

Monitoring the plasma radiation profile with real-time bolometer tomography at JET

Diogo R. Ferreira^{a,*}, Pedro J. Carvalho^b, Ivo S. Carvalho^b, Chris Stuart^b, Peter J. Lomas^b, JET Contributors¹

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

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

^bCulham Centre for Fusion Energy, UK Atomic Energy Authority, Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, UK

Abstract

The use of real-time tomography at JET opens up new possibilities for monitoring the plasma radiation profile and for taking preventive or mitigating actions against impending disruptions. By monitoring the radiated power in different plasma regions, such as core, edge and divertor, it is possible to set up multiple alarms for the radiative phenomena that usually precede major disruptions. The approach is based on the signals provided by the bolometer diagnostic. Reconstructing the plasma radiation profile from these signals is a computationally intensive task, which is typically performed during post-pulse analysis. To reconstruct the radiation profile in real-time, we use machine learning to train a surrogate model that performs matrix multiplication over the bolometer signals. The model is trained on a large number of sample reconstructions, and is able to compute the plasma radiation profile within a few milliseconds in real-time. The implementation has been further optimized by computing the radiated power only in the regions of interest. Experimental results show that, during uncontrolled termination, there is an impurity accumulation at the plasma core, which eventually leads to a disruption. A threshold-based alarm on core radiation, among other options, is able to anticipate a significant fraction of such disruptions.

Keywords: Plasma Tomography, Real-Time Systems, Machine Learning, GPU Computing

1. Introduction

The JET bolometer system [1] is a diagnostic that measures the radiated power through an array of sensors, called bolometers, where each bolometer measures the line-integrated radiation along a specific line of sight. The system includes two cameras – a horizontal camera and a vertical camera – with 24 bolometers each. Using these two cameras, it is possible to reconstruct the plasma radiation profile on a poloidal cross-section of the device. The two cameras and their lines of sight are illustrated in Figure 1. Essentially, the geometry is such that, for each camera, 16 lines of sight cover the entire plasma region, and 8 lines of sight provide a more fine-grained resolution of the divertor region at the bottom of the vessel.

From the bolometer measurements, there are several tomography techniques that can be applied to reconstruct the plasma radiation profile [2]. The method that is used at JET uses an iterative constrained optimization algorithm that minimizes the error with respect to the measurements, while requiring the solution to be non-negative [3]. This iterative procedure takes a significant amount of time, typically on the order of minutes to produce a single reconstruction. More recently, it was shown that it is possible to produce essentially the same results with a deep neural network, at the cost of a small amount of error, but several orders of magnitude faster [4].

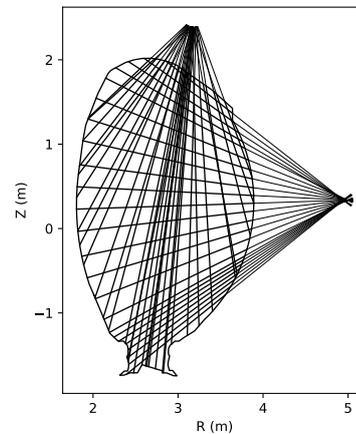


Figure 1: Lines of sight for the bolometer system at JET.

Similarly, there has been a recent surge of interest in using machine learning to develop surrogate models, especially in the form of neural networks [5], to accelerate the processing of diagnostics data. However, such models can hardly run on a resource-constrained environment such as the real-time network at JET. For example, deep learning models [6] often require specialized hardware such as Graphics Processing Units (GPUs) to be trained and run, but there are no such facilities in the real-time environment at JET. Therefore, any approach that aims at reconstructing the plasma radiation profile in real-time must be based on a significantly simpler model.

*Corresponding author

Email address: diogo.ferreira@tecnico.ulisboa.pt (Diogo R. Ferreira)

¹See the author list of E. Joffrin et al. 2019 Nucl. Fusion 59 112021

In this work, we use machine learning to train a simple model for plasma tomography based on matrix multiplication. Although the model is trained on GPU, once it is trained it can run on CPU, and even on a limited CPU, since it requires only a single matrix multiplication step, which can be implemented in a variety of programming languages, and can be performed sufficiently fast (on the order of milliseconds) even on limited hardware. This approach is especially convenient for implementation in the real-time environment at JET, which is based on the MARTe framework [7], and allows new features to be configured flexibly on top of the RTCC system [8].

