.

In general, CNNs have a structure that comprises one or more convolutional layers. Each convolutional layer applies multiple filters (in the form of a small kernel or sliding window) to the input. The purpose of a filter is to detect a specific feature, so its output is called a *feature map*. In a CNN, the first convolutional layer operates directly on the input image; subsequent layers operate on the feature maps produced by previous layers.

In addition to convolutional layers, CNNs contain subsampling layers. There is usually one subsampling layer after each convolutional layer. The purpose of having subsampling is to make the feature maps smaller and allow their number to progressively increase after each convolution, while keeping the network under a manageable size.

After the convolutional and subsampling layers, the end result is usually a large number of small feature maps. These are then flattened and connected to a couple of dense layers, which perform the actual classification. The main idea behind a CNN is to have a first stage of convolutional layers to extract meaningful features from the input image, and a second stage of dense layers to perform classification based on those features.

In summary, the input to a CNN is typically a 2D image and its output is a 1D vector of class probabilities. However, for the purpose of plasma tomography, it would be useful to have a network that takes a 1D vector of bolometer measurements as input, and produces a 2D image of the plasma radiation profile as output.

In the literature, the inverse of a CNN has been referred to as a *deconvolutional network*, and it has found applications in image segmentation [19] and object generation [20]. For example, by specifying the class label and the camera position (1D data), a deconvolutional network is able to generate an object (2D image) of the specified class from the given camera view [21].

In this work, we use a deconvolutional network that receives the bolometer measurement as input (a total of 56 lines of sight from both cameras) and produces a 120×200 reconstruction of the plasma radiation profile. The network architecture is shown in Figure 3.

At the beginning, there are two dense layers with 7500 nodes, which can be reshaped into a 3D structure of size $25 \times 15 \times 20$. This structure can be regarded as comprising 20 feature maps of size 25×15 . By applying a series of transposed convolutions, these feature maps are brought up to a size of 200×120 , from which the output image is generated by one last convolution.

Essentially, the *transposed convolution* is the inverse of the convolution in the sense that, if a sliding window would be applied to the output, the result would be the feature map given as input. It can be shown that learning a transposed convolution is equivalent to learning a weight matrix that is the transpose of a regular convolution, hence the name of this operation [22].

In Figure 3, the upsampling of feature maps from 25×15 up to 200×120 is achieved by having each transposed convolution operate with a stride of two pixels (i.e. one pixel is being skipped between each two consecutive positions of the sliding window). This means that the output is two times larger and taller than the input, except for the very last convolution which uses a stride of one to keep the same size.

To train the network shown in Figure 3, we gathered the tomographic reconstructions that have been computed at JET for all the experimental campaigns between 2011 and 2016. This

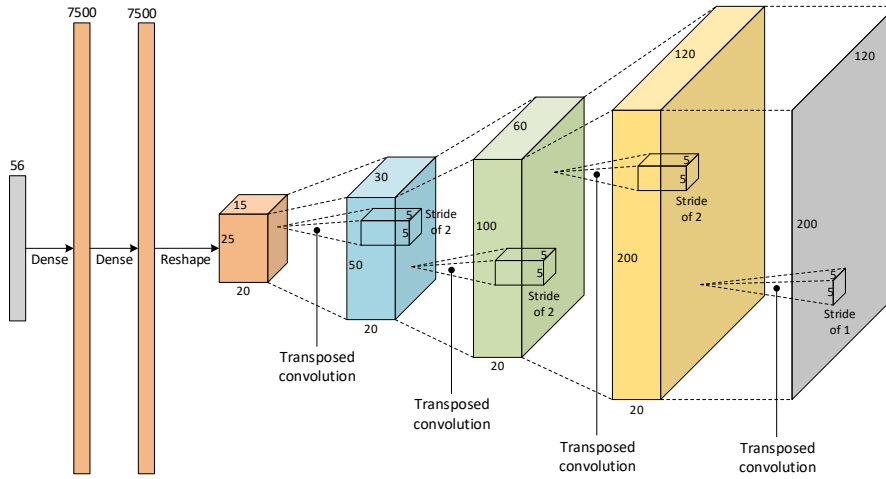


Figure 3: Deconvolutional neural network for plasma tomography.

yielded a total of 25584 reconstructions, which have been divided into 90% (23025) for training, and 10% (2559) for validation.

The network was trained on an NVIDIA Tesla P100 GPU using accelerated gradient descent (Adam [23]) with a small learning rate (10^{-4}) and a batch size of 307 samples. The batch size was chosen to be a perfect divisor of the number of training samples ($23025/307 = 75$) in order to avoid having any partially filled batch in the training set.

Figure 4 shows the evolution of the training loss and of the validation loss across 3530 epochs, which took about 20 hours. The best result was achieved at epoch 1765, with a minimum validation loss of 8.556 kW m^{-3} . This is the mean absolute error per pixel, which can be compared to the dynamic range of Figure 2 ($8.556/250 \approx 3.4\%$ error).

