Diag2Diag: Multi modal super resolution for physics discovery with application to fusion

Azarakhsh Jalalvand[1*], Max Curie[1], SangKyeun Kim[2],
Peter Steiner[1], Jaemin Seo[3], Qiming Hu[2],
Andrew Oakleigh Nelson[4], Egemen Kolemen[1,2*]

[1*]Department of Mechanical and Aerospace Engineering, Princeton University, Olden Street, Princeton, 08540, NJ, USA.
[2]Princeton Plasma Physics Laboratory, Stellarator Road, Princeton, 08543, NJ, USA.
[3]Department of Physics, Chung-Ang University, Seoul, South Korea.
[4]Applied Physics and Applied Mathematics, Columbia University, W. 120th Street, New York, 10027, NY, USA.

*Corresponding author(s). E-mail(s): azarakhsh.jalalvand@princeton.edu; ekolemen@princeton.edu;

Abstract

This paper introduces a groundbreaking multi-modal neural network model designed for resolution enhancement, which innovatively leverages inter-diagnostic correlations within a system. Traditional approaches have primarily focused on uni-modal enhancement strategies, such as pixel-based image enhancement or heuristic signal interpolation. In contrast, our model employs a novel methodology by harnessing the diagnostic relationships within the physics of fusion plasma. Initially, we establish the correlation among diagnostics within the tokamak. Subsequently, we utilize these correlations to substantially enhance the temporal resolution of the Thomson Scattering (TS) diagnostic, which assesses plasma density and temperature. By increasing its resolution from conventional 200 Hz to 500 kHz, we facilitate a new level of insight into plasma behavior, previously attainable only through computationally intensive simulations. This enhancement goes beyond simple interpolation, offering novel perspectives on the underlying physical phenomena governing plasma dynamics.

Keywords: Diagnostics, Machine Learning, Fusion pilot plant, Synthetic Diagnostic

1

1 Introduction

The success of nuclear fusion for energy production strongly depends on controlling the plasma condition to achieve the highest performance while avoiding unstable regimes. It was recently shown that Artificial Intelligence (AI) can be a helpful tool [1] to achieve that goal. The experimental fusion reactors such as DIII-D make use of a variety of diagnostics [2], where each diagnostic provides useful insights into the plasma. All diagnostics primarily consist of time-series data, yet they exhibit distinct characteristics. While some diagnostics record single values within the plasma at specific time points, others, such as Electron Cyclotron Emission (ECE), CO_2 interferometer (Interferometer), Motional Stark Effect (MSE), Charge Exchange Recombination (CER), Magnetic probes (Magnetics) or Thomson Scattering (TS) capture information about the plasma's properties at multiple spatial locations simultaneously.

The different diagnostics form a complementary set for extracting as much information from the plasma as possible. For example, ECE measures electron temperature [3], Interferometer measures electron density and density fluctuations [4], MSE measures the magnetic field [5], and TS measures the electron temperature and density [6]. Obviously, these measurements capture different physical properties of the plasma. Although it is likely that there exists some kind of correlation or coupling between the different diagnostics, the exact relationship cannot be specified analytically. Machine Learning (ML) can be an interesting candidate for identifying hidden relationships in data. Learning the hidden relationships among different diagnostics would be a great asset to enhance their measurements, and it also helps to find a minimal set of diagnostics for the Fusion Pilot Plant (FPP) in which the availability of diagnostics is limited due to the cost and hardware constraints.

A limitation of some diagnostics, such as TS is the low temporal resolution of only 200 Hz, which does not allow for detecting and tracking fast events like Edge Localized Mode (ELM) (≤ 1 ms). Figure 1 shows an example of missing ELM in a discharge, due to the low temporal resolution of TS. Nevertheless, it is still important to detect such events reliably, as they can have a strong impact on plasma behavior. In DIII-D there is a specific operational method to increase the sampling rate of TS up to 10 kHz (aka "Burst Mode") [7]. However due to its limitations including high energy consumption, equipment stress, heat generation, complex data analysis, and limited measurement repetition, burst mode is typically reserved for specific experimental runs or particular phases of reactor operation where high-resolution temporal data is crucial.