The availability of the plasma radiation profile in real-time opens new possibilities to monitor the plasma behavior, especially the radiative phenomena that often precede plasma disruptions. An example is impurity accumulation, which leads to a pattern where most of the radiated power emanates from the plasma core. As we will illustrate below, this pattern is clearly visible in the plasma profile. In extreme cases, core radiation may exceed the input power and cause a disruption by radiative collapse [9], so monitoring the radiated power in the core region is one of the possible uses of real-time tomography, and this is the application that we describe here. We also discuss the results that can be achieved by setting up an alarm on core radiation in the real-time control system at JET.

2. Tomography by matrix multiplication

As described in the previous section, the bolometer system at JET has two cameras with 24 lines of sight each. On the other hand, the plasma profiles that are routinely produced at JET have a resolution of 196×115 pixels. This means that there are 22540 unknowns to be computed from only 48 measurements, which illustrates the ill-posed nature of the problem.

Since each bolometer measures the line-integrated radiation along its line of sight, and the plasma radiation profile has a discrete representation (in the form of a rectangular grid that can be flattened into a column vector), then this setup can be represented as a linear system in the form $\mathbf{y} = \mathbf{A}\mathbf{x}$, where \mathbf{y} are the bolometer measurements and \mathbf{x} is the radiation profile to be reconstructed. This system is typically under-determined because the number of measurements in \mathbf{y} is much less than the number of (unknown) power density values in \mathbf{x} . Therefore, the system cannot be inverted simply as $\mathbf{x} = \mathbf{A}^{-1}\mathbf{y}$.

However, given some sample measurements $\tilde{\mathbf{y}}$ and a sample reconstruction $\tilde{\mathbf{x}}$ that has been obtained from those measurements using some existing tomography technique, it might be possible to train a model \mathbf{M} by minimizing an error measure such as $\|\tilde{\mathbf{x}} - \mathbf{M}(\tilde{\mathbf{y}})\|$. In previous work [4], \mathbf{M} was a deep neural network with several layers, which can be regarded as a nonlinear function applied over the bolometer data \mathbf{y} to produce the plasma profile \mathbf{x} . In the present work, \mathbf{M} is just a matrix, and $\mathbf{M}(\tilde{\mathbf{y}}) = \mathbf{M}\tilde{\mathbf{y}}$ becomes a simple matrix multiplication step. Once trained, the matrix \mathbf{M} can be used to compute the radiation profile from bolometer data as $\mathbf{x} \cong \mathbf{M}\mathbf{y}$.

This approach extends naturally to batches of training data $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$. In this case, both $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$ are matrices where each column corresponds to a different training sample. Then \mathbf{M} is

trained by minimizing $\|\tilde{\mathbf{X}} - \mathbf{M} \cdot \tilde{\mathbf{Y}}\|$, and can also be used to compute a batch of reconstructions via $\mathbf{X} \cong \mathbf{M} \cdot \mathbf{Y}$.

More precisely:

- $\mathbf{Y} = \{Y_{jk}\}$ is a $48 \times N$ matrix representing a batch of bolometer data, where 48 is the number of lines of sight and N is an arbitrary batch size.
- $\mathbf{M} = \{M_{ij}\}$ is a 22540×48 matrix representing the model parameters, where 22540 is the total number of pixels in the plasma profile.
- $\mathbf{X} = \{X_{ik}\}$ with $X_{ik} = \sum_{j=1}^{48} M_{ij}Y_{jk}$ is a $22540 \times N$ matrix representing a batch of N plasma profiles reconstructed from bolometer data.

Given a batch of training data $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) = (\{\tilde{X}_{ik}\}, \{\tilde{Y}_{jk}\})$, the model is trained by gradient descent to minimize the mean absolute error:

$$L(\mathbf{M}) = \frac{1}{22540 \times N} \sum_{i=1}^{22540} \sum_{k=1}^N |\tilde{X}_{ik} - \sum_{j=1}^{48} M_{ij}\tilde{Y}_{jk}| \quad (1)$$

In summary, the model has 22540×48 parameters, where each of its 22540 outputs is generated by a weighted sum of its 48 inputs, and the weights are learned by minimizing the loss function in Eq. (1).

