Investigating the physics of disruptions with real-time tomography at JET

Diogo R. FERREIRA[1], Pedro J. CARVALHO[1], Ivo S. CARVALHO[1], Chris I. STUART[2], Peter J. LOMAS[2] and JET Contributors*

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

[1] Instituto de Plasmas e Fusão Nuclear (IPFN), Instituto Superior Técnico (IST), Universidade de Lisboa, Av. Rovisco Pais 1, 1049-001 Lisboa, Portugal

[2] Culham Centre for Fusion Energy (CCFE), UK Atomic Energy Authority (UKAEA), Culham Science Centre, Abingdon, OX14 3DB, Oxfordshire, UK

Abstract

As JET is developing and testing operational scenarios for higher fusion performance, an increase in pulse disruptivity is being observed. On a deeper analysis, we find that several radiative phenomena play an active role in determining the outcome of the pulse. The analysis is enabled by the use of real-time tomography based on the bolometer diagnostic. Even though plasma tomography is an inverse problem, we use machine learning to train a forward model that provides the radiation profile directly, based on a single matrix multiplication step. This model is used to investigate radiative phenomena including sawtooth crashes, ELMs and MARFE, and their relationship to the radiated power in different regions of interest. In particular, we use real-time tomography to monitor the core region, and to throw an alarm whenever core radiation exceeds a certain threshold. Our results suggest that this measure alone can anticipate a significant fraction of disruptions in the JET baseline scenario.

Keywords: plasma tomography, machine learning, sawtooth crashes, edge localized modes (ELMs), multifaceted asymmetric radiation from the edge (MARFE), plasma disruptions

(Some figures may appear in colour only in the online journal)

1 Introduction

The JET baseline scenario [1] is being developed for deuterium-tritium (DT) operation, with a view towards high fusion performance and sustained fusion power, using a normalized pressure ratio of about $\beta_N \sim 1.8$ and an edge safety factor of $q_{95} \sim 3$. However, as the plasma current and heating power

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

are being increased, a higher rate of pulse disruptivity is also being observed. Currently, the need to bring disruptivity under control is one of the most pressing issues in scenario development.

In the past, several different causes have been identified to explain the occurrence of disruptions at JET [2]. Those causes included multiple physical phenomena as well as a variety of technical issues. The study was conducted when JET still had the carbon wall. More recently, after the introduction of the ITER-like wall [3], it was found that present disruptions are more closely related to radiative phenomena [4], by a sequence of events that includes impurity transport from the edge to the core, impurity accumulation at the core, development of strong core radiation, and radiative collapse when the radiated power eventually exceeds the input power.

In this context, it becomes important to monitor the plasma radiation profile for any behaviour that may indicate that the experiment is heading towards a possible disruption. One of such behaviours is the development of strong core radiation, but there may be other physical phenomena that could be relevant as well.

In this work, we explore a range of physical phenomena that can be observed with real-time tomography based on the bolometer diagnostic at JET [5]. For this purpose, we developed a real-time implementation of plasma tomography, based on a simplification of the original reconstruction process [6]. In contrast with the original method, which uses an iterative procedure to find the plasma radiation profile as the best solution that agrees with the measurements, we simplified the reconstruction process to a single matrix multiplication step that can be employed in real-time.

The approach makes use of a simplified model, in the form of an inversion matrix, that has been trained by machine learning on a large collection of sample reconstructions, which have been previously computed at JET using the original method. Such inversion matrix provides the means to transform a set of input measurements (the bolometer data) into an output image (the plasma radiation profile). Although the idea of training an inversion matrix by machine learning has been introduced in previous work [7], our goal here is to focus on the range of physical phenomena that can be observed and studied by using such model to reconstruct the plasma radiation profile.

The use of such inversion matrix also makes it possible to monitoring the plasma radiation profile in real-time and, in particular, it enables creating alarms on conditions that are thought to be disruption-relevant, such as the development of core radiation. We took advantage of this possibility by setting an alarm in the new PETRA system [8], which provides real-time event detection at JET. The alarm goes off whenever core radiation exceeds 3 MW. This value was found by judicious choice, after a systematic study of core radiation across baseline pulses that we also report here.

The paper is structured as follows. Section 2 introduces the bolometer diagnostic and its lines of sight, as well as the tomographic inversion methods that can be used to reconstruct the plasma radiation profile. Section 3 describes the machine learning approach that has been employed to train a simplified model for plasma tomography to be used in real-time. Section 4 discusses several radiative phenomena that can be observed with real-time tomography. Finally, section 5 provides a study of core radiation across baseline pulses, and explains how we arrived at the threshold of 3 MW.

2 Bolometer tomography

The JET bolometer diagnostic [5] provides a view on a poloidal cross-section of the machine, with two cameras – a vertical camera and a horizontal camera – as illustrated in figure 1(a). Each camera has 24 lines of sight, with 16 lines of sight covering the entire plasma region, and 8 lines of sight providing a more fine-grained resolution of the divertor region. Each line of sight is associated with a separate bolometer, which is essentially a sensor with a thin metal foil (about 10 μm) and a temperature-sensitive resistor attached to it. As the metal foil absorbs radiation, its temperature changes and there is a proportional change in the resistance, which can be measured to provide linear response to radiated power in the range of ultraviolet to soft X-rays [9].

