IMPROVING THE TIME RESOLUTION OF THOMSON SCATTERING VIA MACHINE LEARNING ON REFLECTOMETRY DATA

D. R. Ferreira[1], A. Gillgren[2], A. Ludvig-Osipov[2], P. Strand[2] and JET Contributors[*]

[1]Instituto de Plasmas e Fusão Nuclear, Instituto Superior Técnico, Universidade de Lisboa, Portugal
[2]Department of Space, Earth and Environment, Chalmers University of Technology, Sweden
[*]See the author list of J. Mailloux et al 2022 Nucl. Fusion 62 042026
diogo.ferreira@tecnico.ulisboa.pt

Abstract

At the Joint European Torus (JET), the reference diagnostic to measure electron density is Thomson scattering. However, the low sampling rate of this diagnostic makes it impractical to study the dynamics of the density profile. In this work, we use machine learning to predict the density profile based on data from another diagnostic, namely reflectometry. The proposed model is trained to transform reflectometry data into Thomson scattering profiles, and is able to generate density profiles at a much higher sampling rate than Thomson scattering, and more accurately than reflectometry alone. This enables the study of pedestal dynamics and other edge phenomena.

1. INTRODUCTION

High resolution Thomson scattering (HRTS) is the reference diagnostic for measuring the density profile at JET (Pasqualotto et al., 2004). This diagnostic provides good accuracy in terms of radial position (Frassinetti et al., 2012) but has a low sampling rate (20 Hz), which makes it infeasible to analyze pedestal dynamics in detail, namely phenomena such as edge localized modes (ELMs). These occur on faster time scales and would require a sampling rate on the order of at least 1 kHz to capture the transient processes associated with ELM crashes.

On the other hand, the reflectometry diagnostic at JET (Sirinelli et al., 2010) provides a high temporal resolution, and is able reconstruct the density profile at a sampling rate that is typically in the range of 1–10 kHz. However, this diagnostic is not as accurate in terms of the radial position of its density measurements, and over the years new and improved methods have been developed for the reconstruction of the density profile from reflectometry data (Morales et al., 2017).

The two diagnostics are based on different principles. Thomson scattering, which we will refer to as HRTS, is based on laser scattering, where density values are measured at fixed positions from the intensity of scattered light. On the other hand, the reflectometry diagnostic, which we will refer to as KG10 (its internal name at JET), is based on probing the plasma with electromagnetic waves over a wide range of frequencies (44–150 GHz). Each frequency value corresponds to a density value that is known beforehand, because that is the point where reflection occurs; however, the position of that point will have to be calculated from the travel time of the reflected wave.

In summary, HRTS measures density values at precise positions but with some uncertainty in those measurements, whereas KG10 measures precise density values but with some uncertainty in the positions of those measurements. Our goal is to take advantage of the high temporal resolution of reflectometry in order to derive Thomson scattering profiles at a much higher sampling rate than this diagnostic is able to provide. For this purpose, we use machine learning to train a neural network that maps KG10 to HRTS data. Once trained, the model can generate density profiles similar to HRTS at the sampling rate of KG10, even when HRTS is not available. This can be used, for example, to study pedestal dynamics in detail, as we illustrate in the sections below.

2. MACHINE LEARNING APPROACH

Figure 1 shows two examples of how the two diagnostics may agree or disagree on the density profile. On the left-hand side, Figure 1 shows an example where both diagnostics are in good agreement: KG10 is able to probe the profile further inwards, but both diagnostics agree on the

