

High temporal resolution of pedestal dynamics via machine learning on density diagnostics

Diogo R. Ferreira¹, Andreas Gillgren², Andrei Ludvig-Osipov², Pär Strand² and JET Contributors*

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

² Department of Space, Earth and Environment, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

* See the author list of Mailloux J *et al* 2022 *Nuclear Fusion* **62** 042026

E-mail: diogo.ferreira@tecnico.ulisboa.pt

Abstract. At the Joint European Torus (JET), the reference diagnostic to measure electron density is Thomson scattering. However, this diagnostic has a low sampling rate, which makes it impractical to study the temporal dynamics of fast processes, such as edge localized modes (ELMs). In this work, we use machine learning to predict the density profile based on data from another diagnostic, namely reflectometry. By learning to transform reflectometry data into Thomson scattering profiles, the model is able to generate the density profile at a much higher sampling rate than Thomson scattering, and more accurately than reflectometry alone. This enables the study of pedestal dynamics, by analyzing the time evolution of the pedestal height, width, position and gradient. We also discuss the accuracy of the model when applied on experimental campaigns that are different from the one it was trained on.

Keywords: Plasma Diagnostics, Thomson Scattering, Reflectometry, Machine Learning

1. Introduction

In fusion devices, it is often the case that there are multiple diagnostics to measure the same physical quantity. These diagnostics can be based on different physical principles, and each diagnostic may have specific advantages over another, such as having a higher time resolution or a higher spatial accuracy.

This is precisely the case with plasma density, which is a key quantity to be measured in fusion experiments, and for which there are several alternative methods. At the Joint European Torus (JET), the reference diagnostic for measuring electron density is Thomson scattering [1]. This diagnostic, which is referred to as High Resolution Thomson Scattering (HRTS), is based on laser scattering, and it measures plasma density at fixed radial positions from the intensity of scattered light. In this context,

high resolution refers to the spatial accuracy of those positions, which is on the order of 1 cm [2], on a machine whose minor radius exceeds 1 m.

Despite its accuracy and reliability, perhaps the one and only drawback of the Thomson scattering diagnostic is its low sampling rate, which is on the order of 20 Hz. This makes it impractical to analyze the time evolution of the density profile in detail. For example, to study pedestal dynamics [3], it would be important to track events such as edge localized modes (ELMs), which occur on a faster time scale, and would require a sampling rate of at least 1 kHz to capture the transient behavior of ELM crashes.

For that purpose, there are other diagnostics that can measure the density profile at a higher sampling rate. One of such diagnostics is reflectometry [4], which provides a high temporal resolution, and can reconstruct the density profile at a sampling rate that is typically in the range of 1–10 kHz. This diagnostic, whose internal name at JET is KG10, is based on probing the plasma with electromagnetic waves over a wide range of frequencies (44–150 GHz). Each frequency corresponds to a known density value, where the wave will be reflected. The position of that density value is determined from the travel time of the reflected wave. In other words, reflectometry probes the plasma to find the positions of a set of known density values.

In reflectometry, the process of frequency sweeping can be carried out as fast as 15 μ s, but the process of reconstructing the density profile is typically performed at a rate that does not exceed 10 kHz. This is because calculating the group delay of the reflected wave involves Fourier analysis of the reflectometry signals [5]. In addition, determining the positions from the group delay is typically achieved through a sequential algorithm, such as [6], where errors in the calculated positions, if they occur, tend to accumulate as the density profile is being reconstructed inwards.

In summary, Thomson scattering measures density values at precise positions, but at a low sampling rate and with some uncertainty in the measured values. On the other hand, reflectometry measures precise density values at a high temporal resolution, but with some uncertainty in the positions of those measurements. Ideally, it should be possible to combine the characteristics of these diagnostics in order to obtain density profiles with accurate measurements, accurate positions, and at a high sampling rate.

In this work, we use machine learning to achieve this purpose. By training a machine learning model to predict Thomson scattering profiles from reflectometry data, we can generate density profiles with the spatial accuracy of HRTS and at the sampling rate of KG10. The model has enough learning capacity to overcome the problems observed in some reflectometry profiles, where the radial positions of the measured density values deviate significantly from the ones provided by Thomson scattering. This allows the trained model to produce density profiles that are free from reflectometry artifacts and are close to what HRTS would provide.