For the present work, we collected a training dataset comprising about 10 000 of the most recent reconstructions that are routinely produced at JET. These are reconstructions across a wide range of pulses (75000–97000) and at specific points of interest for the analysis and modeling tasks at JET. The focus on the most recently produced reconstructions is justified by the fact that, since 2017, all reconstructions have been double-checked to fix problems in the bolometer data and to avoid profiles with noticeable artifacts. This ensures that the training profiles are of the highest possible quality.

Using a 32-bit floating point representation (4 bytes), the training data takes $10000 \times 48 \times 4 + 10000 \times 22540 \times 4 \cong 904$ MB, which can fit the memory of a single GPU. Therefore, we loaded the training data into GPU memory and trained the model by running each epoch over the full dataset rather than using mini-batching. With the dataset pre-loaded into GPU memory, this allows the model to be trained efficiently for a large number of epochs and/or with a small learning rate, without any data transfers between CPU and GPU.

We separated the data into 90% for training and 10% for validation, while making sure that there were no sample reconstructions from the same pulse in the training set and in the validation set. The model was trained by gradient descent implemented directly with TensorFlow, and using the Adam optimizer [10] with a learning rate of 10^{-4} . As illustrated in Figure 2, the mean absolute error in the validation set converged to a value of about 0.0355 MW/m^3 after 13300 epochs.

For comparison, a deep neural network, which is a much more sophisticated model, achieved a mean absolute error of 0.0128 MW/m^3 [11]. Although the training/validation dataset is not the same for both models, this provides a sense of how

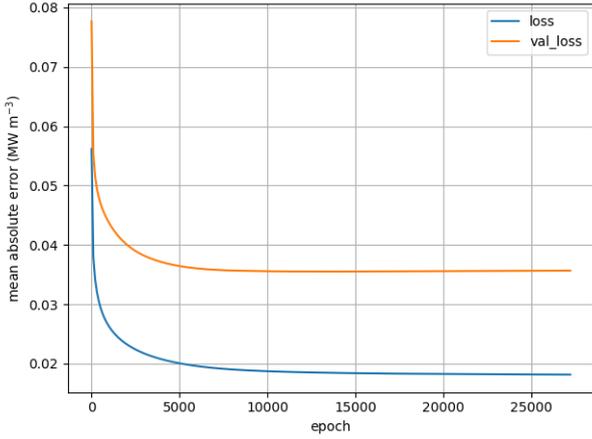


Figure 2: Evolution of loss and validation loss during training.

much error we incur by reducing the tomographic reconstruction process to a single matrix multiplication step. As an example, Figure 3 shows an average case where the mean absolute error between the original reconstruction and the reconstruction produced by the model is exactly 0.0355 MW/m^3 .

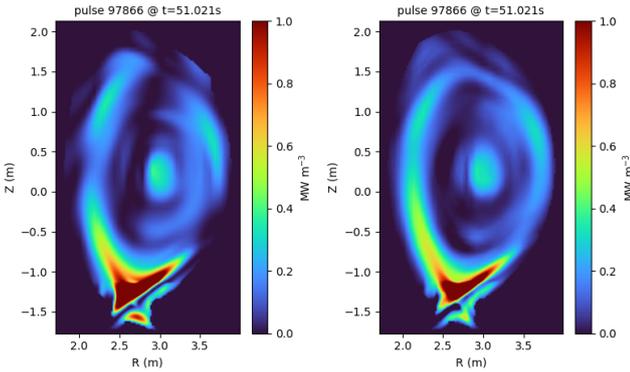


Figure 3: Sample reconstruction from validation set (*left*) compared to the reconstruction produced by the model (*right*).

Figure 4 shows the average error per pixel computed using the validation set. We can see that the model tends to slightly overestimate the radiated power on certain regions, and underestimate it others. The most troublesome region seems to be the divertor, where the model overestimates the radiation above the X-point, and underestimates it below that point. Some of the underestimation in the divertor region appears to be related to a slight overestimation in the core region, since there are vertical lines of sight traversing both regions.

For real-time deployment, matrix \mathbf{M} is exported as an array and written to a file that can be imported by other programs, written in different programming languages. Nowadays, many languages include standard libraries with matrix multiplication routines, but even if that is not the case, matrix multiplication can be implemented using low-level instructions or, in some cases, even vectorized operations that take advantage of the parallel instructions available in modern processors.