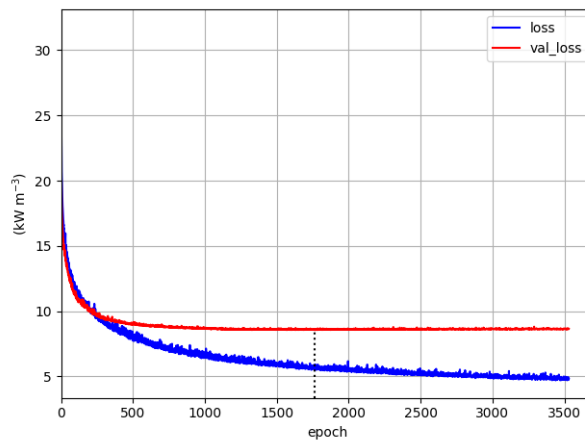


Figure 4: Loss and validation loss during training of the deconvolutional network.

Once trained, the network can be used to generate the reconstruction for any given pulse at any given point in time. In fact, given the bolometer data from an entire pulse, the network can generate the reconstructions for all points in time. This way, it becomes possible to analyze the evolution of the plasma radiation profile across the entire experiment.

Figure 5 shows the reconstruction of pulse 92213 from $t = 49.62$ s onwards, with a time increment of 0.1 s (due to space restrictions only). The first row shows a focus of radiation developing on the outer wall, which seems to slowly fade away (rows 1–2), only to reappear later at the plasma core with particularly strong intensity (rows 2–3).

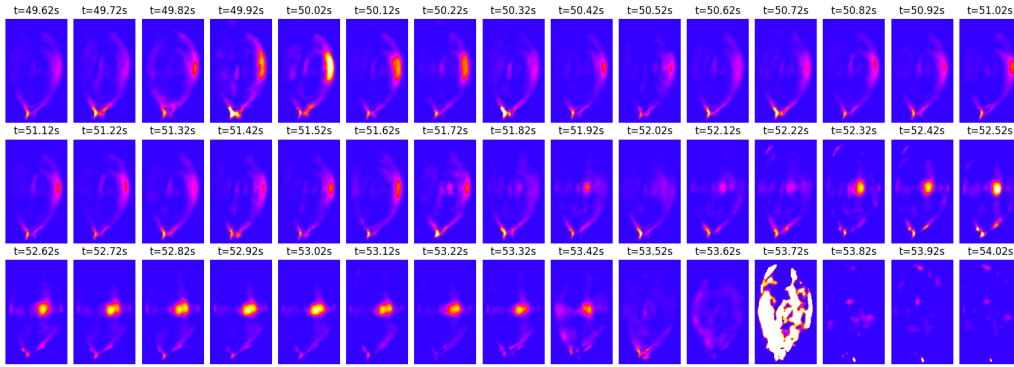


Figure 5: Reconstruction of pulse 92213 from $t = 49.62$ s to $t = 54.02$ s with a time step of 0.1 s.

The radiation peaking stays at the core for a relatively long time (at least from $t = 52.32$ s to $t = 53.42$ s) while changing slightly in shape during that interval. Eventually, it also fades away as the heating systems are being turned off. However, just as it seemed that the plasma was about to soft land, there is a disruption at $t = 53.72$ s.

Disruptions are one of the major problems affecting tokamaks, and they are one of the major impediments towards fusion energy today. The dynamics of disruptions are not yet fully understood, but one of the main causes seems to be impurity accumulation at the plasma core, which decreases the core temperature and eventually leads to core collapse [24].

To some extent, this phenomenon is clearly visible in Figure 5. It is therefore not surprising that the bolometer system has played a key role in several disruption studies [25, 26, 27, 28], either by providing a measurement of total radiation or by providing the tomographic reconstruction at a few, specific points in time. With the proposed deconvolutional network, it is possible to quickly obtain the plasma radiation profile for all points in time.

3. Deep Learning for Disruption Prediction

Over the years, a number of different approaches to disruption prediction have been developed, including the use of neural networks [29, 30, 31, 32], support vector machines [33, 34, 35] and decision trees [36, 37], among other techniques. In all of these approaches, the prediction model takes a set of global plasma parameters as input (such as the plasma current, the locked mode amplitude, the total input power, the safety factor, the poloidal beta, etc.) and outputs a probability of disruption, or the remaining time to a predicted disruption.

More recently, the use of random forests [38] and even deep learning [39] have shown promising results, but the input signals that are provided to these models still consist of about ten global plasma parameters derived from different diagnostics.

In the previous section, we have seen that a bolometer system with multiple lines of sight is able to capture some of the physical phenomena associated with plasma disruptions. From such bolometer data, it is possible to reconstruct the internal shape of the plasma radiation profile and observe, for example, an impurity concentration developing at the plasma core, which could be the precursor of an impending disruption.

This suggests that the array of sensor measurements coming from multiple cameras of a bolometer system could, in principle, be used for disruption prediction. The main difference to previous works is that the prediction model will be working on the basis of a diagnostic with multiple views into the plasma, rather than on a set of global plasma parameters.