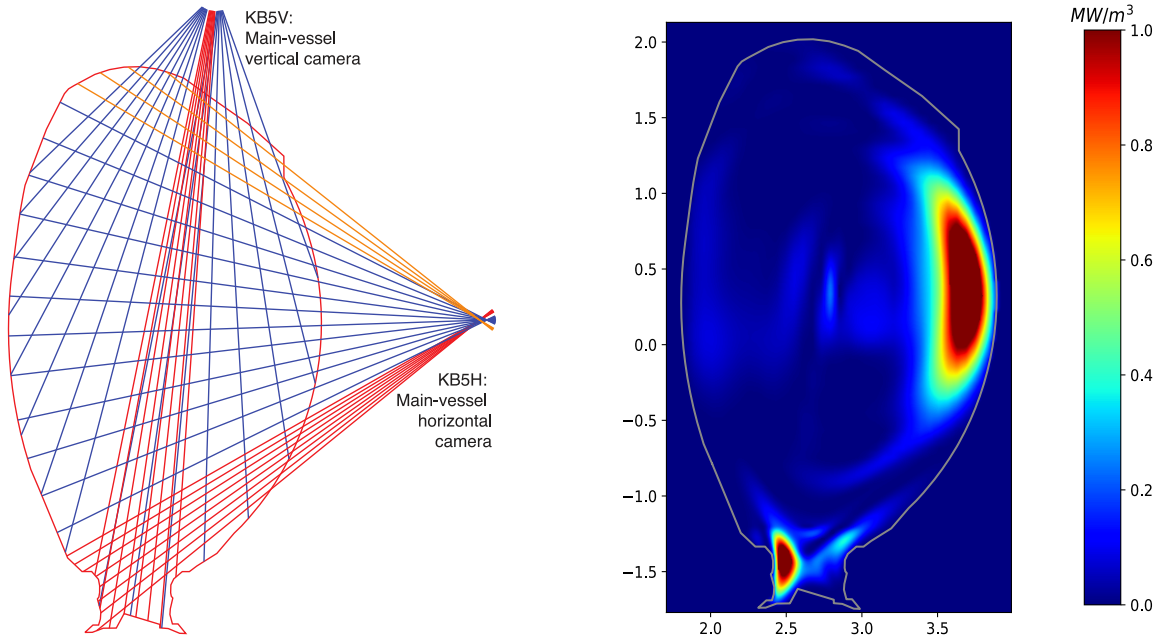

(a) (b)

Figure 1: Geometry for the lines of sight of the bolometer diagnostic at JET (a) together with a sample reconstruction of the plasma radiation profile (b).

Technically, each bolometer measures the line-integrated radiation along a particular line of sight. Therefore, if $\boldsymbol{f}$ are the bolometer measurements (in the form of a 1D column vector), and $\boldsymbol{g}$ is the plasma radiation profile (in the form of a 2D image flattened into a 1D column vector), then $\boldsymbol{f} = \mathbf{K}\boldsymbol{g}$ describes how the bolometer measurements $\boldsymbol{f}$ are obtained from the plasma profile $\boldsymbol{g}$ by a geometry matrix $\mathbf{K}$, which defines how each point (or pixel) in $\boldsymbol{g}$ contributes to each measurement in $\boldsymbol{f}$.

The equation $\boldsymbol{f} = \mathbf{K}\boldsymbol{g}$ describes the forward problem of finding the measurements $\boldsymbol{f}$ from the plasma profile $\boldsymbol{g}$ but, in general, we are interested in the inverse problem of finding $\boldsymbol{g}$ from $\boldsymbol{f}$. Conceptually, one could think of a solution in the form $\boldsymbol{g} = \mathbf{K}^{-1}\boldsymbol{f}$ but $\mathbf{K}$ is not invertible because the problem is severely ill-posed [10]. Specifically, the standard resolution of a plasma radiation profile at JET is 195×115 pixels, i.e. 22540 unknows for only 48 measurements, so $\mathbf{K}$ is a 48×22540 matrix.

A more feasible approach is to find a solution that minimizes some sort of error, such as $\|\boldsymbol{f} - \mathbf{K}\boldsymbol{g}\|^2 \leq \|\varepsilon\|^2$. To ensure some desirable properties of the solution $\boldsymbol{g}$, such as the smoothness of the plasma radiation profile, one can add a regularisation term in the form $\|\boldsymbol{f} - \mathbf{K}\boldsymbol{g}\|^2 + \alpha O(\boldsymbol{g}) \leq \|\varepsilon\|^2$ where α is the regularisation weight and $O(\cdot)$ is some function of $\boldsymbol{g}$ that decreases when $\boldsymbol{g}$ approaches the desired property. In addition, one can enforce some constraints, such as $g_i \geq 0$ for every element of $\boldsymbol{g}$ in order to rule out unphysical solutions, such as negative radiation.

The tomographic reconstruction method that is used at JET [6] is an iterative, constrained optimisation procedure that minimizes the error with respect to the measurements, while requiring the solution to be non-negative. This iterative procedure takes a significant amount of time, typically on the order of minutes to produce a single reconstruction. Figure 1(b) shows a sample result, plotted using a dynamic range from zero (clipping any spurious negative values) to 1.0 MW/m^3.

Naturally, there are faster implementations of the same or similar tomographic techniques [11–13]. There are also new methods, including maximum likelihood [14] and Gaussian process tomography [15], but these are more geared towards estimating the uncertainty of the results rather than improving their accuracy or run-time performance. In this respect, a promising approach would be to use deep learning to generate the plasma profile [16], but this requires the use of GPUs, which are not available in the real-time network at JET. Offloading the computation of a deep learning model to the CPU is also not a viable option for real-time.

This means that any attempt at implementing plasma tomography in real-time must be prepared to make significant concessions in accuracy to favour of run-time performance. In the next section, we describe how plasma tomography was reduced to a single matrix multiplication step in order to bring it to real-time.

3 Machine learning for real-time tomography

In the previous section, we have seen that one of the main ideas in plasma tomography is to come up with a hypothetical solution for $\boldsymbol{g}$ such that $\mathbf{K}\boldsymbol{g}$ is close to the actual measurements $\boldsymbol{f}$ according to some error measure. Even if $\mathbf{K}\boldsymbol{g}$ is not close to $\boldsymbol{f}$, the solution can be improved iteratively until the error measure is minimised. This is a basic strategy for solving the inverse problem without actually inverting $\mathbf{K}$; instead, we resort to the repeated application of the forward problem $\boldsymbol{f} = \mathbf{K}\boldsymbol{g}$.