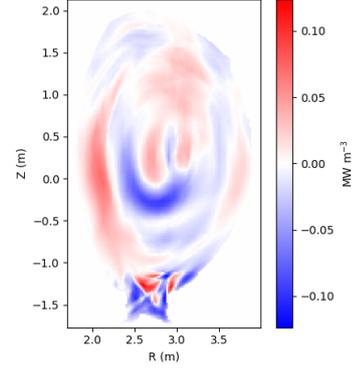


Figure 4: Average error per pixel in the validation set.

3. Tomography in real-time

Figure 5 shows a sequence of frames illustrating the real-time reconstructions produced by the model for a sample pulse. Despite the model simplicity, all the usual features that often appear in the radiation profile are clearly recognizable. In this example, it is possible to observe the development of a radiation blob at the outboard edge, followed by the development of a radiation blob at the plasma core, which eventually leads to a disruption at $t = 53.7\text{s}$. The two radiation blobs are possibly (but not necessarily) related by impurity transport from edge to core, since at around $t = 52.0\text{s}$ the outboard radiation decreases, followed by a sudden increase in core radiation.

The main advantage of real-time tomography is that it allows monitoring the radiated power in different regions of interest, namely at the edge, at the core, and at the divertor. In particular, core radiation turns out to be an important precursor of plasma disruptions [12], so one of the main features to be monitored is the radiated power in that region.

When the purpose is to monitor a specific region, a possible approach is to apply a mask over the radiation profile in order to consider only the points that lie inside that region. Figure 6 shows the masks that have been defined for the core, edge and divertor regions. Another approach is to reconstruct only the region of interest, while avoiding to compute the entire radiation profile. This is the approach that has been implemented at JET to monitor the radiated power in the core region.

Figure 7 illustrates the signals that correspond to the total radiated power and to the radiated power in the core, outboard edge, and divertor regions. Again, it is possible to observe that at some point around $t = 52.0\text{s}$ the edge radiation decreases, followed by a sharp increase in core radiation. The tomographic reconstruction on the right-hand side of Figure 7 illustrates the buildup of core radiation at that point.

From these signals, it becomes apparent that when core radiation exceeds a certain threshold, this corresponds to the formation of a radiation blob in the plasma core. If this blob is allowed to stay there, then the radiation losses will lead to a decrease of the core temperature and the development of a hollow temperature profile. In turn, this will cause changes to the resistivity, to the current profile and to the q -profile, setting the stage