In addition to having a spatial view into the plasma, another dimension that will be important to consider is time. As suggested by Figure 5, the phenomena that precede a disruption will develop and endure for some time, on the order of seconds. Therefore, it will be important to consider not only the bolometer measurements at a certain point in time, but also the preceding measurements in order to have a sense of how the plasma behavior is developing (e.g. to determine whether it is moving closer to, or away from, a possible disruption).

With both dimensions, the input to the predictor will be sequence of vectors, where each vector contains the bolometer measurements from the multiple lines of sight for a given point in time. This closely resembles the way in which recurrent neural networks are used in natural language processing (NLP), where each word in a sentence (the sequence) is mapped into a vector space of a certain dimensionality, through a process called *word embedding* [40].

A recurrent neural network (RNN) is able to handle sequences of vectors by processing each vector at a time. The key distinctive feature of RNNs is that they maintain an internal state, and this internal state is fed back to the input along with the next vector in the sequence. A RNN can therefore “remember” features from past vectors in the sequence.

In simple RNNs, such “memory” is necessarily short-lived, since the network keeps updating its internal state as it receives new vectors, and gradually “forgets” about vectors that are in the distant past. To provide RNNs with the ability to maintain long-term dependencies, more advanced architectures have been proposed, namely the *long short-term memory* (LSTM) [41, 42]. In an LSTM, information can be carried across multiple time steps, and that information can be updated or forcefully forgotten at each time step, through the use of an *input gate* and a *forget gate* respectively, while also maintaining an internal state.

More recently, the combination of CNNs and RNNs has been shown to be very effective in tasks such as video processing [43, 44] and text classification [45, 46]. In essence, the CNN is used as a feature extractor before passing the data to a RNN for sequential processing.

In this work, we use a similar approach to preprocess the bolometer data through a series of 1D convolutional layers before handing them over to an LSTM. Besides feature extraction, this has the advantage of further reducing the noise in the bolometer data. As in a traditional CNN, we use subsampling after each convolutional layer to progressively reduce the length of the time sequence, while increasing the number of filters that are applied to it.

The resulting network is shown in Figure 6. The input has a shape of 200×56 , with the first being the sequence length, and the latter being the vector size. We use sequences of 200 time points, which corresponds to the 200 Hz sampling rate of the bolometer system, meaning that a full second of bolometer measurements is analyzed by the network in order to produce a prediction. The vector size of 56 corresponds to the number of lines of sight available.

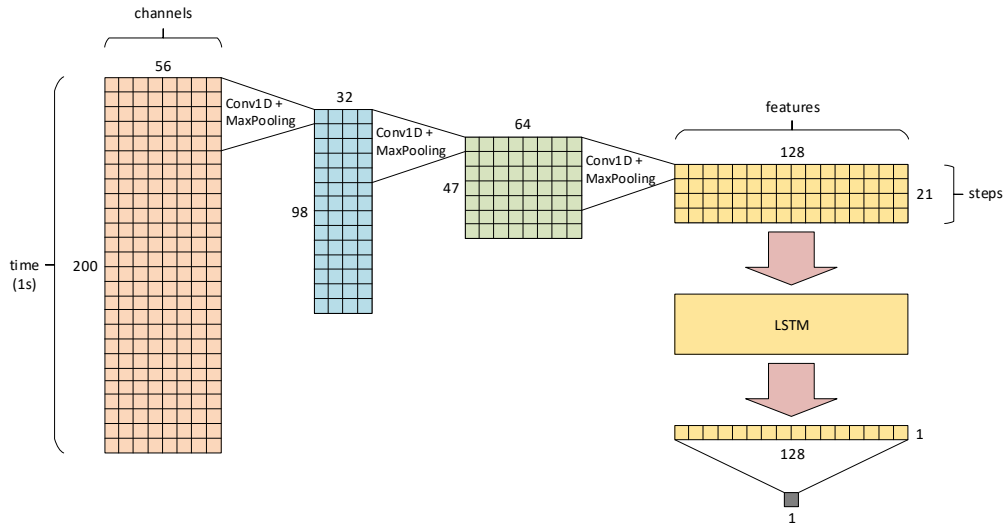


Figure 6: Recurrent neural network for disruption prediction.

Some of these lines of sight yield always a measurement of zero, since they are reserve channels that typically remain unused. A few other lines of sight are associated with bolometers that are known to be faulty, providing erratic measurements. In these cases, the network is expected to filter out those channels, so the number of filters used in the first convolutional layer (32) is actually less than the number of input channels (56).

In subsequent layers, the number of filters doubles after each convolution. A kernel size of 5 time steps is used in all convolutions, together with a max-pooling operation in order to reduce the sequence length to half. Eventually, the convolutional part of the network yields a sequence of 21 steps, where each step contains a 128-feature vector.

This sequence of feature vectors is handed over to an LSTM with 128 units, which returns the last output after having processed the whole sequence. One final dense layer combines these vector elements into a single output, which is the network prediction.