On the other hand, diagnostics like Interferometer, ECE, have much higher temporal resolution with sampling frequencies around MHz, which allows for a much more detailed analysis of the plasma. However, these diagnostics have different characteristics compared to TS. While TS offers detailed insights into both electron density and temperature with high accuracy, it requires complex setups and is usually more resource-intensive. A Interferometer provides a more straightforward approach to measuring electron density, excelling in situations that require rapid response and continuous monitoring. Furthermore, ECE and TS are both pivotal diagnostic tools used in tokamaks for measuring electron temperature, yet they operate on distinctly different principles and offer unique advantages. ECE utilizes the natural microwave emissions from electrons gyrating around magnetic field lines to provide excellent temporal
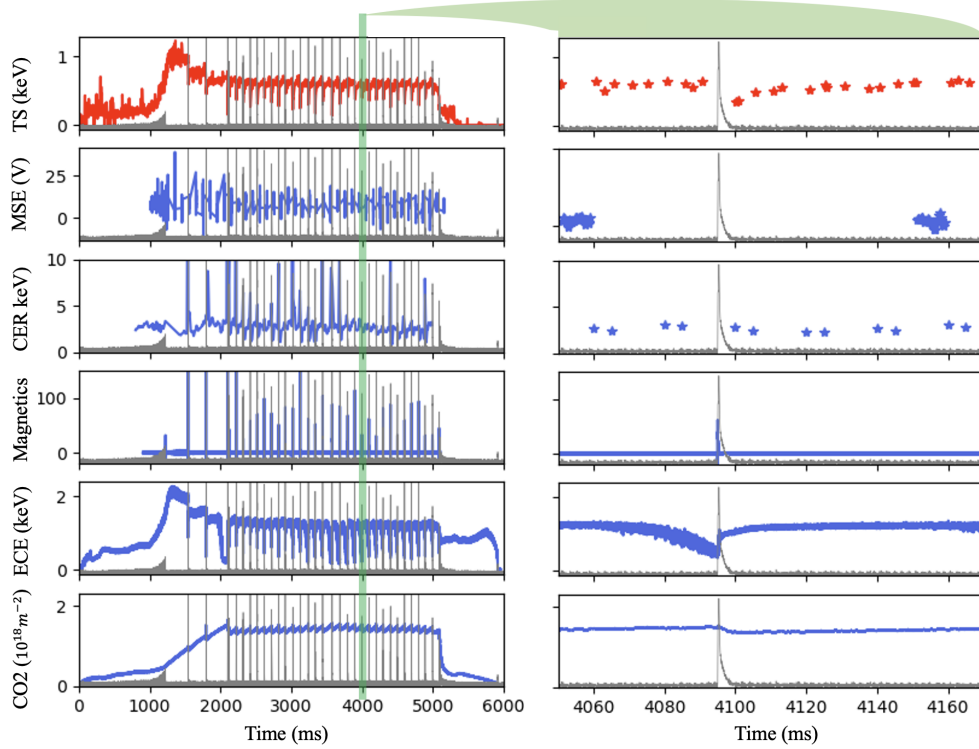
Fig. 1: (Left) Example of TS signal for DIII-D discharge 174823, in red, along with the D_α measurements, in gray, as an indicator of ELM, and a collection of other diagnostics that will be used to increase the resolution of TS. (Right) The same data as left but zoomed in to show the sampling points of the diagnostics around one example of ELM event. Due to the low sampling rate, this ELM is not observed by TS. However, thanks to their high temporal resolution, diagnostics including ECE, Interferometer, and Magnetics capture that.

resolution, allowing for the monitoring of rapid plasma changes and instabilities, though its effectiveness can be limited by variations in magnetic field strength. On the other hand, TS involves firing a laser into the plasma and analyzing the scattered light, which provides robust, absolute measurements of both electron temperature and density with less susceptibility to magnetic influences. While ECE excels in continuous data collection and fine temporal analysis, TS offers superior spatial resolution and is less dependent on external conditions, making it invaluable for comprehensive, though typically less frequent, plasma evaluations. If it would be possible to find a correlation between those high-resolution diagnostics and TS, this would be useful for developing new physical analyses.

To address these limitations of TS, we hypothesize that a data-driven model with a multimodal input comprising CER, Interferometer, ECE, Magnetics, and MSE can

effectively make use of internal correlations in order to estimate TS. We furthermore hypothesize that we can use such a model to generate additional samples between two consecutive TS measurements and thus generate a so-called **super-resolution** TS diagnostic as a tool for deeper physical analysis of plasma behavior. Our proposed model can generate high spatiotemporal resolution mappings of key plasma properties like electron temperature and density, enabling enhanced analysis capabilities beyond what is possible from the existing low-resolution TS diagnostic alone.