The improved sampling rate that can be achieved by predicting the HRTS profile from reflectometry data enables the study of pedestal dynamics and ELM behavior. While it would be possible to consider predicting the same profile from other signals, diagnostics, or control parameters of the machine, a challenge with the latter is that even

in stationary phases where those parameters are approximately constant, the pedestal generally shows dynamic or cyclic patterns for most high-confinement mode (H-mode) plasma scenarios [7]. Therefore, in this work we resort to the use of measurements that can directly capture the pedestal dynamics, namely reflectometry.

The paper is organized as follows. In Section 2, we explain how the data – used for training, validating and testing our model – was collected from deuterium (D), deuterium-tritium (DT) and tritium (T) campaigns at JET. In Section 3, we present the model architecture, the training process, and prediction examples. The advantages of using the model are illustrated in Section 4, by analyzing the inter-ELM pedestal dynamics of the record-breaking pulse achieved at JET in the last DT campaign [8]. The impact of training the model on one campaign and testing on another is discussed in Section 5. Finally, Section 6 concludes the paper.

2. Training data

Figure 1 shows an example of a time point where data from both diagnostics are available. The Thomson scattering profile has 63 data points, which correspond to the fixed radial positions at which density values are measured. On the other hand, the reflectometry profile has 100 data points, whose positions are determined by frequency sweeping and computing the group delay. In Figure 1, it can be seen that the reflectometry profile extends further inwards in terms of radial position, and is also smoother than the one provided by Thomson scattering.

In this example, both diagnostics agree on the density profile, and particularly on the pedestal shape and position. However, this is not always the case. In a




Figure 1. Example of an instance where Thomson scattering (HRTS) and reflectometry (KG10) agree on the density profile.

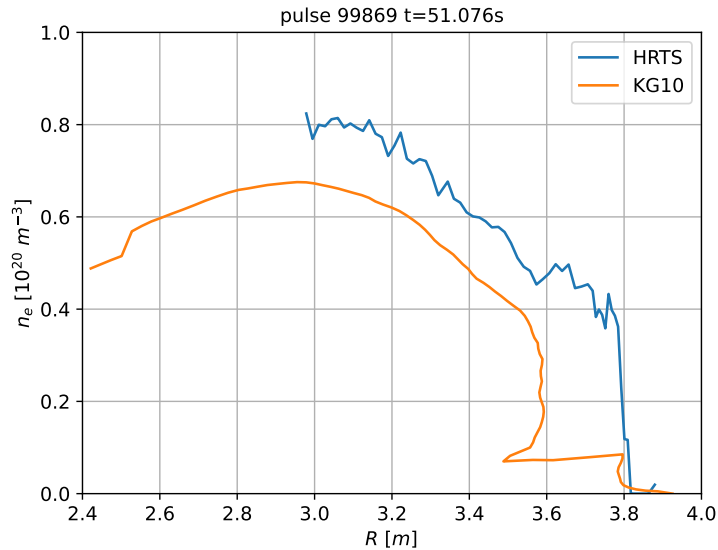


Figure 2. Example of an instance where the reflectometry profile (KG10) deviates from the one provided by Thomson scattering (HRTS).

small percentage of cases (estimated to be around 5%), reconstruction errors prevent reflectometry from arriving at an accurate density profile. Figure 2 presents such a case, where an error in determining the position along the steep region, which is referred to as the pedestal, results in all subsequent inward positions going off track.

In these rare cases, of which Figure 2 is an extreme example, the reflectometry profile deviates significantly from the one obtained by Thomson scattering. However, if we use machine learning to learn the mapping of KG10 to HRTS data, it might be possible to transform one into the other. For this purpose, it is necessary to gather training data in the form of examples where both HRTS and KG10 data are available at the same time point. It is possible to find such training examples in the large amounts of data from JET experiments. For this work, we selected the two main experimental scenarios at JET – namely, the *baseline* and the *hybrid* scenarios [9] – which are being developed to achieve high fusion performance in stationary conditions.