As mentioned earlier, there are two ways of predicting disruptions:

- One way is to predict the *probability of disruption*, i.e. the probability that the current pulse will end in a disruption. For this purpose, the last layer of the network will contain a sigmoid activation, to produce a value between 0 and 1. The loss function used to train the network will be the binary cross-entropy [47], and the network will be trained with sample data from both disruptive and non-disruptive pulses.
- Another way is to predict the *time to disruption*, i.e. the remaining time up to a incoming disruption. For this purpose, the last layer of the network will have a linear activation (i.e. no activation) so that it can produce positive as well as negative values within any range. The loss function used to train the network is the mean absolute error, and the network is trained with sample data from disruptive pulses only.

To train these two variants, we gathered the plasma current and the bolometer data for all pulses in the experimental campaigns from 2011 to 2016. The time of disruption was defined

as the point in time where the current gradient exceeded 20 MA/s, as is standard practice at JET [48]. From a universe of 9798 pulses in those experimental campaigns, only 1683 pulses (about 17%) reached such threshold and were therefore considered to be disruptive. Hence:

- To train the network for probability of disruption, we divided the 9798 pulses into 90% for training and 10% for validation. We used a data generator to draw random samples from each pulse and feed them to the network, together with a binary label to indicate whether the sample was drawn from a disruptive or non-disruptive pulse.
- To train the network for time to disruption, we divided the 1683 disruptive pulses into 90% for training and 10% for validation. We used a data generator to draw random samples from each pulse and feed them to the network, together with a label to indicate the timespan between the sample time and the disruption time for that pulse.

In both cases, a batch size of 2000 samples was used, where each sample is a 200×56 sequence of vectors, and the sample time is defined as the latest time point in the sequence. Each training epoch was defined as a run through 100 batches. For validation purposes, a similar data generator was used to draw random samples from the validation pulses, with the same batch size, but using only 10 batches for validation at the end of each training epoch.

Figure 7 shows the evolution of the training loss and of the validation loss for both variants:

- For probability of disruption, training ran for a total of 1118 epochs (14 hours on an NVIDIA Tesla P100 GPU), with the best result being achieved at epoch 559 with a validation loss (binary cross-entropy) of 0.202.
- For time to disruption, training ran for 1026 epochs (13 hours on the same GPU) with the best result being achieved at epoch 513 with a validation loss (mean absolute error) of 1.986 seconds.

Although these loss values appear to be relatively high, it should be noted that they are averaged over all samples used for training/validation. For samples that are far away from the disruption, their loss value will be quite high, because it is difficult to predict when or if a disruption will occur. This will be compensated by a small loss value for samples that are near the disruption, when it will be more important to have a correct prediction.

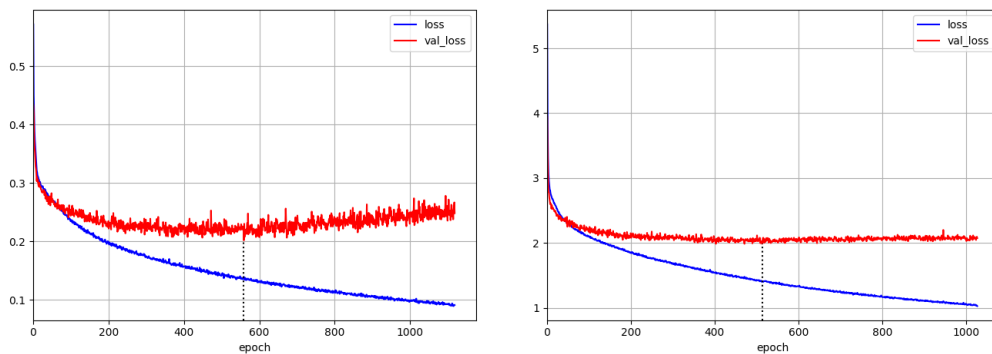


Figure 7: Loss and validation loss for probability of disruption (*left*) and time to disruption (*right*).

Figure 8 shows an example of how the predicted time to disruption (tt_d) and the predicted probability of disruption (prd) evolve across an entire pulse. There is an upward trend in the probability of disruption and a downward trend in the time to disruption. Eventually, as the pulse gets close to the disruption, the probability gets close to 1.0 and the time to disruption approaches zero, indicating that a disruption is imminent.

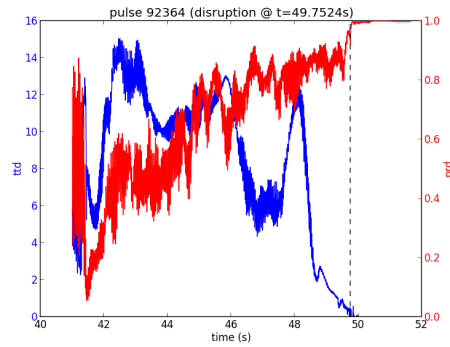


Figure 8: Time to disruption and probability of disruption for pulse 92364.

As soon as the probability of disruption exceeds a certain threshold *and* the time to disruption falls below a certain threshold, it is possible to raise an alarm that will trigger preventive actions, such as activating the *disruption mitigation valve* (DMV) [28] which, in a matter of a few milliseconds, will inject a massive amount of gas to cool down the plasma.