position of the pedestal, for example. On the right-hand side, Figure 1 shows an example where there are miscalculations in the positions of the density values measured by KG10. When those errors occur, they tend to accumulate over subsequent positions, resulting in a density profile that does not match the one provided by HRTS. Fortunately, these problems occur in a small percentage of cases (~5%) and it would be desirable, if possible, to correct them via machine learning.
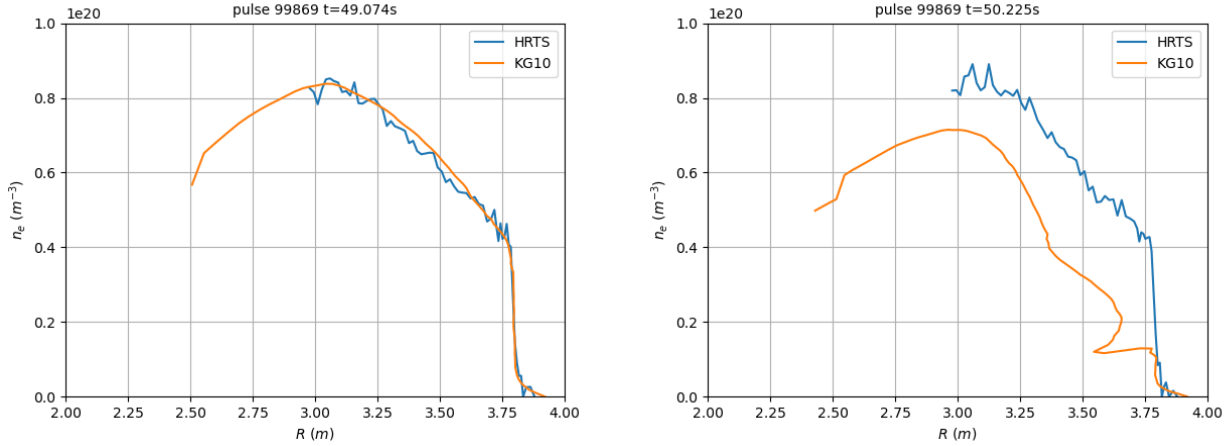

Figure 1. Comparison of HRTS and KG10 profiles at two different times in the same pulse.

To train a machine learning model to map KG10 to HRTS data, it is necessary to have instances where data from both diagnostics are available. In other words, we need training examples where both the HRTS and KG10 profiles are known. Since the sampling rate of KG10 is much higher than that of HRTS, in general for any given HRTS profile it is possible to find a KG10 profile that is close in time. The difference in time will be at most half the sampling period of KG10, which is on the order of 10^{-4} s. We refer to this process as the time-syncing of HRTS and KG10.

After the model is trained, it can be used to predict the HRTS profile when only the KG10 profile is known. This yields a virtual HRTS diagnostic with the sampling rate of KG10.

The model itself is a 4-layer neural network with a 100-dimensional input, which corresponds to the size of a KG10 profile, and a 63-dimensional output, which corresponds to the size of an HRTS profile. As shown in Figure 2, each dense layer has 1024 units, and it includes a ReLU (rectified linear unit) activation. The total number of parameters exceeds 3×10^6. The generous number of layers and parameters are motivated by the need to transform KG10 into HRTS data as accurately as possible, while overcoming the occasional issues with KG10 profiles. For this purpose, it turns out that having more layers, and hence more non-linearities, is more beneficial than having fewer and wider layers, and the model is also easier to train.
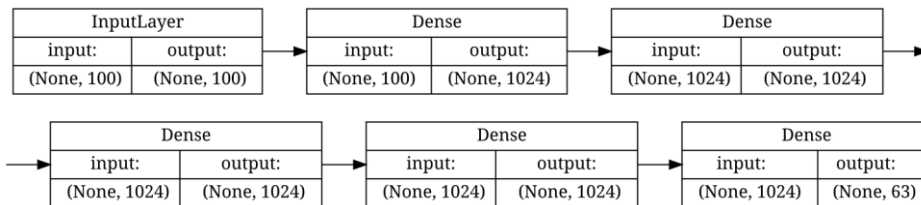
Figure 2. Structure of the machine learning model.

3. TRAINING DATA

Before collecting the training data, we ensured that both diagnostics were fully operational for any experimental session that we selected for the training examples; this is the case from December 2020 onwards. In addition, the training examples should be representative of the type of experiments performed at JET; for this purpose, we selected experimental sessions from the baseline and hybrid scenarios (Garzotti et al., 2019). Specifically, we focused on the baseline and

hybrid experiments performed in the C40 campaign (March–July 2021) and in the C41 campaign (August–December 2021), which was the second Deuterium-Tritium campaign at JET (DTE2), after the first D-T campaign back in 1997.

From these experiments, we identified a total of 170 pulses, in the range of pulse numbers from 98794 to 99953. Since HRTS has a sampling rate of 20 Hz, and each pulse lasts for about 30 s, we expected a number of training examples on the order of $170 \times 30 \times 20 \approx 10^5$. In practice, we obtained 43 531 training examples of KG10–HRTS profile pairs. This was due to a number of reasons, including the fact that, while HRTS is operational during the entire pulse, KG10 has a shorter operating window because of limited data buffer capacity.