An essential requirement for building such dataset is that the data should come from experimental sessions where both HRTS and KG10 were fully operational. This was the case from December 2020 onwards. Therefore, we identified the following campaigns as being the most interesting for collecting data:

- the C40 campaign (March–July 2021), which was a tritium campaign;
- the C41 campaign (August–December 2021), which was the second deuterium-tritium campaign at JET, after the first DT campaign back in 1997;
- the C42 campaign (July–August 2022), which was a deuterium-based cleaning campaign.

From these campaigns, we identified 52 experimental sessions belonging to the

baseline and hybrid scenarios, with a total of about 200 usable pulses. Since each pulse lasts for about 30 seconds, and the sampling rate of HRTS is 20 Hz, we expected to collect a number of samples on the order of $200 \times 30 \times 20 \simeq 1.2 \times 10^5$. However, in practice we got 62140 examples of HRTS–KG10 paired profiles. There are several reasons for this: the first is that, for some pulses, the data from either one or both of those diagnostics were partially missing; the second is that, while HRTS is operational during the entire pulse, KG10 has a shorter operating window due to limited data buffer capacity in the data storage system at JET.

The dataset was split into 90% for training and 10% for testing, in order to have a separate test set on which to evaluate the model, as we will see in Section 5. The data was split by pulse to avoid having examples from the same experiment in the training set and in the test set. To train the model, the training set was split, again by pulse, into 90% for training and 10% for validation, in a 10-fold cross-validation scheme. In addition, we held out a specific pulse (99869) from training, since we use this pulse as an example to study pedestal dynamics in this work, as we will see in Section 4.

3. Machine learning model

To transform KG10 profiles into HRTS profiles, we devised a neural network that receives a KG10 profile as input and produces an HRTS profile as output. Each profile is a set of density values measured at certain positions. For HRTS, the positions are known and fixed; therefore, it is unnecessary to reproduce them as an output of the model. However, for KG10, the positions vary from profile to profile; hence, these positions are provided as an additional input to the model. This additional input is especially useful in cases where, as in Figure 2, the profile suffers from errors in assigning the correct positions to the measured density values.

Figure 3 illustrates the layout of the neural network, where circles represent nodes, numbers above the circles represent the number of nodes in each layer, and lines represent connections between nodes. The model is based on fully-connected (dense) layers, where each intermediate layer has 512 units and includes a rectified linear unit (ReLU) [10] as activation function. There are two input branches, and the first couple of dense layers are meant to preprocess each input before combining them in a linear addition layer. A second stage of dense layers produces the actual prediction. The total number of parameters, in terms of weights and biases, is on the order of 1 million (1 186 367 parameters). This relatively large number of parameters is motivated by the need to transform KG10 into HRTS profiles as accurately as possible, while consistently overcoming the occasional issues with KG10 data.

The model was trained by stochastic gradient descent, using the Adam optimizer [11] with a learning rate of 10^{-4} and a batch size of 256 samples. This choice of batch size resulted in 200 updates of the model parameters for each run (epoch) over the training set, hence the use of a relatively low learning rate to avoid large jumps in the values of training/validation loss from epoch to epoch. To prevent overfitting, we




Figure 3. Layout of the neural network to transform reflectometry profiles into Thomson scattering profiles. The network has several fully connected layers with 512 nodes using ReLU activation. The two input branches are merged by linear addition.




Figure 4. Evolution of loss and validation loss during training. The loss is the mean absolute error of the predicted density values with respect to the ground truth. A minimum value of the validation loss is achieved during training, after which the model begins to overfit.

monitored the evolution of the learning curves in order to find the epoch at which the validation loss stops improving. As illustrated in Figure 4, the validation loss reaches a minimum before epoch 600, after which it begins to increase, while the training loss keeps decreasing. Since the model is not learning anything meaningful past that point (at least not anything that can be useful on the validation set), we stop training and use the model weights that have been obtained at that point.

Once trained, the model can be used to predict the HRTS profile from KG10 data, at time points where Thomson scattering measurements are not available. In particular, the HRTS profile can be predicted at the sampling rate of the KG10 data that is provided as input. To obtain robust predictions and also a basic estimate of model uncertainty, we trained 10 instances of the model via 10-fold cross-validation. Each model instance takes on the order of minutes to train on a single GPU.