In recent years different kinds of a Neural Network (NN) have been used for upsampling video data [8–11] and for radar data [12–15]. These approaches are typically some kind of non-linear interpolation to add frames between existing video frames. More examples for ML-based upsampling were proposed for medical data [16] and for audio data [17–20]. Similar to the video upsampling approaches, these approaches are some kind of non-linear interpolation as well. In [21], an alternative to interpolation is suggested to estimate missing data in temporal data streams. It is kind of a multi-modal approach, because it fuses different kinds of information. However, the algorithm is limited towards estimating missing data or dealing with irregularly sampled data. Approaches like these work well for enhancing existing sequences, which are quasi-stationary in a way such that consecutive frames or samples do not change very fast. However, in nuclear fusion, many spurious events like ELM can happen between two TS samples. By interpolating between consecutive TS samples, regardless of linearly or non-linearly, it is likely that we would miss such spurious events. In this paper, we thus develop a novel method to generate additional TS samples based on other diagnostics. This is roughly inspired by other multi-modal ML approaches, such as [22], where it was proposed to fuse Radar and camera data for an enhanced distance estimation. This is a multi-modal approach and thus related to our approach, or [23], where machine learning was used to reveal the control mechanics of an insect wing hinge. This was also a multimodal approach in a way that the ML algorithm received different features recorded from flying insects. However, similar to the other approaches, no attempts to upsampling or estimating missing/in-between data are made.

In particular, we train a NN on aligned samples from CER, Interferometer, ECE, glsmag, MSE as input to estimate TS diagnostic. Later, we use the trained NN model to generate additional TS samples with a high temporal resolution. We evaluate this model in several ways:

- First of all, we show that we can estimate one diagnostic from another one, particularly Interferometer from ECE, because there must exist intrinsic correlations between these diagnostics, although they cannot be described analytically. We compare the measured and estimated Interferometer spectrograms based on similarity scores and based on the task of Alfvén Eigenmode (AE) detection with the model [24].
- Next, we train a model to estimate TS samples from CER, Interferometer, ECE, Magnetics and MSE samples. We again compare the measured and estimated TS diagnostics based on similarity scores.
- In a final step, we use the trained TS estimation model to generate super-resolution TS diagnostic and present its capability for investigating the ELM cycles [25] and the mechanism of Resonant Magnetic Perturbation (RMP) effect on the pedestal degradation via RMP modulation [26].

We postulate that this method can be transferred to further applications ranging from astrophysics to other multi-modal time-series data.
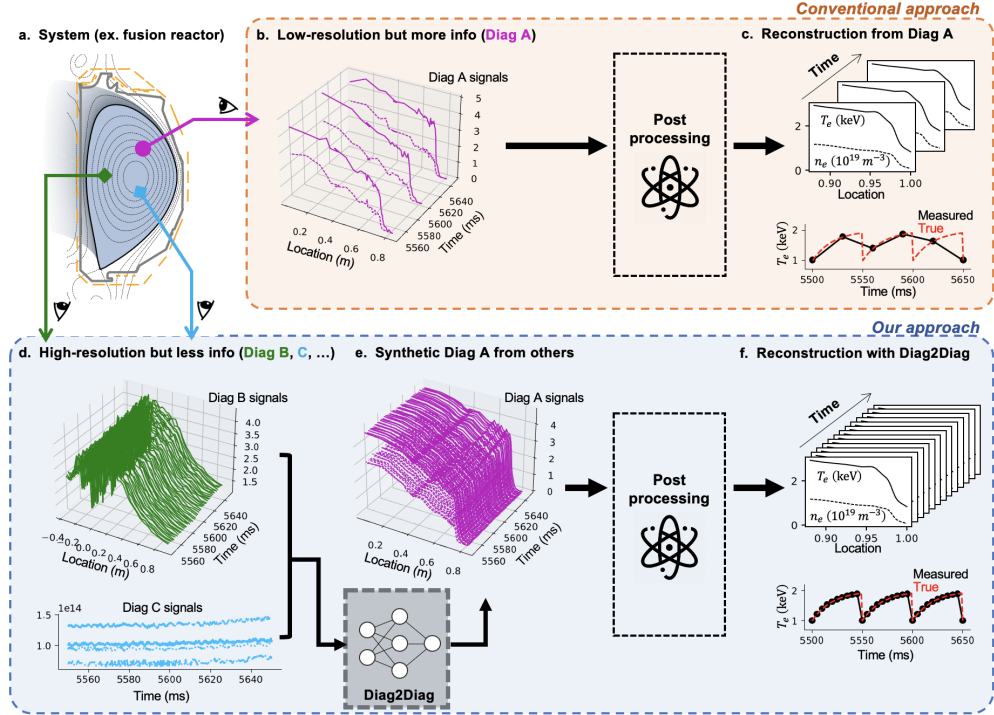