Such alarm should not be triggered when the time to disruption is low but the probability of disruption is also low, or when the probability of disruption is high but the time to disruption is also high, as illustrated in the examples of Figure 9. Both of these pulses are non-disruptive.

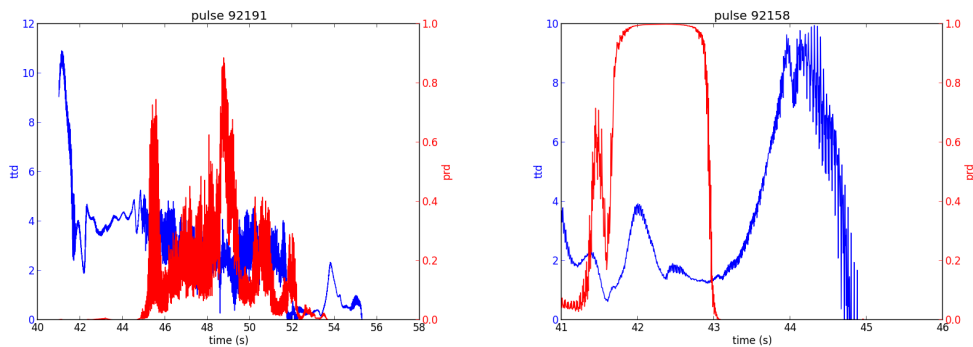


Figure 9: Time to disruption and probability of disruption for pulses 92191 (left) and 92158 (right).

A surprising result from this approach is that, for several disruptive pulses, the probability of disruption gets close to 1 early on, and stays there for almost the entire pulse until a disruption eventually occurs. Figure 10 shows two examples. Of course, an alarm would only be triggered when the time to disruption also approaches zero, but it is still puzzling how the network was

able to recognize signs of disruption so early in the pulse. We hypothesize that the network has learned that certain radiation patterns are always associated with disruptive pulses, even if those radiation patterns are not related to the actual cause of disruption.

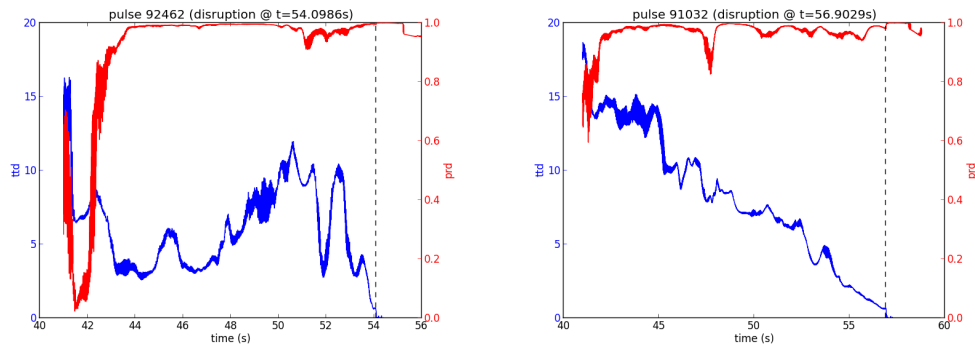


Figure 10: Time to disruption and probability of disruption for pulses 92462 (*left*) and 91032 (*right*).

In Figure 11, we show an example of how false alarms may occur. During the pulse, the probability of disruption becomes very high, when the time to disruption is also very low. This would raise an alarm, and therefore trigger mitigating actions, when in fact the pulse is non-disruptive. To some extent, such false alarms are difficult to avoid completely, but should be kept to a minimum so as not to jeopardize legitimate experiments.

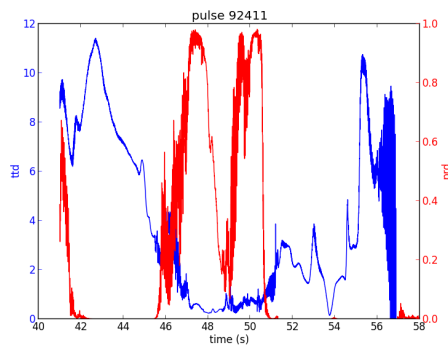


Figure 11: Time to disruption and probability of disruption for pulse 92411.

Finally, in Figure 12 we show an example of the (fortunately) rare situation of missed alarms. At the time of disruption, the probability is low and the time to disruption is still high, which would not raise an alarm. Even if the time to disruption drops to zero and the probability hits 1 at disruption time, this would probably be too late for any mitigating action.

In general, when an alarm is raised and the pulse is disruptive, this is a *true positive* (TP). False alarms correspond to *false positives* (FP) and missed alarms correspond to *false negatives* (FN). When no alarm is raised and the pulse is non-disruptive, it is a *true negative* (TN).