However, this iterative procedure is not compatible with the requirements of real-time monitoring and control because, in real-time, we cannot afford the run-time needed to find a solution iteratively. Instead, it would be ideal, from a purely performance perspective, to run plasma tomography as a forward problem, where the plasma profile would be generated from bolometer data in a single computational step.

Suppose, for a moment, that there is a matrix $\mathbf{M}$ such that $\mathbf{M}\boldsymbol{f} = \boldsymbol{g}$. This means that, by multiplying $\mathbf{M}$ with the bolometer measurements $\boldsymbol{f}$, the plasma profile $\boldsymbol{g}$ can be obtained immediately. In fact, we do not have matrix $\mathbf{M}$, but if we have some training data $(\tilde{\boldsymbol{f}}^{(1)}, \tilde{\boldsymbol{f}}^{(2)}, ..., \tilde{\boldsymbol{f}}^{(n)})$ and $(\tilde{\boldsymbol{g}}^{(1)}, \tilde{\boldsymbol{g}}^{(2)}, ..., \tilde{\boldsymbol{g}}^{(n)})$ – with $\tilde{\boldsymbol{g}}^{(k)}$ being the pre-computed plasma profile for measurements $\tilde{\boldsymbol{f}}^{(k)}$ – then it is possible to train $\mathbf{M}$ by minimizing an error measure, such as the mean absolute error $\frac{1}{n}\sum_k \|\mathbf{M}\tilde{\boldsymbol{f}}^{(k)} - \tilde{\boldsymbol{g}}^{(k)}\|_1$. The solution is not unique but it should be possible to reduce the set of admissible solutions if a large and diverse dataset is available for training.

Once the training process is complete, then matrix $\mathbf{M}$ will be readily available for use in real-time to compute the plasma profile $\boldsymbol{g}$ for any given bolometer measurements $\boldsymbol{f}$. The matrix can also be used offline to compute the plasma profiles across an entire pulse in one sweep, through the matrix multiplication $\boldsymbol{G} = \mathbf{M}\boldsymbol{F}$, where the columns of $\boldsymbol{F}$ contain the bolometer measurements for successive points in time and the columns of $\boldsymbol{G}$ will provide the corresponding plasma profiles.

To train the model $\mathbf{M}$, we collected about ∼12000 sample profiles that have been computed across ∼700 pulses at JET. Typically, these reconstructions were computed at a few, specific points of interest in each pulse. The sample profiles were manually checked for reconstruction quality and, whenever possible, the measurements from known malfunctioning bolometers have been removed.

In other words, we collected each sample profile $\tilde{\boldsymbol{g}}^{(k)}$ together with the corresponding bolometer measurements $\tilde{\boldsymbol{f}}^{(k)}$ for the same pulse, at the same point in time. The model $\mathbf{M}$ was trained by minimizing the mean absolute error between $\mathbf{M}\tilde{\boldsymbol{f}}^{(k)}$ and $\tilde{\boldsymbol{g}}^{(k)}$. In this context, the mean absolute error is more robust to noise than, for example, the mean squared error. Since the model is linear, no data normalization was employed.

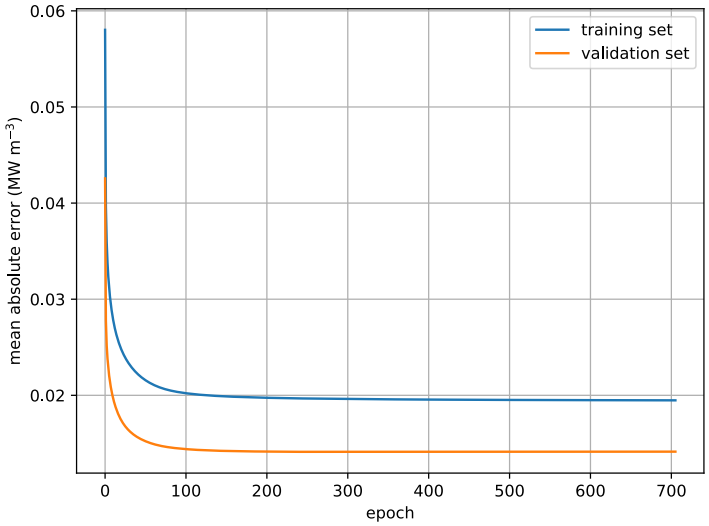

Figure 2: Evolution of the mean absolute error during the training of matrix **M**.

At a resolution of 195×115 pixels, and using a 32-bit floating point representation (4 bytes), the dataset is about $12000 \times 195 \times 115 \times 4 \simeq 1$ GB, which can easily fit the memory of a single GPU. Therefore, we loaded the training data onto a single GPU and trained the model until convergence.

To avoid overfitting, we kept 10% of data apart for validation, and we stopped training when the error ceased to improve on the validation set. The data split was done by pulse (1 in every 10 pulses was set apart for validation) in order to avoid having profiles from the same pulse across the training and validation sets. Figure 2 shows the evolution of the error in the training set and in the validation set during training, where we applied the Adam optimizer [17] starting from a zero-initialization of **M**. The model was trained for about 700 epochs, but the minimum error on the validation set did not improve beyond epoch 350. The training process was implemented with TensorFlow and took about an hour to run on an Nvidia P100 GPU.

As shown in figure 2, the error in the training set converges to about 0.02 MW/m^3 (one third of its initial value), which is relatively low when compared to the dynamic range of figure 1(b). The error in the validation set is even lower, which comes as a result of the data split, meaning that the profiles in the validation set happen to be somewhat easier to reconstruct than those in the (larger) training set.