Fig. 2: Main methodology. (a) The configuration of diagnostics in the system. (b) The low temporal resolution diagnostics. (c) Low resolution profile extracted from the diagnostic. (d) High temporal resolution diagnostics. (e) Synthetic high resolution diagnostics generated using Diag2Diag. (f) High resolution profile extracted from the synthetic diagnostic.

Figure 2 summarizes the main methodology for this work. DIII-D utilizes hundreds of diagnostics for observing the plasma. It is illustrated that these diagnostics have different temporal resolutions, and they have intrinsic correlations that are difficult to describe analytically. A potential ML model can map between the different diagnostics and thus learn exactly these intrinsic correlations. This works for both, time-series and spectrograms. For the experiments on time-series and spectrogram data, different variants of NN are used. The design choices and the optimization and training strategies are described in the following section.

2 ML-based mapping between different diagnostics

For developing an ML-based super-resolution TS diagnostic from other diagnostics, it is essential to verify the existence, strength and robustness of correlation among them. We therefore approach this in several steps as described subsequently.
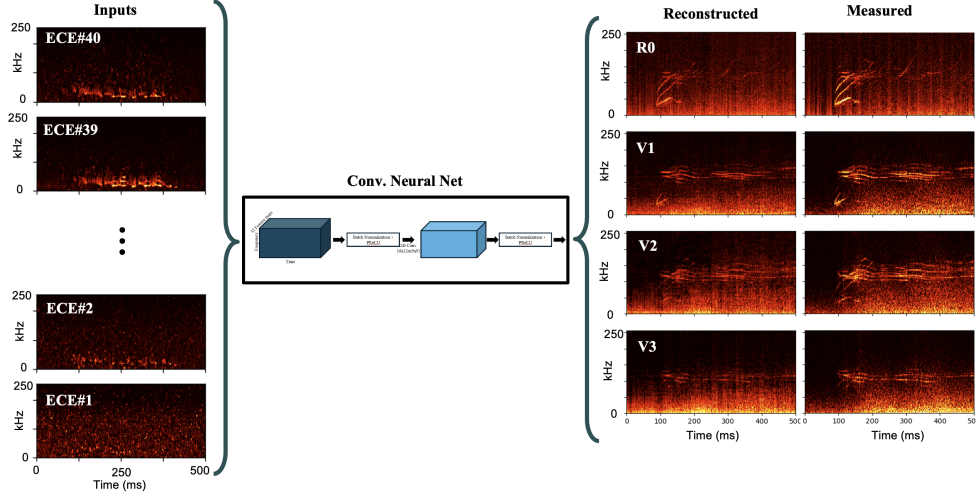

Fig. 3: Mapping between spectrograms from different diagnostics. ECE spectrograms form a 40-channel tensor as the input to the Convolutional Neural Network (CNN). The outputs are the four channels of Interferometer spectrograms.

The aim of the first step is to show that we can estimate the spectrograms of one diagnostic based on another one. As it was discussed in Section 1, it is very likely that different diagnostics have intrinsic correlations. For example, based on the physics of thermodynamics there should be a correlation between electron temperature, which is measured by ECE, and density fluctuations measured by Interferometer. We now show that a NN is able to learn exactly this relationship by mapping from ECE spectrograms to Interferometer spectrograms as it is illustrated in Figure 3.

The starting point for this experiment is the raw time-series of the diagnostics, i.e., ECE with 40 channels, and Interferometer with 4 channels. For each channel, spectrograms were computed independently. The ECE spectrograms were computed with a window size of 512 samples, and with a hop size of 256 samples. The Interferometer spectrograms were computed with a window size of 1536 samples, and with a hop size of 768 samples. In this way, it was ensured that the different magnitude spectrograms were aligned in time. Afterwards, the linear magnitudes are converted to a logarithmic scale, and the spectrograms were clipped and rescaled to the range of $[0, 1]$.