The training examples have been split into 90% for training and 10% for validation, and the model took 5 minutes to train on a single GPU (graphics processing unit). To analyze the variance of the results, we repeated the training procedure across a 10-fold cross validation.

4. RESULTS

Figure 3 shows the predictions of the density profile for the same pulse and times that were illustrated before in Figure 1. Here, the red line is the mean prediction of the 10 models trained by cross validation, and the light-red band around it denotes the standard deviation. On the left-hand side in Figure 3, where KG10 and HRTS agree, the prediction also coincides with HRTS. The band around the prediction is hardly noticeable because all the models yield almost identical results. In the plot on the right-hand side, the model again provides a prediction that coincides with HRTS, despite the large gap that exists between the HRTS and KG10 profiles. We conclude that all those models can overcome the problems in the KG10 profile by providing predictions that are consistently close to the HRTS profile.
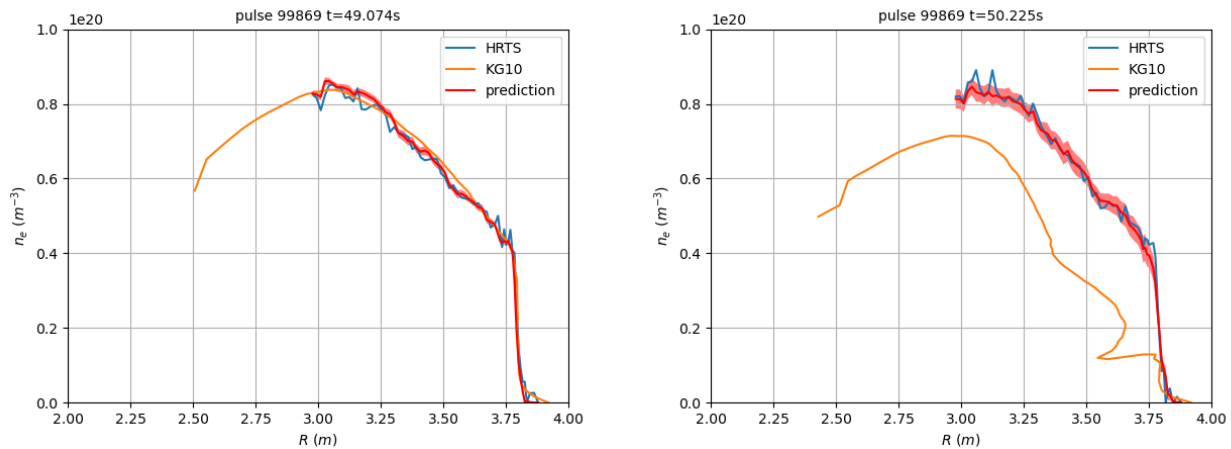
Figure 3. Prediction of HRTS profile from KG10 data at two different times in the same pulse.

The pulse that has been used to produce these results is in fact the record-breaking JET pulse of the Deuterium-Tritium campaign (DTE2). This pulse has not been used for training or validation, so the results in this section are meant to illustrate how the model performs on unseen data.