Overall, the error is still higher (about two times higher) than the error achieved with a deep neural network [16], but this penalty in accuracy is compensated by the computational speedup and the possibility of using the model in real-time. As we will show in the next section, the accuracy is sufficient to recognize a wide range of plasma phenomena in the reconstructions produced by the model.

4 Observation of radiative phenomena in real-time

Figure 3 illustrates the type of reconstructions that can be obtained in real-time by multiplying **M** with the bolometer data at each point in time during a pulse. In this particular 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 $t = 53.7$ s. The two radiation blobs are possibly related by impurity transport from edge to core, since at around $t = 52.0$ s the outboard radiation decreases, followed by a sudden increase in core radiation.

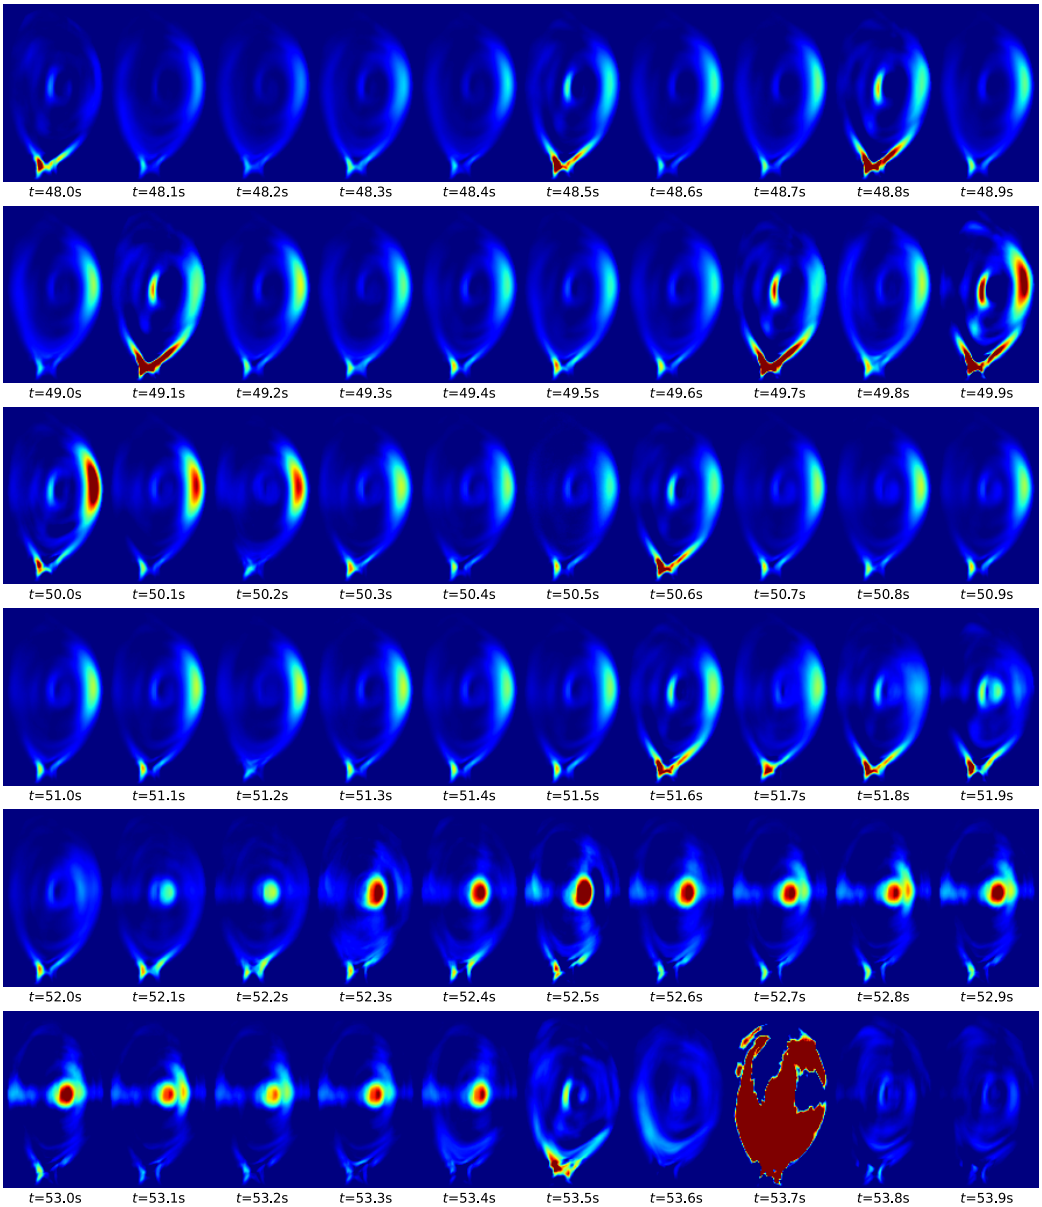

Figure 3: Plasma radiation profiles for pulse 92213 from $t = 48.0$ s to $t = 53.9$ s with a time step of 0.1 s and a dynamic range of 1.0 MW/m^3.

Other radiative phenomena can also be observed in great detail. Figure 4 shows the example of a sawtooth crash [18], where, at around t = 51.965 s, there is an event on a fast time scale that flushes away the impurities from the core. Core radiation then slowly recovers until it reaches, and eventually exceeds, its previous levels. In this case, the sawtooth crash appears to have the beneficial effect of countering impurity accumulation and reducing core radiation, at least for a brief moment. Unfortunately, this does not seem to change the course of events in this pulse.