The resulting multi-channel ECE spectrograms were used as the input to a CNN, and the multi-channel Interferometer spectrograms were used as the target outputs. We optimized all important hyper-parameters based on the R^2 score to maximize the similarity between the ground truth and the estimated outputs on the training

and validation sets. Details about the optimization can be found in Section 5.3. The best-performing model is a CNN that transforms the ECE spectrograms with 40 channels subsequently to 32, 16 and 8 feature maps and finally to the Interferometer spectrograms with 4 channels. For each feature map, 2D filter kernels with a size of 7×7 are used. Batch normalization was used separately for each channel, and parametric ReLU activation functions were used after each batch normalization layer. The model had in total $95\,823$ trainable parameters (i.e., filter kernels for each feature map, batch normalization parameters, and negative slope of the parametric ReLU activation function).

Besides the image comparison, we are also interested to study how much of the effective physics information is preserved using this method. Therefore we evaluate the best performing model on the downstream task AE detection. We use the model proposed in [24] to detect AE based on the computed Interferometer spectrograms and based on the estimated Interferometer spectrograms. From Figure 4, we can visually observe that the CNN model was able to reconstruct the pronounced coherent activities in the spectrogram, and the Toroidal AE (TAE) scores also show that there is enough information in the reconstructed spectrogram to detect the AE mode after applying the predefined threshold [24].
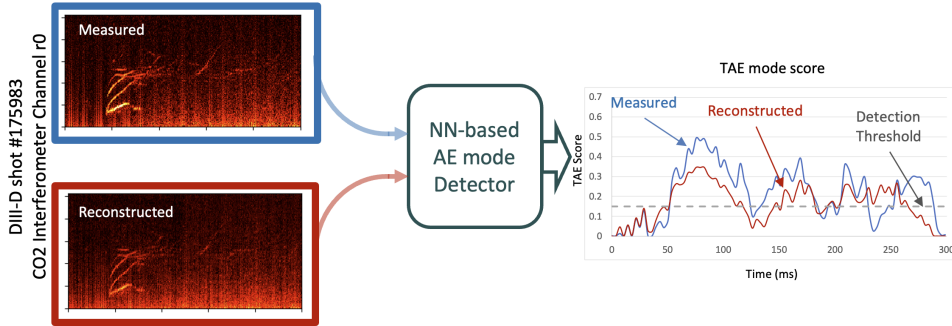

Fig. 4: AE detector output for measured and reconstructed spectrograms.

The results of this phase support that ML models can learn the correlation amongst diagnostics. We can now go one step further and face the estimation of TS amplitude from other diagnostics based on raw time-series.

3 Multi-modal super-resolution diagnostic

In the previous section, we showed the capability of data-driven models in learning the correlation among Interferometer and ECE channels, by showing that the coherent and physically meaningful patterns in a spectrogram of a diagnostic can be reconstructed from other diagnostics. In this section, we switch from spectrograms to time-series signals and show that the amplitude of a diagnostics can be reconstructed from other diagnostics. More importantly, we will show that if the input diagnostics are of much

7

higher temporal resolution compared to the target one, such a model can be used to increase the time resolution of the time resolution of the target signals in a much more intelligent way compared to the conventional uni-modal interpolations. As a use case, we target TS as one of the most important diagnostics to measure the density and pressure profile of plasma. However as mentioned earlier, its low sampling rate of 200 Hz is a bottleneck in studying the evolution of plasma in the rapidly changing events such as RMP effect on the pedestal degradation and ELM.

As shown in Figure 5(a), we consider a suit of diagnostics including Interferometer, ECE, Magnetics, CER, and MSE with sampling frequencies of 1.66 MHz, 500 kHz, 200 kHz, 200 Hz, and 4 kHz, respectively.

To obtain a dataset suitable for this task, all the included diagnostics are aligned with the TS sampling time steps by matching their most recent measured sample. In this way, we create a dataset with which we train the neural network for this task. Since the sampling steps of TS are not uniform (See Fig. 1), we did not use a conventional NN for time-series data analysis such as Long-Short term memory cell (LSTM) or Recurrent Neural Network (RNN). Instead, we opted for memory-less models, namely Multilayer Perceptron (MLP). However, we included the first and second derivatives of the high-resolution input diagnostics, ECE and Interferometer, to include the temporal evolution information.