With an upper threshold of 0.85 for the probability of disruption, and a lower threshold of 1.5

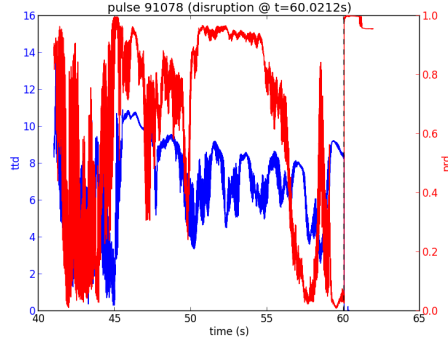


Figure 12: Time to disruption and probability of disruption for pulse 91078.

seconds for the time to disruption, we can correctly classify (as positive or negative) 85.5% of the validation pulses. The remaining 14.5% are divided into false alarms (10.1%) and missed alarms (4.4%). It is possible to tweak these thresholds in order to further reduce the missed alarms, but at the cost of increasing the false alarms. It is also possible to tweak the thresholds so that the rate of false alarms and missed alarms is approximately the same.

Table 1 presents the results when varying the thresholds within reasonable limits. The first row and the first column specify the conditions which will trigger an alarm. It should be noted that only alarms that have been raised up to 30 ms before the disruption are considered, since this is the minimum time to perform mitigation actions at JET [49]. The results in bold (diagonal) illustrate the trade-off between false alarms (FP) and missed alarms (FN).

	prd \geq 0.95	prd \geq 0.90	prd \geq 0.85	prd \geq 0.80
ttd \leq 0.5 s	TP: 4.4% TN: 81.4% FP: 2.7% FN: 11.5%	TP: 6.0% TN: 79.5% FP: 4.6% FN: 9.9%	TP: 7.2% TN: 78.2% FP: 5.9% FN: 8.7%	TP: 8.5% TN: 76.9% FP: 7.1% FN: 7.4%
ttd \leq 1.0 s	TP: 6.8% TN: 81.0% FP: 3.1% FN: 9.1%	TP: 9.0% TN: 78.0% FP: 6.1% FN: 6.9%	TP: 10.3% TN: 75.9% FP: 8.2% FN: 5.6%	TP: 11.0% TN: 74.0% FP: 10.1% FN: 4.9%
ttd \leq 1.5 s	TP: 8.1% TN: 80.3% FP: 3.8% FN: 7.9%	TP: 10.2% TN: 76.6% FP: 7.4% FN: 5.7%	TP: 11.5% TN: 74.0% FP: 10.1% FN: 4.4%	TP: 12.2% TN: 71.6% FP: 12.4% FN: 3.7%
ttd \leq 2.0 s	TP: 8.8% TN: 79.5% FP: 4.6% FN: 7.1%	TP: 11.1% TN: 75.5% FP: 8.6% FN: 4.8%	TP: 12.4% TN: 72.3% FP: 11.7% FN: 3.5%	TP: 13.3% TN: 69.8% FP: 14.3% FN: 2.7%

Table 1: Prediction results on the validation set (980 pulses) with different alarm-triggering thresholds

The results obtained here are comparable to those reported in the literature. For example, the disruption predictor that is currently being used at JET (APODIS) has been reported to have a

success rate of 85.38% and a false alarm rate of 2.46% [50]. This seems to be within reach of the results presented in Table 1, with a choice of thresholds such as $prd \geq 0.95$ and $ttd \leq 0.5$ s.

Furthermore, when APODIS triggers an alarm for a disruptive pulse, its average warning time is 350 ms in advance of the disruption [50]. With $prd \geq 0.95$ and $ttd \leq 0.5$ s, we get an average warning time of 929 ms, almost a full second ahead of the disruption.

The average warning time is even higher for every other configuration in Table 1, because increasing the ttd threshold or decreasing the prd threshold has the effect of triggering the alarm sooner. This will, on the other hand, increase the rate of false alarms.

4. Conclusion

Deep learning finds several applications in the analysis data coming from fusion diagnostics. Here we have showcased the use of deep neural networks for plasma tomography and disruption prediction. In the first case, a CNN has been trained to replace the computationally expensive task of reconstructing the plasma profile after each experiment. In the second case, a RNN has been trained to detect the onset of plasma disruptions, which can be extremely useful for tokamak operation by raising timely alarms in advance of incoming disruptions.

This means that deep neural networks can be used for both post-processing and real-time processing of fusion data. For real-time processing, fast inference (using special-purpose GPUs, such as the NVIDIA Tesla T4) will be of utmost importance.

Acknowledgements

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014–2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. IPFN activities received financial support from Fundação para a Ciência e Tecnologia (FCT) through project UID/FIS/50010/2013.