Our main goal, however, is to use the model to study the pedestal dynamics, particularly the processes at the time scale of an ELM cycle. Figure 4 illustrates two views over an ELM cycle at around t = 50.0 s in the same pulse. The density profile is predicted at the sampling rate of KG10, and is averaged over a sliding window of 1 ms. Figure 4 shows how the top of the pedestal (at around position R = 3.8 m) gradually builds up, until it crashes abruptly and a new build-up process is initiated. This particular ELM cycle lasts for about 30 ms, which is consistent with the range of 15–40 Hz for ELM frequencies at JET (Lennholm et al., 2015).
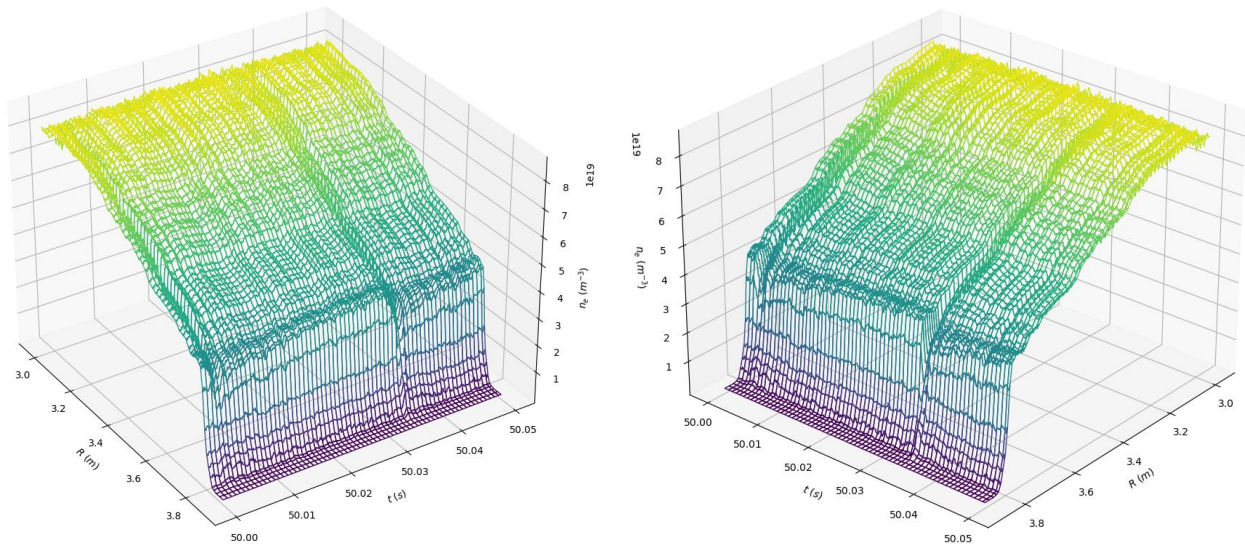
Figure 4. Two views of the course of an ELM cycle at around t=50.0s in pulse 99869.

5.   CONCLUSION

The proposed model predicts HRTS-like profiles from KG10 data assuming that HRTS provides the ground truth for the density profile, but it should be noted that HRTS itself is sometimes calibrated on KG10. Although our focus is on predicting the HRTS profile, we have not relegated KG10 to secondary importance; on the contrary, KG10 is essential to enable this approach, even with its occasional issues in position accuracy. These issues were mitigated by using a non-linear model with multiple layers and a large number of parameters. Still, the model is fast to train and use, and can be applied on any pulse with KG10 data, enabling the analysis of pedestal dynamics, ELMs, and other features of the density profile.

6.   REFERENCES

L. Frassinetti, M. N. A. Beurskens, R. Scannell, T. H. Osborne, J. Flanagan, M. Kempenaars, M. Maslov, R. Pasqualotto, M. Walsh and JET-EFDA Contributors (2012). Spatial resolution of the JET Thomson scattering system. Rev Sci Instrum 83: 013506.

R. Pasqualotto, P. Nielsen, C. Gowers, M. Beurskens, M. Kempenaars, T. Carlstrom, D. Johnson and JET-EFDA Contributors (2004). High resolution Thomson scattering for Joint European Torus (JET). Rev Sci Instrum 75: 3891–3893.

A. Sirinelli, B. Alper, C. Bottereau, F. Clairet, L. Cupido, J. Fessey, C. Hogben, L. Meneses, G. Sandford, M. J. Walsh and JET-EFDA Contributors (2010). Multiband reflectometry system for density profile measurement with high temporal resolution on JET tokamak. Rev Sci Instrum 81: 10D939.

R. B. Morales, S. Hacquin, S. Heuraux, R. Sabot (2017). New density profile reconstruction methods in X-mode reflectometry. Rev Sci Instrum 88: 043503.

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, Hyun-Tae 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 and JET Contributors (2019). Scenario development for D-T operation at JET. Nucl Fusion 59: 076037.

M. Lennholm, P.S. Beaumont, I.S. Carvalho, I.T. Chapman, R. Felton, D. Frigione, L. Garzotti, A. Goodyear, J. Graves, D. Grist, S. Jachmich, P. Lang, E. Lerche, E. de la Luna, R. Mooney, J. Morris, M.F.F. Nave, F. Rimini, G. Sips, E. Solano, M. Tsalas and JET EFDA Contributors (2015). ELM frequency feedback control on JET. Nucl Fusion 55: 063004