5(b) shows the output of the Diag2Diag model for one channel of the TS density when the model has been trained and validated on the measured TS datapoints. With an average R^2 score of 0.92 on the validation set, one can argue that the Diag2Diag NN is indeed able to reconstruct TS from the input diagnostics. In 5(b-2) we highlighted one of the discrepancies and showed that it is in fact still in the error bar of the measured data. 5(b-3) and (b-4) present the agreement of the measured and reconstructed TS on a specific time for all 40 channels, across the spatial domain, also with a focus on the pedestal region.

Now that the model has been successfully trained to reproduce the measured TS, we can feed it with higher resolution inputs and generate super-resolution TS. To do so, we aligned all the input diagnostics with the sampling rate of ECE data, i.e., 500 kHz. Figure 5(c) shows the same TS channel as Figure 5(b), but now in much higher resolution. Figure 5(c) illustrates how the new TS data points help to fill the gaps between the measured points.

Obviously, there is no ground truth to evaluate the accuracy of the super-resolution TS. However, we can investigate whether the generated diagnostics can help to verify the physics hypotheses about plasma profiles that have been proposed theoretically, or by simulations, but have never been visualized with the experimental data due to the lack of plasma profile resolution. In the following subsections we study two phenomena namely, the mechanism of RMP effect on the pedestal degradation via RMP modulation, and ELM cycle analysis.

3.1 Case 1: Investigating ELM cycles in DIII-D

In order to let fusion energy be a viable energy source, it must achieve significant fusion gain through continuous fusion reactions. A prominent method to reach this objective is operating a tokamak in high confinement mode (H-mode), which has a narrow
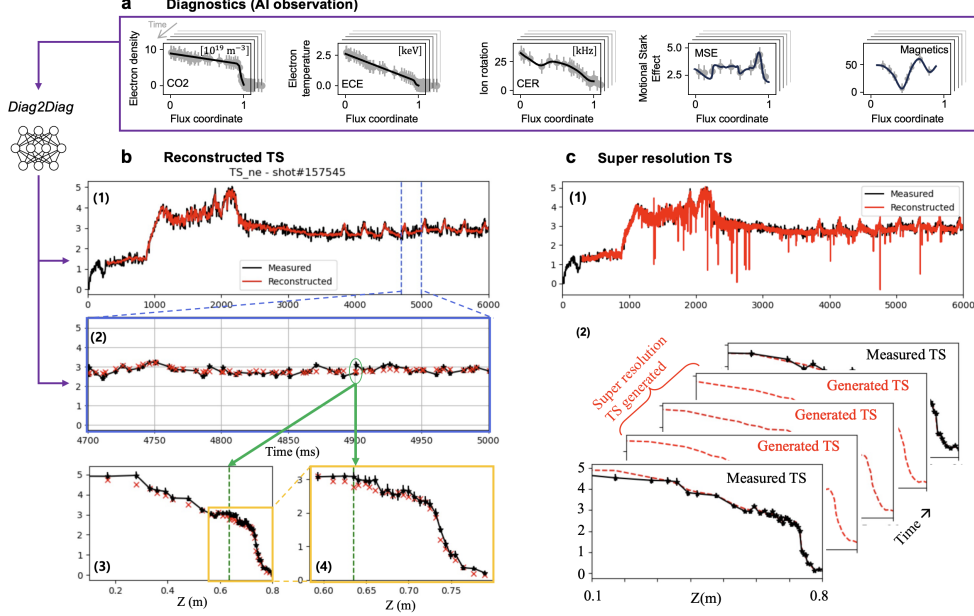
Fig. 5: (a) The input diagnostics to Diag2Diag NN. (b-1) Validation of the reconstructed TS-density when the model is supplied with the downsampled inputs at 200Hz to match the measured TS sampling rate for shot 157545 at Z=0.7 m in time. (b-2) The same comparison as (b-1) but zoomed in to the range of 4000 ms and 4300 ms. (b-3) Comparison of the measured and reconstructed TS-density for shot 157545 at time 4050.5 ms. (b-4) The same comparison as (b-3) but zoomed in to the pedestal part. (c-1) Super resolution TS generated by Diag2Diag NN using the full-resolution inputs at 500 kHz. (c-2) Illustration of the generated TS between to consecutive measurements.