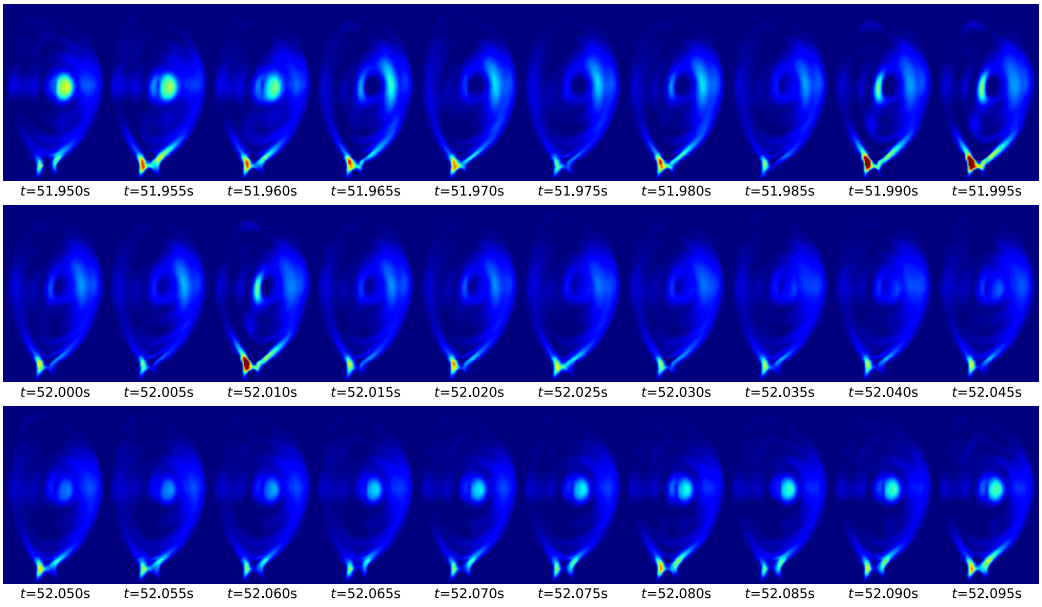

Figure 4: Plasma radiation profiles for pulse 92213 from $t = 51.950$ s to $t = 52.095$ s with a time step of 5 ms, illustrating the occurrence of a sawtooth crash.

Another phenomena that can be observed with real-time tomography are ELMs (edge-localised modes) [19]. These are standard behavior in H-mode (high-confinement mode) [20], and they appear as short bursts of radiation around the divertor region in the plasma radiation profile. In general, there are many occurrences of ELMs during H-mode, and figure 5 illustrates only one of such occurrences. The burst lasts for about 5 ms and, from what can be seen in the plasma profile, it appears to interact with the radiation blob on the outboard edge. On the other hand, the hint of core radiation is an artifact of the reconstruction, due to the fact that the same line of sight passes through the core and divertor.

When observing the occurrence of ELMs across entire pulses, it becomes apparent that, as long as the plasma is ELMing regularly, the radiation blob on the outboard edge is not allowed to become too strong. On the other hand, in the JET baseline scenario it has been observed that if ELMs are interrupted or the spacing between ELMs increases, the outboard radiation intensifies dramatically, and contributes to a global rise in the total radiated power.

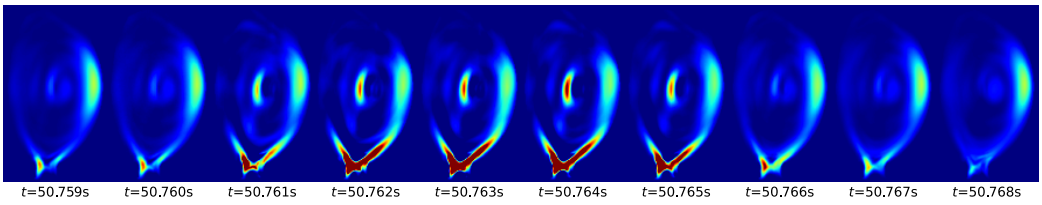

Figure 5: Plasma radiation profiles for pulse 92213 from $t = 50.759$ s to $t = 50.768$ s with a time step of 1 ms, illustrating the occurrence of an ELM.

For this reason, ELMs (in particular regularly paced ELMs, and moderate ELMs so as not to damage the plasma-facing materials) are regarded as having a beneficial effect in removing impurities and keeping the total radiated power under control.

A third and last phenomenon that we will illustrate here is the occurrence of MARFE (multifaceted asymmetric radiation from the edge) [21]. This is a phenomenon that occurs much later in the pulse, during the ramp-down phase when the input power is being turned off and plasma radiation is still relatively high. In these conditions, the plasma starts to exhibit bursts of radiation on the high-field side that appear to have their origin near the X-point and rotate clockwise around the edge.