References

- [1] F. F. Chen, Introduction to Plasma Physics and Controlled Fusion, 3rd Edition, Springer, 2016.
- [2] J. Freidberg, Ideal MHD, Cambridge University Press, 2014.
- [3] K. Lackner, S. Günter, P. Lauber, G. Pautasso, B. Scott, Equilibrium and macroscopic stability of tokamaks, in: M. Kikuchi, K. Lackner, M. Q. Tran (Eds.), Fusion Physics, IAEA, 2012, Ch. 3, pp. 225–359.
- [4] P. M. Bellan, Fundamentals of Plasma Physics, Cambridge University Press, 2006.
- [5] H. Goedbloed, S. Poedts, Principles of Magnetohydrodynamics, Cambridge University Press, 2004.
- [6] L. D. Horton, Joint European Torus, in: G. H. Neilson (Ed.), Magnetic Fusion Energy: From Experiments to Power Plants, Elsevier, 2016, Ch. 8, pp. 215–259.
- [7] J. Wesson, Tokamaks, 4th Edition, Oxford University Press, 2011.
- [8] A. Huber, K. McCormick, P. Andrew, P. Beaumont, S. Dalley, J. Fink, J. Fuchs, K. Fullard, W. Fundamenski, L. Ingesson, F. Mast, S. Jachmich, G. Matthews, P. Mertens, V. Philipps, R. Pitts, S. Sanders, W. Zeidner, Upgraded bolometer system on JET for improved radiation measurements, Fusion Engineering and Design 82 (5) (2007) 1327–1334.
- [9] G. Matthews, P. Coad, H. Greuner, M. Hill, T. Hirai, J. Likonen, H. Maier, M. Mayer, R. Neu, V. Philipps, R. Pitts, V. Riccardo, Development of divertor tungsten coatings for the JET ITER-like wall, Journal of Nuclear Materials 390–391 (2009) 934–937.
- [10] L. C. Ingesson, B. Alper, B. J. Peterson, J.-C. Vallet, Chapter 7: Tomography diagnostics: Bolometry and soft-x-ray detection, Fusion Science and Technology 53 (2) (2008) 528–576.
- [11] L. Ingesson, B. Alper, H. Chen, A. Edwards, G. Fehmers, J. Fuchs, R. Giannella, R. Gill, L. Lauro-Taroni, M. Romanelli, Soft X ray tomography during ELMs and impurity injection in JET, Nuclear Fusion 38 (11) (1998) 1675.

- [12] G. T. Herman, *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*, 2nd Edition, Springer, 2009.
- [13] G. C. Fehmers, L. P. J. Kamp, F. W. Sluijter, An algorithm for quadratic optimization with one quadratic constraint and bounds on the variables, *Inverse Problems* 14 (4) (1998) 893.
- [14] D. R. Ferreira, P. J. Carvalho, H. Fernandes, JET Contributors, Full-pulse tomographic reconstruction with deep neural networks, *Fusion Science and Technology* 74 (1-2) (2018) 47–56.
- [15] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (2015) 436–444.
- [16] Y. Lecun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition, *Proceedings of the IEEE* 86 (11) (1998) 2278–2324.
- [17] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Advances in Neural Information Processing Systems* 25 (2012) 1097–1105.
- [18] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, L. Fei-Fei, Large-scale video classification with convolutional neural networks, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 1725–1732.
- [19] H. Noh, S. Hong, B. Han, Learning deconvolution network for semantic segmentation, in: *IEEE International Conference on Computer Vision*, 2015, pp. 1520–1528.
- [20] A. Dosovitskiy, J. T. Springenberg, T. Brox, Learning to generate chairs with convolutional neural networks, in: *IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1538–1546.
- [21] A. Dosovitskiy, J. T. Springenberg, M. Tatarchenko, T. Brox, Learning to generate chairs, tables and cars with convolutional networks, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39 (4) (2017) 692–705.
- [22] V. Dumoulin, F. Visin, A guide to convolution arithmetic for deep learning, arXiv:1603.07285.
- [23] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980.
- [24] P. C. de Vries, M. F. Johnson, B. Alper, P. Buratti, T. C. Hender, H. R. Koslowski, V. Riccardo, Survey of disruption causes at JET, *Nuclear Fusion* 51 (5) (2011) 053018.
- [25] G. Arnoux, A. Loarte, V. Riccardo, W. Fundamenski, A. Huber, Heat loads on plasma facing components during disruptions on JET, *Nuclear Fusion* 49 (8) (2009) 085038.
- [26] V. Riccardo, G. Arnoux, P. Cahyna, T. C. Hender, A. Huber, S. Jachmich, V. Kiptily, R. Koslowski, L. Krlin, M. Lehnen, A. Loarte, E. Nardon, R. Paprok, D. Tshakaya, JET disruption studies in support of ITER, *Plasma Physics and Controlled Fusion* 52 (12) (2010) 124018.
- [27] A. Huber, G. Arnoux, M. Beurskens, S. Bozhnikov, S. Brezinsek, T. Eich, C. Fuchs, W. Fundamenski, S. Jachmich, U. Kruezi, M. Lehnen, A. Loarte, G. Matthews, P. Mertens, P. Morgan, V. Philipps, R. Pitts, V. Riccardo, U. Samm, B. Schweer, G. Sergienko, M. Stamp, Radiation loads onto plasma-facing components of JET during transient events – experimental results and implications for ITER, *Journal of Nuclear Materials* 415 (1) (2011) S821–S827.
- [28] M. Lehnen, A. Alonso, G. Arnoux, N. Baumgarten, S. A. Bozhnikov, S. Brezinsek, M. Brix, T. Eich, S. N. Gerasimov, A. Huber, S. Jachmich, U. Kruezi, P. D. Morgan, V. V. Plyusnin, C. Reux, V. Riccardo, G. Sergienko, M. F. Stamp, Disruption mitigation by massive gas injection in JET, *Nuclear Fusion* 51 (12) (2011) 123010.
- [29] G. Pautasso, C. Tichmann, S. Egorov, T. Zehetbauer, O. Gruber, M. Maraschek, K.-F. Mast, V. Mertens, I. Perchermeier, G. Raupp, W. Treutterer, C. Windsor, A. U. Team, On-line prediction and mitigation of disruptions in ASDEX Upgrade, *Nuclear Fusion* 42 (1) (2002) 100.
- [30] B. Cannas, A. Fanni, E. Marongiu, P. Sonato, Disruption forecasting at JET using neural networks, *Nuclear Fusion* 44 (1) (2004) 68.
- [31] C. Windsor, G. Pautasso, C. Tichmann, R. Buttery, T. Hender, JET EFDA Contributors, the ASDEX Upgrade Team, A cross-tokamak neural network disruption predictor for the JET and ASDEX Upgrade tokamaks, *Nuclear Fusion* 45 (5) (2005) 337.
- [32] R. Yoshino, Neural-net predictor for beta limit disruptions in JT-60U, *Nuclear Fusion* 45 (11) (2005) 1232.
- [33] B. Cannas, R. Delogu, A. Fanni, P. Sonato, M. Zedda, Support vector machines for disruption prediction and novelty detection at JET, *Fusion Engineering and Design* 82 (5) (2007) 1124–1130.
- [34] G. Rattá, J. Vega, A. Murari, G. Vagliasindi, M. Johnson, P. de Vries, JET EFDA Contributors, An advanced disruption predictor for JET tested in a simulated real-time environment, *Nuclear Fusion* 50 (2) (2010) 025005.
- [35] J. M. López, J. Vega, D. Alves, S. Dormido-Canto, A. Murari, J. M. Ramírez, R. Felton, M. Ruiz, G. de Arcas, Implementation of the disruption predictor APODIS in JET’s real-time network using the MARTe framework, *IEEE Transactions on Nuclear Science* 61 (2) (2014) 741–744.
- [36] A. Murari, G. Vagliasindi, P. Arena, L. Fortuna, O. Barana, M. Johnson, JET-EFDA Contributors, Prototype of an adaptive disruption predictor for JET based on fuzzy logic and regression trees, *Nuclear Fusion* 48 (3) (2008) 035010.
- [37] A. Murari, J. Vega, G. Rattá, G. Vagliasindi, M. Johnson, S. Hong, JET-EFDA Contributors, Unbiased and non-supervised learning methods for disruption prediction at JET, *Nuclear Fusion* 49 (5) (2009) 055028.
- [38] C. Rea, R. S. Granetz, Exploratory machine learning studies for disruption prediction using large databases on DIII-D, *Fusion Science and Technology* 74 (1–2) (2018) 89–100.