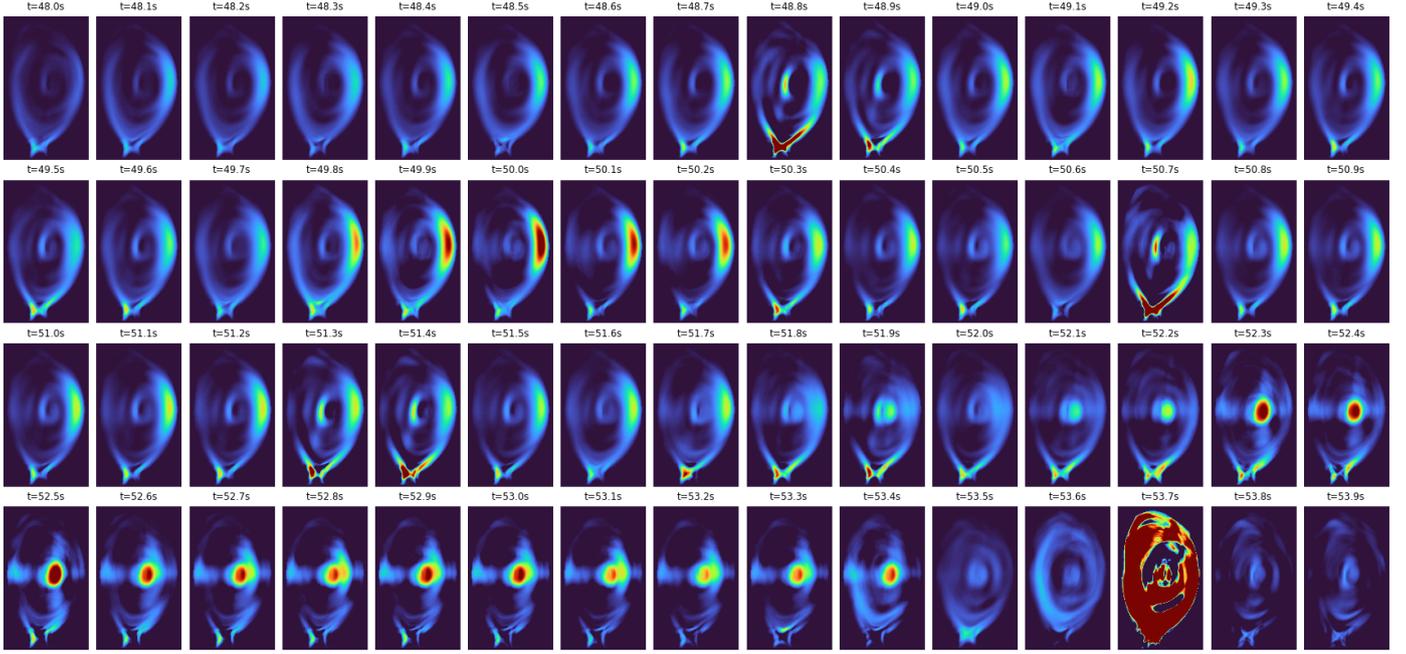


Figure 5: Plasma radiation profiles for pulse 92213 from $t=48.0\text{s}$ to $t=53.9\text{s}$ with a time step of 0.1s and a dynamic range of 1 MW m^{-3} .

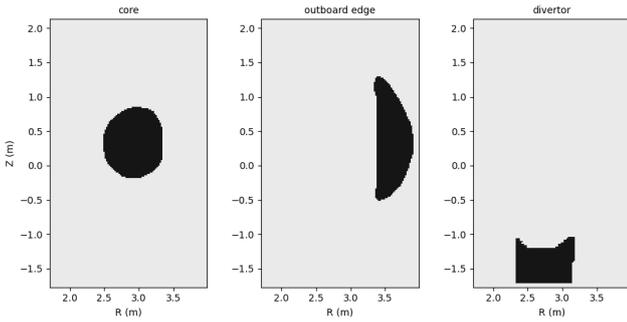


Figure 6: Masks for core, outboard edge and divertor regions.

for MHD (magnetohydrodynamic) instabilities to appear. One way to counteract these effects is to apply auxiliary heating, namely ICRH (ion cyclotron resonance heating), to re-heat the core and recover from a hollow temperature profile, creating the conditions for a safe termination [13]. On other devices, different means of auxiliary heating may be employed to achieve the same purpose; for example, on ASDEX Upgrade it is possible to counteract impurity accumulation through the use of ECRH (electron cyclotron resonance heating) [14].

In any case, our main purpose is to be able to detect the onset of such conditions, and for this purpose we used the PETRA (Plasma Event TRiggering and Alarms) system to configure an alarm on core radiation, such that the alarm is triggered when core radiation exceeds a threshold of 3 MW . In the example of Figure 7, such threshold is exceeded at $t = 52.29\text{s}$, which is about 1.4s before the disruption, giving ample time for disruption avoidance or mitigating actions.

4. Core radiation threshold

The threshold of 3 MW on core radiation has been chosen by analyzing the baseline pulses from two recent campaigns at JET, which amounts to about 400 pulses in the last 5 years. The aim of the JET baseline scenario [15] is to develop a recipe for achieving high fusion performance and sustained fusion power with a view towards D-T (Deuterium-Tritium) operation. However, as the plasma current and heating power are being increased, a higher rate of pulse disruptivity is also being observed (about 40%). This makes it especially important to address disruptivity in this operational scenario.

In this context, the core radiation threshold can be regarded as playing a similar role to a disruption predictor. For disruptive pulses, if core radiation exceeds the threshold before the disruption, we have a successful prediction (true positive); on the other hand, if core radiation exceeds the threshold too late or does not exceed the threshold at all, we have a missed alarm (false negative). For non-disruptive pulses, if core radiation exceeds the threshold at any point in time, we have a false alarm (false positive); if core radiation does not exceed the threshold, we have a correct prediction (true negative).