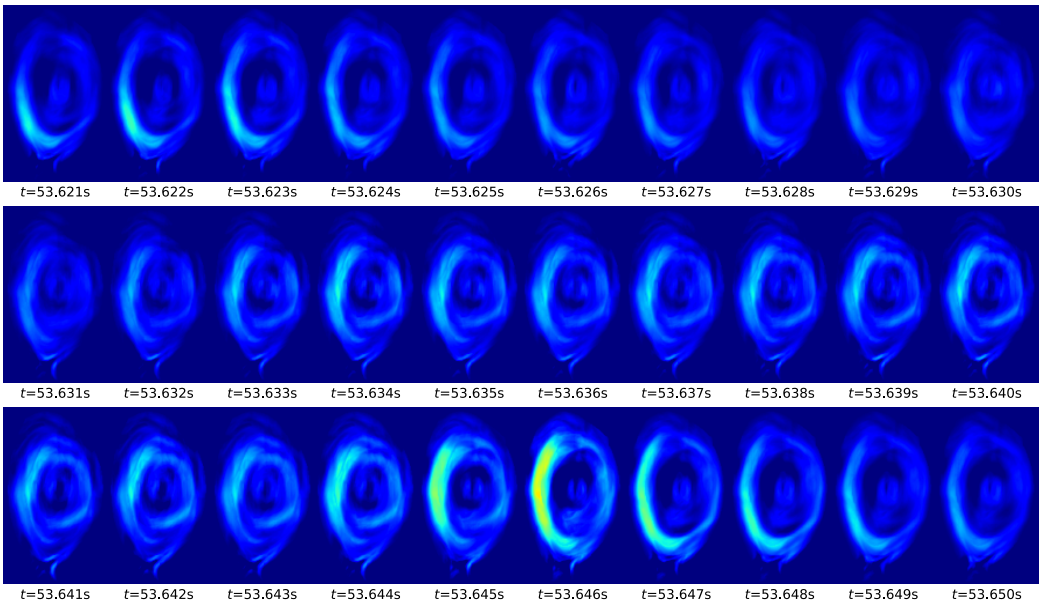

Figure 6: Plasma radiation profiles for pulse 92213 from $t = 53.621$ s to $t = 53.650$ s with a time step of 1 ms, illustrating the occurrence of MARFE.

Figure 6 illustrates this phenomenon, which tends to repeat itself a few times, with growing intensity, before the end of the pulse. At JET, MARFEs are usually observed at a point when the total radiated power exceeds the input power. If this excess radiation is continued, the pulse invariably ends with a disruption. Since MARFEs occur a few hundred milliseconds before the disruption, they pro-

vide an indication of what is about to happen, with sufficient time for taking action, such as reapplying heating.

This realization has led to a renewed interest in MARFE as disruption precursors, and recently there has been an effort to develop MARFE detectors for other machines such as ASDEX Upgrade and TCV [22]. Even though MARFE is a well-known phenomena that has been identified long ago, the relevance of such phenomena for anticipating disruptions has been realized only more recently, and real-time tomography can provide a possible way to detect it. At JET, and as described in the next section, we focus on core radiation because it provides an indication that something is going wrong much earlier in the pulse. The development of strong core radiation associated with impurity accumulation has become an important issue in scenario development.

5 Core radiation in JET baseline pulses

To analyze core radiation in the JET baseline scenario, we focus on three regions of interest in the plasma radiation profile: (1) the outboard edge region, which is always a focus of intense radiation before core radiation develops, when it occurs; (2) the core region itself, where impurities tend to accumulate and create a hollow temperature profile; and (3) the divertor region, which, through the effect of ELMs, can play an active role in keeping the outboard radiation under control.

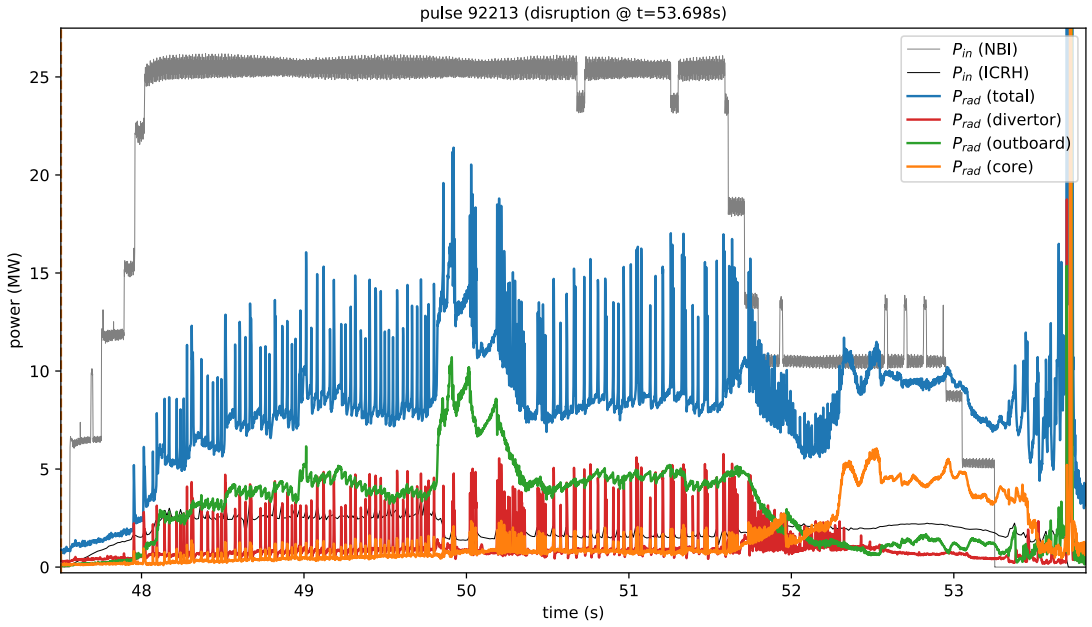

Figure 7: Input power (NBI and ICRH), total radiated power, and radiated power in the three regions of interest (core, outboard and divertor) for pulse 92213.

To monitor, in real-time, the radiated power in each of these three regions, we apply a specific mask over the plasma profile [7]. The signals that are produced by those masks are plotted in figure 7, together with the total radiated power, and also the input power, from both neutral beam injection (NBI) and ion cyclotron resonance heating (ICRH), that was applied in this pulse. We can see that the spikes in the total radiated power are due to the activity in the divertor region (i.e. ELMs). When such activity is interrupted or becomes more spaced, there is a dramatic increase in outboard radiation, which occurs in this plot (and in figure 3) at around $t = 50$ s.

Later on, at around $t = 52$ s, after a major drop in NBI power, there is a transition of radiated power from edge to core, which is not very apparent in the total radiated power (blue), but is evident when looking at the ouboard (green) and core (orange) regions. In fact, what follows after $t = 52$ s is a period of about 1 s or more of intense core radiation. When the input power is brought down after $t = 53$ s, core radiation is still quite high, and the pulse starts MARFEing, as can be seen by the oscillations in total radiated power, just before the disruption.