- [39] A. Svyatkovskiy, J. Kates-Harbeck, W. Tang, Training distributed deep recurrent neural networks with mixed precision on GPU clusters, in: *Proceedings of Machine Learning on HPC Environments (MLHPC'17)*, 2017, pp. 10:1–10:8.
- [40] Y. Goldberg, *Neural Network Methods in Natural Language Processing*, Vol. 37 of *Synthesis Lectures on Human Language Technologies*, Morgan & Claypool, 2017.
- [41] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Computation* 9 (8) (1997) 1735–1780.
- [42] F. A. Gers, N. N. Schraudolph, J. Schmidhuber, Learning precise timing with LSTM recurrent networks, *Journal of Machine Learning Research* 3 (2002) 115–143.
- [43] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, in: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 2625–2634.
- [44] J. Donahue, L. A. Hendricks, M. Rohrbach, S. Venugopalan, S. Guadarrama, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39 (4) (2017) 677–691.
- [45] J. Wang, L.-C. Yu, K. R. Lai, X. Zhang, Dimensional sentiment analysis using a regional CNN-LSTM model, in: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Vol. 2: Short Papers)*, 2016, pp. 225–230.
- [46] X. Wang, W. Jiang, Z. Luo, Combination of convolutional and recurrent neural network for sentiment analysis of short texts, in: *Proceedings of the 26th International Conference on Computational Linguistics (COLING 2016)*, 2016, pp. 2428–2437.
- [47] I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, 2016.
- [48] J. Morris, *Disruption Chain1 Code Documentation*, JET-EFDA, Culham Science Centre, Abingdon, OX14 3DB, UK (May 2012).
- [49] P. de Vries, M. Johnson, I. Segui, J. E. Contributors, Statistical analysis of disruptions in JET, *Nuclear Fusion* 49 (5) (2009) 055011.
- [50] R. Moreno, J. Vega, S. Dormido-Canto, A. Pereira, A. Murari, JET Contributors, Disruption prediction on JET during the ILW experimental campaigns, *Fusion Science and Technology* 69 (2) (2016) 485–494.