Figures 5 and 6 show the model predictions for the same pulse and time points as in Figures 1 and 2, respectively. The red line indicates the average prediction of the 10 model instances, while a red band shows the standard deviation across their predictions. In the case of Figure 5, where KG10 agrees with HRTS, the band is hardly noticeable, since the predictions of all model instances are almost identical. In Figure 6, however, the uncertainty band is slightly larger, but still all predictions fall on top of the HRTS profile. In both cases, the model is capable of generating an output that resembles the HRTS profile, regardless of how far the KG10 profile is.

The analysis of the test pulse 99869, which was held out from the training data, is especially interesting, since this was the record-breaking pulse in terms of energy production achieved in the DTE2 campaign at JET (C41). The experimental conditions were such that the heating and confinement were sustained for long enough, with a good




Figure 5. Prediction on an example where Thomson scattering (HRTS) and reflectometry (KG10) agree.




Figure 6. Prediction on an example where the reflectometry profile (KG10) deviates from the one provided by Thomson scattering (HRTS).

neutron yield, to establish the new record. Specifically, the main engineering parameters of the pulse 99869 were: $B_T = 3.4$ T, $I_P = 2.3$ MA, $P_{NBI} = 27$ MW, $P_{ICRH} = 4.2$ MW, with a fueling rate, in the time frames considered here, of $\Gamma = 1.3 \cdot 10^{22}$ e/s. Among other features, ELMs were occurring at a regular pace, which makes this pulse an interesting candidate to study pedestal dynamics.

4. Pedestal dynamics

In the density profiles predicted by the model over the course of a pulse, one of the features that immediately stands out is the occurrence of ELM crashes. These appear as a sudden decrease in the pedestal height, followed by gradual recovery, until the next ELM crash occurs, and the process repeats itself. Figure 7 shows a plot of the time evolution of the predicted density profile for the test pulse 99869, illustrating how events unfold over an ELM cycle that lasts for about 30 ms.

The ELM crash and its subsequent recovery can be observed by looking at what happens around the radial position $R = 3.8$ m. In addition, the effect of the ELM crash is present not only at the outer edge but, to a lesser extent, towards inner positions as well, as can be seen between $t = 50.03$ s and $t = 50.04$ s in Figure 7, where there is a gap or discontinuity that extends inwards along the density profile.

In practice, the pedestal structure is typically analyzed by means of a curve fit of the profile edge using a modified hyperbolic tangent function (*mtanh*), which is defined as [12]:

$$\text{mtanh}(R) = \frac{h - h_0}{2} \left(\frac{(1 + sx)e^x - e^{-x}}{e^x + e^{-x}} + 1 \right) + h_0, \text{ with } x = \frac{p - R}{w/2}.$$




Figure 7. Time evolution of the predicted density profile across an ELM cycle that begins at around $t = 50.0$ s.

The function is parameterized by pedestal height (h), pedestal position at the largest gradient (p), and pedestal width (w). The pedestal gradient can be approximated by h/w . Besides the pedestal, the fitting is also parameterized by a core slope s and offset (h_0), but in this work we limit the scope to studying h , p and w . Additionally, we consider the position of the top of the pedestal, which can be approximated as $p_{ped,top} = p - w/2$ since p represents the position of the largest gradient, which is near the middle of the steep region. Figure 8 illustrates the main parameters of this hyperbolic tangent fit.

To study the time evolution of these parameters for pulse 99869 – specifically, the pedestal height, position, width and gradient – we applied a modified hyperbolic tangent fit to the edge of the density profiles predicted by the model at each time point. Specifically, we include the data where $R > 3.62$ m in each fit. Figure 9, on the left-hand side, shows the results obtained over a time window of approximately five ELM cycles. For comparison, Figure 9 also shows, on the right-hand side, the results that would be obtained if such curve fit would be applied on the original KG10 profiles. A moving average of 7 entries is used for each signal, both for predicted HRTS and for KG10. The conclusion is that transforming KG10 profiles into predicted HRTS profiles makes it easier to analyze the time evolution of the pedestal parameters. This provides a further




Figure 8. Pedestal parameters of the modified hyperbolic tangent fit to the density profile.

illustration of the advantages of using the proposed model.