To study the deleterious effect of core radiation in the baseline scenario, we gathered all baseline pulses from recent experimental campaigns at JET, since 2015. This yielded a total of 486 pulses, which we separated into disruptive and non-disruptive ones. For each of these classes, we computed the maximum core radiation as follows:

- For disruptive pulses, we computed the maximum core radiation achieved before the disruption time.
- For non-disruptive pulses, we computed the the maximum core radiation achieved across the entire pulse.

We then plotted the maximum core radiation for each pulse as shown in figure 8. Overall, there are 40% of disruptive pulses and 60% of non-disruptive ones. This corresponds the disruptivity rate observed in the JET baseline scenario. As expected, we can see that non-disruptive pulses are concentrated towards low values of core radiation, while disruptive pulses are spread over a wider range of, in general, higher values. In other words, disruptions tend to be associated with higher core radiation, but that is not always the case, as there are disruptions where this factor does not play a role. In any case, there will be a threshold on core radiation that separates disruptive and non-disruptive pulses in an optimal way, and we tried to find this threshold.

For this purpose, we employed several metrics for binary classification, namely accuracy, F1 score, and critical success index, as defined below. All of these metrics are based on the number of

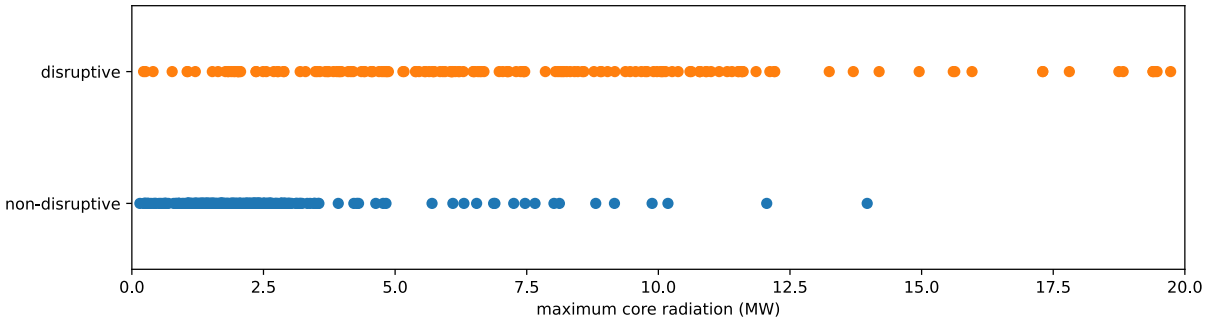

Figure 8: Maximum core radiation achieved for each baseline pulse in two different classes (disruptive and non-disruptive).

true positives (TP), the number of true negatives (TN), the number of false positives (FP) and the number of false negatives (FN).

In our context:

- A true positive occurs when core radiation exceeds the threshold before the disruption time in a disruptive pulse.
- A true negative occurs when core radiation does not exceed the threshold at any point in time during a non-disruptive pulse.
- A false positive occurs when core radiation exceeds the threshold at any point in time during a non-disruptive pulse. This is also referred to as a false alarm.
- A false negative occurs when core radiation does not exceed the threshold before the disruption time in a disruptive pulse. If core radiation exceeds the threshold on or after the disruption time, it is also a false negative. This is also referred to as a missed alarm.

Accuracy is calculated as ACC = (TP+TN)/(TP+TN+FP+FN), the F1 score is calculated as F1 = 2×TP/(2×TP+FP+FN), and the critical success index is calculated as CSI = TP/(TP+FP+FN). By scanning possible values for the threshold in a sensible range (from 1 to 6 MW), we obtained the results in figure 9.

Here, there are two critical points:

- One is the point at which the metrics are maximised, and this point is the same for all three metrics; it would yield a threshold of about 3.56 MW.
- The other is the point at which the false positive rate (% of false alarms) is equal to the false negative rate (% of missed alarms). This would yield a threshold of 3.0 MW, which is lower,

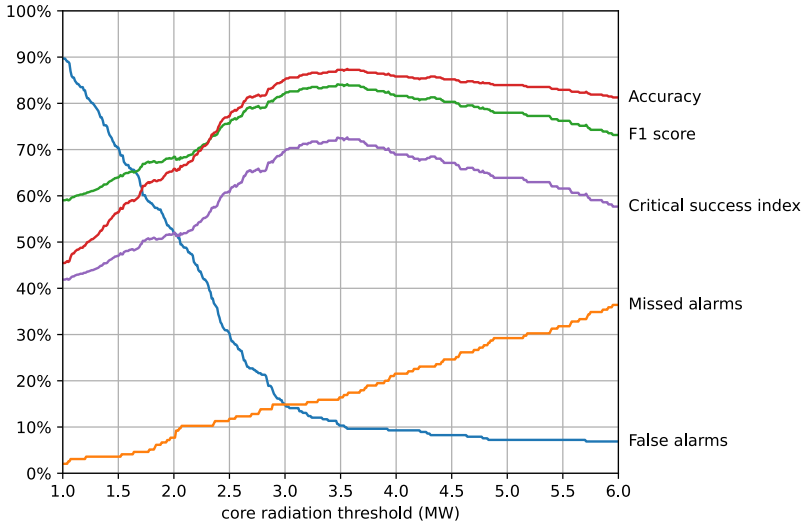

Figure 9: Binary classification metrics as a function of threshold on core radiation.

but it makes sense to be more conservative, since missed alarms involve the risk of damaging the machine, whereas false alarms only bring the inconvenience of ending a healthy experiment prematurely. Therefore, we opted for the more conservative choice of 3.0 MW.