With this interpretation, it is possible to assess the effectiveness of the core radiation threshold in terms of binary classification metrics such as precision, recall, etc. However, for our purposes, an analysis based on true positives, true negatives, false positives and false negatives will suffice.

Figure 8 illustrates the effectiveness of the core radiation threshold to anticipate disruptions in the JET baseline scenario. Each dot in the chart represents a baseline pulse. The dot is positioned vertically according to whether the pulse is disruptive or not, and horizontally according to the maximum core radiation achieved during that pulse (and before the disruption, if

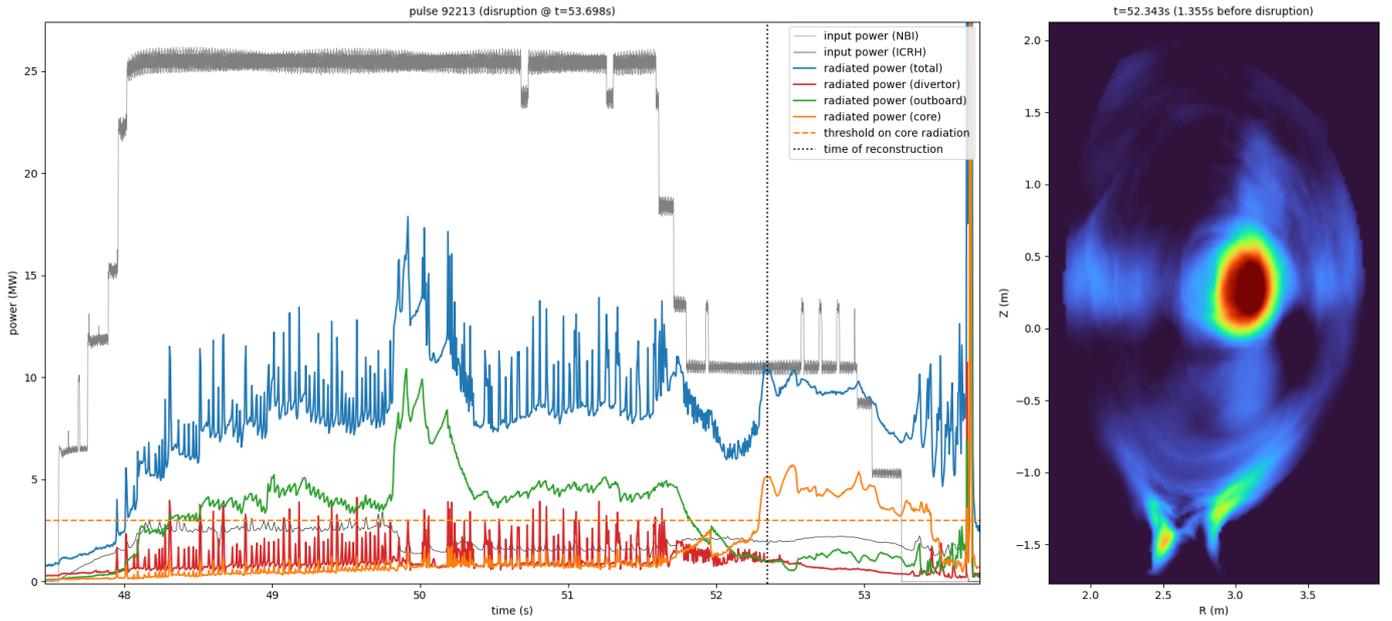


Figure 7: Radiated power in different regions of interest (*left*) and a tomographic reconstruction of the plasma profile at $t = 52.3s$ showing a peak in core radiation (*right*) for JET pulse 92213.