In terms of pedestal dynamics, the plots on the left-hand side of Figure 9 illustrate the following behavior, from top to bottom:

- As expected, the first plot of the predicted HRTS signal shows that the pedestal height (h) increases before an ELM crash occurs, which rapidly lowers the pedestal. For the majority of the buildup, the increase follows a linear trend. The KG10 signal shows a similar pattern and a similar density range, although with a higher vertical variation of the signal, which is associated with noise and artifacts in the reflectometry reconstruction.
- The second plot of the predicted HRTS signal shows that the pedestal position (midpoint of pedestal, p , which is approximately the same as the position of the maximum gradient for density pedestals) shifts outwards during the ELM crash and then shifts inwards and saturates during the buildup. The position of the top of the pedestal $p_{ped,top}$ shows the opposite pattern as it shifts inwards during the ELM crash and then moves outwards before it saturates early in the ELM cycle during the buildup. This highlights the importance of emphasizing the definition of the pedestal position during analysis. The reason why the midpoint of the pedestal is shifted outwards during the ELM crash is because the width increases during the ELM crash, which can be seen in the third plot of Figure 9. The width changes more than twice as much during the ELM crash compared to the position of the pedestal top. This moves the midpoint of the pedestal outwards, which, when analysed alone, might lead to the wrong conclusion that the entire pedestal is moving outwards during the ELM crash. A closer look reveals that the top of the pedestal is moving inwards during the ELM crash, and that the increasing position of the midpoint of the pedestal is the result of an increasing width. The

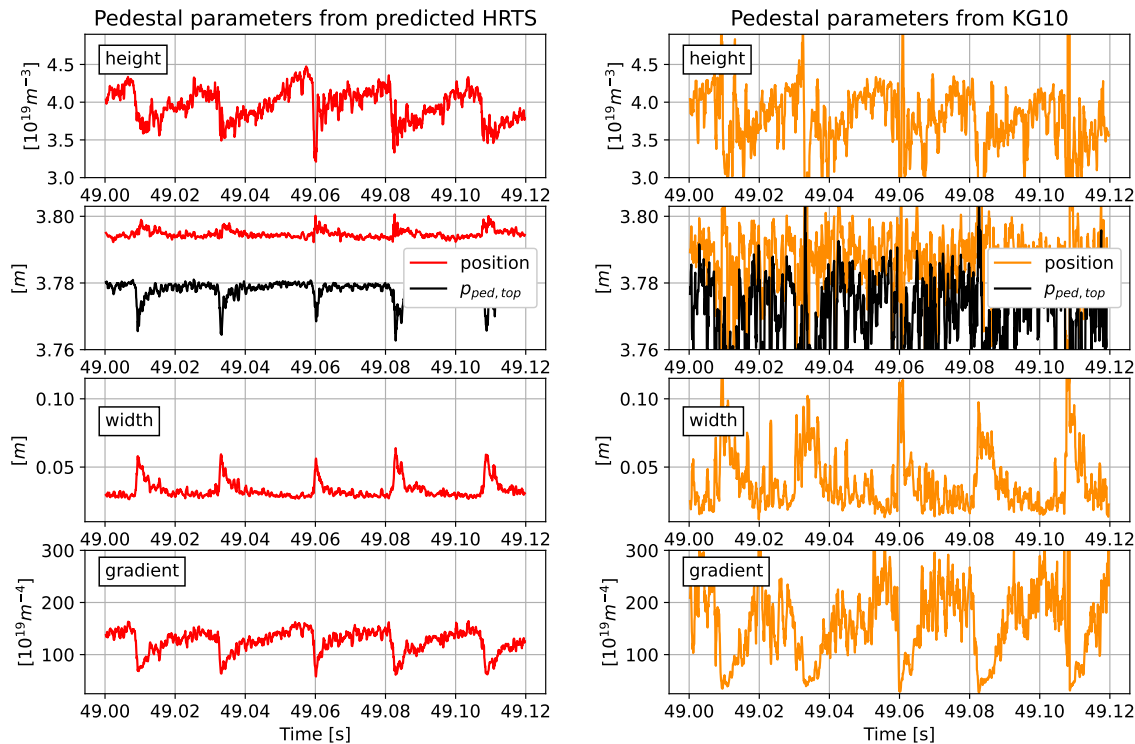


Figure 9. Time evolution of the pedestal parameters for pulse 99869 obtained by curve fitting on the predicted HRTS profiles (*left*) and on the original KG10 profiles (*right*).