This absolute threshold of 3.0 MW was chosen for the particular operating conditions of the JET baseline scenario. The threshold divides the baseline disruptive pulses into 85% true positives and 15% false negatives (missed alarms). This means that there are at least 15% of baseline pulses which disrupt due to causes other than excessive core radiation. On the other hand, the same threshold divides the non-disruptive pulses into 85% of true negatives and 15% of false positives (false alarms). The fact that these percentages (TP-vs.-FN and TN-vs.-FP) are the same is a result of our choice of threshold to make the false alarm rate equal to the missed alarm rate.

Although our goal is not to compete with other real-time disruption predictors [23–26], the results can still be compared with some of the results obtained at JET. For example, APODIS [24] is a disruption predictor that uses multiple signals, including the total radiated power (but not the plasma radiation profile as we do here). Trained on a different dataset, APODIS is reported to have a similar true positive rate (85%), but a much lower false alarm rate (2.5%) [27]. On the other hand, our threshold on core radiation is able to provide a warning much sooner, on average 1.4 s before the disruption, compared the average warning time for APODIS (350 ms). Such warning time, in excess of 1 s, may even be considered premature in some cases.

Once the alarm is triggered, a possible way to respond is to apply auxiliary heating, especially ICRH (ion cyclotron resonance heating), which has been shown to be particularly effective for tungsten control in the JET baseline scenario [28].

6 Conclusion

In this work, we described an approach to real-time tomography that is based on the idea of reducing the reconstruction process to a single matrix multiplication, where the matrix has been trained by machine learning on a large collection of sample profiles. The computational speedup that is achieved by using such matrix to reconstruct the plasma radiation profile opens up several new possibilities.

One possibility is to use this approach to investigate the radiative behavior of the plasma in great detail. In this paper we have illustrated some of the most common radiative phenomena that have been observed in JET baseline pulses. The same approach can be used to investigate plasma behavior in other operational scenarios. The reconstruction of the plasma radiation profile is not perfect and may suffer from the same kind of artifacts that are present in the training data, but it is sufficiently accurate to serve as a replacement of the original method for most practical purposes.

Another possibility is to use this approach to continuously monitor the plasma radiation profile in real-time. In particular, one can focus on special regions of interest where plasma radiation should be closely monitored, such as core, outboard and divertor. The detection of excessive radiation in some of these regions may justify triggering an alarm in the real-time control system of the experiment. Whether such alarm should be acted upon immediately, or whether one should wait for further symptoms before taking action, is another issue. In any case, real-time tomography is fulfilling its role of providing an additional instrument to monitor plasma behavior.

As ongoing work, other alarms are being prepared based on the same regions of interest. Besides core radiation, there is a threshold on outboard radiation, and on core-to-outboard power ratio. Additionally, two alarms based on core-to-input power and outboard-to-input power are also available. Finally, as future work, we will be looking into the options of defining additional regions of interest, such as an inboard region to monitor radiative phenomena on the high-field side, for example.

Acknowledgments

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. IPFN (Instituto de Plasmas e Fusão Nuclear) received financial support from FCT (Fundação para a Ciência e a Tecnologia) through projects UIDB/50010/2020 and UIDP/50010/2020. The authors are thankful for the granted use of computational resources provided by CCFE/UKAEA at Culham, UK.

References

[1] Garzotti L *et al* 2019 *Nucl. Fusion* **59** 076037

[2] de Vries P *et al* 2011 *Nucl. Fusion* **51** 053018

[3] Matthews G F *et al* 2011 *Phys. Scr.* **T145** 014001

[4] Joffrin E *et al* 2013 *Nucl. Fusion* **54** 013011

[5] Huber A *et al* 2007 *Fusion Eng. Des.* **82** 1327–1334

[6] Ingesson L C *et al* 1998 *Nucl. Fusion* **38** 1675–1694

[7] Ferreira D R *et al* 2021 *Fusion Eng. Des.* **164** 112179

[8] Stuart C I *et al* 2021 *Fusion Eng. Des.* **168** 112412

[9] Mast K F *et al* 1991 *Rev. Sci. Instrum.* **62** 744–750

[10] Mlynar J *et al* 2019 *J. Fusion Energy* **38** 458–466

[11] Mlynar J *et al* 2010 *Fusion Sci. Technol.* **58** 733–741

[12] Mazon D *et al* 2013 *Nucl. Instrum. Methods Phys. Res. A* **720** 78–82

[13] Svoboda J *et al* 2019 *J. Instrum.* **14** C11001

[14] Craciunescu T *et al* 2018 *Rev. Sci. Instrum.* **89** 053504

[15] Wang T *et al* 2019 *J. Fusion Energy* **38** 445–457

[16] Ferreira D R *et al* 2018 *Fusion Sci. Technol.* **74** 47–56

[17] Kingma D P and Ba J 2014 *arXiv* 1412.6980

[18] von Goeler *et al* 1974 *Phys. Rev. Lett.* **33**(20) 1201–1203

[19] Keilhacker M *et al* 1984 *Plasma Phys. Control. Fusion* **26** 49–63

[20] Wagner F *et al* 1982 *Phys. Rev. Lett.* **49**(19) 1408–1412

[21] Lipschultz B *et al* 1984 *Nucl. Fusion* **24** 977–988

[22] Maraschek M *et al* 2017 *Plasma Phys. Control. Fusion* **60** 014047

[23] Cannas B *et al* 2010 *Nucl. Fusion* **50** 075004

[24] Vega J *et al* 2013 *Fusion Eng. Des.* **88** 1228–1231

[25] Rea C *et al* 2019 *Nucl. Fusion* **59** 096016

[26] Bao N N *et al* 2020 *IEEE Trans. Plasma Sci.* **48** 715–720

[27] Moreno R *et al* 2016 *Fusion Sci. Technol.* **69** 485–494

[28] Goniche M *et al* 2017 *Plasma Phys. Control. Fusion* **59** 055001