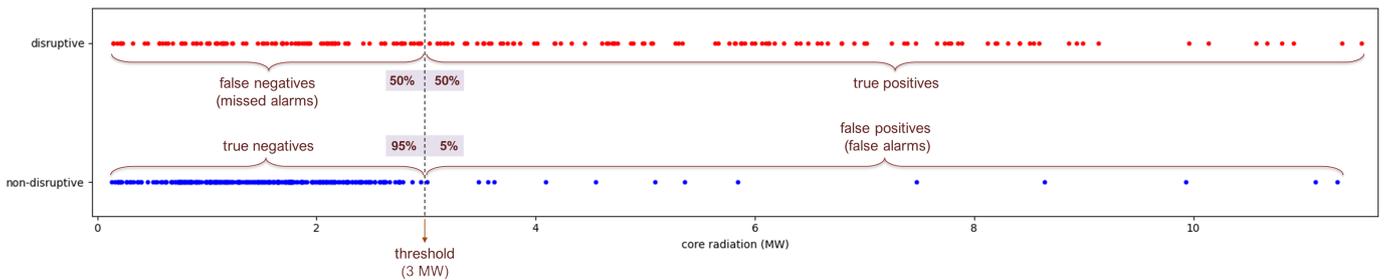


Figure 8: Classification accuracy when using a threshold on core radiation as a disruption predictor.

disruptive). We can observe that most non-disruptive pulses are concentrated towards the left on the horizontal axis (low core radiation), while disruptive pulses are distributed over a much wider range (low to high core radiation).

The analysis shows that, with a threshold of 3 MW on core radiation, it is possible to anticipate 50% of disruptive pulses, at the expense of a false alarm rate of only 5%. The remaining 50% of disruptions must be attributed to factors other than core radiation, and should be further investigated.

Regarding the false alarms (false positives), some of these are genuine cases of pulses with high core radiation that eventually recover, but there also cases of pulses that have been misclassified as non-disruptive, or where high core radiation is due to spurious measurements.

Regarding the true positives, the average warning time is 0.8 seconds, so the example in Figure 7 is an especially favorable one, where core radiation appears relatively early. For other pulses, core radiation may show up closer to the disruption, typically 200 to 600 ms before the disruption time.

5. Conclusion

The buildup of core radiation, which is attributed to impurity accumulation in that region, was found to be responsible for about 50% of disruptions in the JET baseline scenario. Such core radiation shows up in the plasma radiation profile as a distinct pattern with a radiation blob centered at the plasma core. To monitor the development of this and other patterns in the plasma radiation profile, we developed an approach to bolometer tomography based on a single matrix multiplication step, which was implemented in real-time, within the computational constraints of the real-time environment at JET.

Real-time tomography allows monitoring the radiated power in specific regions of interest, and provides the possibility of setting up custom alarms depending on the physics that are most relevant for a given operational scenario. In the present work, we focused on the JET baseline scenario, and on setting up an alarm on core radiation. A similar kind of alarm may become useful in other operational scenarios, or even in other devices. We plan to investigate this in future work.

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 and 2019-2020 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

Instituto de Plasmas e Fusão Nuclear (IPFN) received financial support from *Fundação para a Ciência e Tecnologia* (FCT) through projects UIDB/50010/2020 and UIDP/50010/2020.

The authors are thankful for the granted use of computational resources provided by CCFE / UKAEA at Culham, UK.