KG10 signal is too noisy to allow the analysis of the midpoint of the pedestal as well as the position at the top of the pedestal. This is likely due to the occasional issues with positional accuracy for KG10, which makes these plots more difficult to interpret, although some similarities in pattern can be observed between predicted HRTS and KG10 for the position of the pedestal.

- The third plot of the predicted HRTS signal shows that the pedestal width (w) narrows as the pedestal height increases, and each ELM crash widens the pedestal. Similarly to the pedestal position, this pattern is less linear compared to the height as the width saturates early in the ELM-cycle. A decreasing width of the pedestal has previously been observed, for instance, in JET at low fueling [13], as well as in DIII-D [14]. The KG10 signal of the width shows a similar pattern, however it shows a higher degree of noise as well as a larger range of the width. Additionally, it is not equally obvious that the width saturates at a certain level early in the ELM-cycle for the KG10 signal. Again, the difference between the predicted HRTS signal and the KG10 signal is likely due to the occasional positional accuracy issues with KG10 since the width is strongly related to spatial information.
- The fourth plot of the predicted HRTS signal shows that the combined pattern of

the height and width (h/w) results in a pedestal gradient that increases during the buildup, before each ELM crash occurs. In the early stage of the ELM-cycle, the increase of the gradient is more rapid since both the height increases and the width decreases. When the width saturates, the gradient increases more slowly since only the change of the pedestal height alters the gradient. The same pattern, where the width saturates and the height increases monotonically, has been observed in DIII-D [14] and in TCV [15], although in the latter it is not fully clear if the width is approximately constant throughout the entire ELM cycle, or if the width saturates faster than what could be captured with the time resolution of the used diagnostics. The dynamic pattern of the changing pedestal gradient in this work suggests that either particle transport or diffusivity is changing in the pedestal region throughout the ELM cycle [16]. The KG10 signal shows a similar pattern but a larger range of the gradient mainly due to the larger range of the pedestal width. Thus, an analysis of the KG10 signal alone might lead to an overestimation of the critical pressure gradient triggering the ELM crash, if combined with temperature data.

In summary, by studying the pedestal parameters derived from the predicted density profile, it is possible to analyze the pedestal structure and dynamics in detail. When many predictions in the time sequence are analyzed together, a few occasional occurrences of poor predictions are less degrading to the overall picture of the ELM cycle. For pulse 99869, we observe the inter-ELM behaviour of the density pedestal for a specific operational regime. In JET-ILW, the inter-ELM pattern varies depending on, for instance, heating and gas fueling, as discussed in [3]. Further discussion about the results in Figure 9 and exploration of other regimes are beyond the scope of this work, since the main purpose here is to illustrate an application of the proposed model. Nevertheless, it should be mentioned that the difference in the results between predicted HRTS and KG10 in Figure 9 are due not only to the difference in quality of the different diagnostics, but also to how suitable $mtanh$ fitting is for each diagnostic.

5. Dependence on experimental campaign

As described earlier in Section 2, we trained the model on data from multiple campaigns. It is natural to ask how the model would perform on a campaign that it has not been trained on. In particular, if the model is trained on D experiments, for example, how would it perform on DT or T experiments? To answer this question, we investigate what happens when the model is trained on one campaign and then tested on other campaigns with a different ion content.

Table 1 provides a summary of the results, where each row indicates the campaign used for training, and the columns indicate the campaigns used for testing. The pulses used for training are the ones that belong to the training dataset, and the pulses used for testing are the ones that have been held out from the training dataset.

A row-wise comparison of the results in Table 1 suggests that when a model is transferred to a different test campaign, the error tends to increase when compared to

Table 1. Mean absolute error and standard deviation (in units of 10^{20} particles/m³) obtained when the model is trained on one campaign and tested on others.

		testing campaign			
		C40 (T) 5 pulses	C41 (DT) 11 pulses	C42 (D) 5 pulses	all (T, DT, D) 21 pulses
training campaign	C40 (T) 29 pulses	0.029 \pm 0.009	0.031 \pm 0.005	0.033 \pm 0.008	0.031 \pm 0.007
	C41 (DT) 106 pulses	0.025 \pm 0.007	0.023 \pm 0.007	0.034 \pm 0.018	0.026 \pm 0.012
	C42 (D) 42 pulses	0.034 \pm 0.012	0.036 \pm 0.011	0.022 \pm 0.005	0.032 \pm 0.012
	all (T, DT, D) 177 pulses	0.024 \pm 0.007	0.022 \pm 0.006	0.024 \pm 0.010	0.023 \pm 0.007

the results obtained on the campaign it was trained on. This is what happens when the model is trained on T, and then tested on DT and D (first row); or when the model is trained on DT, and then tested on T and D (second row); or, finally, when the model is trained on D and then tested on T and DT (third row). Although the standard deviation often makes the results overlap, the increase in the mean absolute error for different campaigns has been observed as a consistent trend.

On the other hand, a column-wise analysis of Table 1 indicates that the model trained on DT performs better on T than the model that was solely trained on T (first column). This is attributed to the fact that there is significantly more training data available to the former than to the latter (106 vs. 29 pulses). Therefore, to optimize the performance on a given campaign, one should use training data from the same campaign or, if such data are scarce, then collect as much training data as possible from campaigns with a similar ion content.

However, it should be noted that ion content is not the only factor at play. Often, there are changes in the experimental conditions between campaigns, such as instrument re-calibrations and improvements to data processing, which could affect the results. One example is the re-calibration of HRTS between C41 and C42: in the C42 campaign, the HRTS positions have been shifted 1 cm inwards, in comparison to the HRTS positions from the C40 and C41 campaigns. This does not change the results significantly, nor does it invalidate the analysis of pedestal dynamics, which is based on curve fitting on top of the density profile, but it may contribute to differences between campaigns that go beyond major factors such as ion content.

In any case, recent results suggest that the pedestal has indeed a different structure in D, DT, and T experiments [17], which could explain, even in the absence of other

factors, why the model will perform differently across these campaigns. This highlights how important it is to collect training data from multiple campaigns, as described earlier in Section 2. As shown at the bottom row in Table 1, a model trained on all three campaigns, as we presented in Section 3, in general has better performance than any of the models trained separately on each campaign.

6. Conclusion

This work demonstrates the benefits of using machine learning to combine the advantages of two plasma diagnostics operating on different physical principles. Here, we used machine learning to transform data from one diagnostic (reflectometry) into the output that is typically produced by another diagnostic (Thomson scattering). This cross-diagnostic machine learning approach can be used to great effect when it is possible to use data from one diagnostic to emulate another, with significant gains in terms of accuracy and/or sampling rate.

From this point of view, the redundancy between different diagnostics is actually a key enabler for gaining more insight into what happens inside the plasma. In our case, it was the combination of the positional accuracy of Thomson scattering with the high sampling rate of reflectometry that enabled studying the time evolution of ELM cycles in detail. Our results show that the model, trained on a combination of reflectometry and Thomson scattering, gives a much clearer picture of the transient pedestal parameters than any of those diagnostics alone would be able to provide.

In future work, we plan to extend the approach to other physical quantities, namely electron temperature, for which there are multiple diagnostics available as well. In this case, it would be interesting to study what happens not only at the edge but also at the core, where impurity accumulation often leads to hollow temperature profiles. For the pedestal, access to high temporal resolution temperature data with HRTS-like accuracy will enable studying the inter-ELM behaviour at JET in more detail across multiple operational regimes.

Acknowledgements

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

The research activities at 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 research activities at Chalmers were supported by the Swedish Research Council – *Vetenskapsrådet* (2021-00182).