References

- [1] A. Huber, K. McCormick, P. Andrew, P. Beaumont, S. Dalley, J. Fink, J. C. Fuchs, K. Fullard, W. Fundamenski, L. C. Ingesson, F. Mast, S. Jachmich, G. F. Matthews, P. Mertens, V. Philipps, R. A. Pitts, S. Sanders, W. Zeidner, Upgraded bolometer system on JET for improved radiation measurements, *Fusion Engineering and Design* 82 (5) (2007) 1327–1334.
- [2] J. Mlynar, T. Craciunescu, D. R. Ferreira, P. Carvalho, O. Ficker, O. Grover, M. Imrisek, J. Svoboda, Current research into applications of tomography for fusion diagnostics, *Journal of Fusion Energy* 38 (3) (2019) 458–466.
- [3] L. C. Ingesson, B. Alper, H. Chen, A. W. Edwards, G. C. Fehmers, J. C. Fuchs, R. Giannella, R. D. Gill, L. Lauro-Taroni, M. Romanelli, Soft X ray tomography during ELMs and impurity injection in JET, *Nuclear Fusion* 38 (11) (1998) 1675–1694.
- [4] D. R. Ferreira, P. J. Carvalho, H. Fernandes, Full-pulse tomographic reconstruction with deep neural networks, *Fusion Science and Technology* 74 (1-2) (2018) 47–56.
- [5] A. Pavone, J. Svensson, S. Kwak, M. Brix, R. C. W. and, Neural network approximated bayesian inference of edge electron density profiles at JET, *Plasma Physics and Controlled Fusion* 62 (4) (2020) 045019.
- [6] J. Kates-Harbeck, A. Svyatkovskiy, W. Tang, Predicting disruptive instabilities in controlled fusion plasmas through deep learning, *Nature* 568 (7753) (2019) 526–531.
- [7] A. C. Neto, F. Sartori, F. Piccolo, R. Vitelli, G. De Tommasi, L. Zabeo, A. Barbalace, H. Fernandes, D. F. Valcarcel, A. J. N. Batista, MARTe: A multiplatform real-time framework, *IEEE Transactions on Nuclear Science* 57 (2) (2010) 479–486.
- [8] M. Lennholm, I. Carvalho, K. Cave-Ayland, A. Chagnard, C. Challis, R. Felton, D. Frigione, L. Garzotti, A. Goodyear, J. Graves, C. Guillemaut, J. Harrison, E. Lerche, P. Lomas, R. Mooney, F. Rimini, A. Sips, C. Sozzi, D. Valcarcel, J. Vega, Real time control developments at JET in preparation for deuterium-tritium operation, *Fusion Engineering and Design* 123 (2017) 535–540.
- [9] 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.
- [10] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: 3rd International Conference on Learning Representations, 2014.
- [11] D. R. Ferreira, P. J. Carvalho, H. Fernandes, Deep learning for plasma tomography and disruption prediction from bolometer data, *IEEE Transactions on Plasma Science* 48 (1) (2020) 36–45.
- [12] 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, The influence of an ITER-like wall on disruptions at JET, *Physics of Plasmas* 21 (5) (2014) 056101.
- [13] E. Lerche, M. Goniche, P. Jacquet, D. V. Eester, V. Bobkov, L. Colas, C. Giroud, I. Monakhov, F. Casson, F. Rimini, C. Angioni, M. Baruzzo, T. Blackman, S. Brezinsek, M. Brix, A. Czarnecka, K. Crombé, C. Challis, R. Dumont, J. Eriksson, N. Fedorczak, M. Graham, J. Graves, G. Gorini, J. Hobirk, E. Joffrin, T. Johnson, Y. Kazakov, V. Kiptily, A. Krivska, M. Lennholm, P. Lomas, C. Maggi, P. Mantica, G. Mathews, M.-L. Mayoral, L. Meneses, J. Mlynar, P. Monier-Garbet, M. Nave, C. Noble, M. Nocente, I. Ongena, J. Ongena, G. Petravich, V. Petrzilka, T. Pütterich, M. Reich, M. Santala, E. Solano, A. Shaw, G. Sips, M. Stamp, M. Tardocchi, M. Tsalas, M. Valisa, Optimization of ICRH for core impurity control in JET-ILW, *Nuclear Fusion* 56 (3) (2016) 036022.
- [14] J. Stober, M. Reisner, C. Angioni, A. B. Navarro, V. Bobkov, A. Bock, G. Denisov, E. Fable, R. Fischer, G. Gantenbein, L. Gil, T. Görler, V. Igochine, W. Kasperek, F. Leuterer, A. Litvak, R. McDermott, A. Meier, F. Monaco, M. Munich, V. Nichiporenko, B. Plaum, U. Plank, E. Poli, L. Popov, T. Pütterich, T. Scherer, M. Schubert, W. Suttrop, E. Tai, M. Thumm, D. Wagner, H. Zohm, Exploring fusion-reactor physics with high-power electron cyclotron resonance heating on ASDEX Upgrade, *Plasma Physics and Controlled Fusion* 62 (2) (2020) 024012.
- [15] L. Garzotti, C. Challis, R. Dumont, D. Frigione, J. Graves, E. Lerche, J. Mailloux, M. Mantsinen, F. Rimini, F. Casson, A. Czarnecka, J. Eriksson, R. Felton, L. Frassinetti, D. Gallart, J. Garcia, C. Giroud, E. Joffrin, H.-T. Kim, N. Krawczyk, M. Lennholm, P. Lomas, C. Lowry, L. Meneses, I. Nunes, C. M. Roach, M. Romanelli, S. Sharapov, S. Silburn, A. Sips, E. Stefániková, M. Tsalas, D. Valcarcel, M. Valovič, Scenario development for D–T operation at JET, *Nuclear Fusion* 59 (7) (2019) 076037.