The graphics processing unit (GPU) used in this work was donated by NVIDIA Corporation.

References

- [1] Pasqualotto R, Nielsen P, Gowers C, Beurskens M, Kempenaars M, Carlstrom T, Johnson D and JET-EFDA Contributors 2004 *Review of Scientific Instruments* **75** 3891–3893
- [2] Frassinetti L, Beurskens M N A, Scannell R, Osborne T H, Flanagan J, Kempenaars M, Maslov M, Pasqualotto R, Walsh M and JET-EFDA Contributors 2012 *Review of Scientific Instruments* **83** 013506
- [3] Maggi C, Frassinetti L, Horvath L, Lunniss A, Saarelma S, Wilson H, Flanagan J, Leyland M, Lupelli I, Pamela S, Urano H, Garzotti L, Lerche E, Nunes I, Rimini F and JET Contributors 2017 *Nuclear Fusion* **57** 116012
- [4] Sirinelli A, Alper B, Bottureau C, Clairet F, Cupido L, Fessey J, Hogben C, Meneses L, Sandford G, Walsh M J and JET-EFDA Contributors 2010 *Review of Scientific Instruments* **81** 10D939
- [5] Laviron C, Donn e A J H, Manso M E and Sanchez J 1996 *Plasma Physics and Controlled Fusion* **38** 905
- [6] Mazzucato E 1998 *Review of Scientific Instruments* **69** 2201–2217
- [7] Ham C, Bokshi A, Brunetti D, Ramirez G, Chapman B, Connor J, Dickinson D, Field A, Frassinetti L, Gillgren A, Graves J, Kiviniemi T, Leerink S, McMillan B, Newton S, Pamela S, Roach C, Saarelma S, Simpson J, Smith S, Solano E, Strand P, Virtanen A and JET Contributors 2021 *Nuclear Fusion* **61** 096013
- [8] Garcia J 2022 Deuterium-tritium experiments in JET with the ITER-like wall *6th Asia-Pacific Conference on Plasma Physics*
- [9] Garzotti L, Challis C, Dumont R, Frigione D, Graves J, Lerche E, Mailloux J, Mantsinen M, Rimini F, Casson F, Czarnecka A, Eriksson J, Felton R, Frassinetti L, Gallart D, Garcia J, Giroud C, Joffrin E, Kim H T, Krawczyk N, Lennholm M, Lomas P, Lowry C, Meneses L, Nunes I, Roach C, Romanelli M, Sharapov S, Silburn S, Sips A, Stef anikov a E, Tsalas M, Valcarcel D, Valovi c M and JET Contributors 2019 *Nuclear Fusion* **59** 076037
- [10] Agarap A F 2018 *arXiv preprint arXiv:1803.08375*
- [11] Kingma D P and Ba J 2014 *arXiv preprint arXiv:1412.6980*
- [12] Frassinetti L, Saarelma S, Verdoolaege G, Groth M, Hillesheim J, Bilkova P, Bohm P, Dunne M, Fridstr om R, Giovannozzi E, Imbeaux F, Labit B, de la Luna E, Maggi C, Owsiak M, Scannell R and contributors J 2021 *Nuclear Fusion* **61** 016001
- [13] Saarelma S, Beurskens M, Dickinson D, Frassinetti L, Leyland M, Roach C and EFDA-JET Contributors 2013 *Nuclear Fusion* **53** 123012
- [14] Banerjee S, Mordijck S, Barada K, Zeng L, Groebner R, Osborne T, Rhodes T, Snyder P, Grierson B and Diallo A 2021 *Nuclear Fusion* **61** 056008
- [15] Behn R, Alfier A, Medvedev S Y, Zhuang G, Pasqualotto R, Nielsen P, Martin Y and the TCV team 2007 *Plasma Physics and Controlled Fusion* **49** 1289
- [16] Callen J, Groebner R, Osborne T, Canik J, Owen L, Pankin A, Rafiq T, Rognl en T and Stacey W 2010 *Nuclear Fusion* **50** 064004
- [17] Frassinetti L 2023 Understanding pedestal structure, stability and turbulent transport in D, DT and T plasmas in JET ILW *49th European Conference on Plasma Physics*