edge transport barrier, also known as the pedestal. This feature significantly boosts plasma confinement within the reactor, enhancing fusion power and efficiency. However, operating in H-mode introduces a steep pressure gradient at the pedestal, leading to substantial operational risks. This gradient drives hazardous edge energy bursts due to a plasma instability known as ELMs. These bursts lead to sudden drops in the energy at the pedestal, causing severe, transient heat fluxes on the reactor walls. This results in damaging material erosion and potential surface melting, with heat energy reaching approximately $20\,\mathrm{MJ\,m^{-2}}$, which is an unacceptable level for fusion reactors. Therefore, to advance tokamak designs toward practical application in fusion energy, it is crucial to develop dependable methods to consistently suppress these edge burst events.

As presented in Fig. 1, in DIII-D discharge 174823 [25], there are over 20 highly similar repeated cycles of energy growth ending in explosive bursts (ELMs). This activity modifies the electron temperature in highly regular manner, leading to growth and decay with a period around 100 ms. Since the ELM itself is a rapid energy burst, the temperature changes drastically at during each ELM event. As such, one can easily find the ELM time by identifying times when $|\mathrm{ECE}(t + \delta t) - \mathrm{ECE}(t)|$ peaks, using

9

$\delta t = 0.2\,\text{ms}$. Importantly, while the temperature fluctuations caused by ELMs in this discharge can be seen with the high-temporal resolution diagnostics, TS measurements are too slow to capture the whole behavior.

Interestingly, the extremely regular and slow nature of the ELMs in this discharge, which is relatively rare in tokamak plasmas, allows experimental investigation into the inter-ELM dynamics of the electron temperature and density profiles [25] that would otherwise be impossible with TS diagnostics. The highly reproducible nature of ELM activity enables an overlay of all the cycles on top of each other by aligning the ELM events. By overlaying the accumulated TS measurements in this manner, the dynamics of ELM cycle shown on the electron temperature and density profiles can be studied, as shown in the left two plots in Fig. 6.

In more general plasmas, the ELMs are more irregular, faster and harder to define, preventing robust inter-ELM profile analysis with TS. This can already be observed in our test discharge, as shown in the right two plots of Fig. 6 which show the TS measurements during one ELM cycle. Notably, the ELM events themselves present a rapid pressure decrease that cannot be captured with the slow sampling rate of TS. To develop analysis tools that can be exploited in generalized discharges, we can use the unique effective up-sampled TS measurements obtained by accumulating measurements over repeated ELM cycles in this discharge as ground truth to compare with the super-resolution TS generated by Diga2Diag NN in more general cases.

The super-resolution TS enables us to have a high temporal resolution of n_e and T_e as shown in Fig. 6. The accumulated Thompson scattering time series (left plots) and the up-sampled TS (right plots) present a similar trend, while the up-sampled TS (dots) match perfectly with TS measurements (crosses). This plot showcases the ability of up-sampled TS to capture the dynamics of the plasma with high accuracy.

This exercise shows that up-sampled TS (light dots right plots in Fig. 6) tracking nicely with accumulated TS (left plots in Fig. 6) which can be used as the ground truth. The actual TS (crosses with vertical bars right plots in Fig. 6) almost sits on top of the up-sampled TS shows that the TS not only provides the information of repeated trends, but also the dynamics of the electron that is unique to each cycle.

3.2 Case 2: Unveiling diagnostic evidence of RMP mechanism on the plasma boundary

One promising strategy to suppress ELMs is employing RMPs [27] generated by external 3D field coils. The typical setup of external coils encircling the plasma is depicted in Fig. 7(i). These fields effectively reduce the temperature and density at the confinement pedestal, stabilizing the energy bursts in the edge region. This stabilization is crucial as it actively reduces the intense, bursty heat flux and mitigates seed perturbations that could drive additional core instabilities. Consequently, the International Thermonuclear Experimental Reactor (ITER) relies on RMPs to maintain a burst-free burning plasma in a tokamak. However, the detailed mechanism of pedestal degradation by RMPs still remains a challenge. Because the understanding of this physics is vital for the robust projection of the RMP effect on future devices, it needs urgent exploration.

The leading theory [? ?] for explaining the reduced pedestal by RMPs is the formation of magnetic islands by an external 3D field. As illustrated in Figure 7(a),
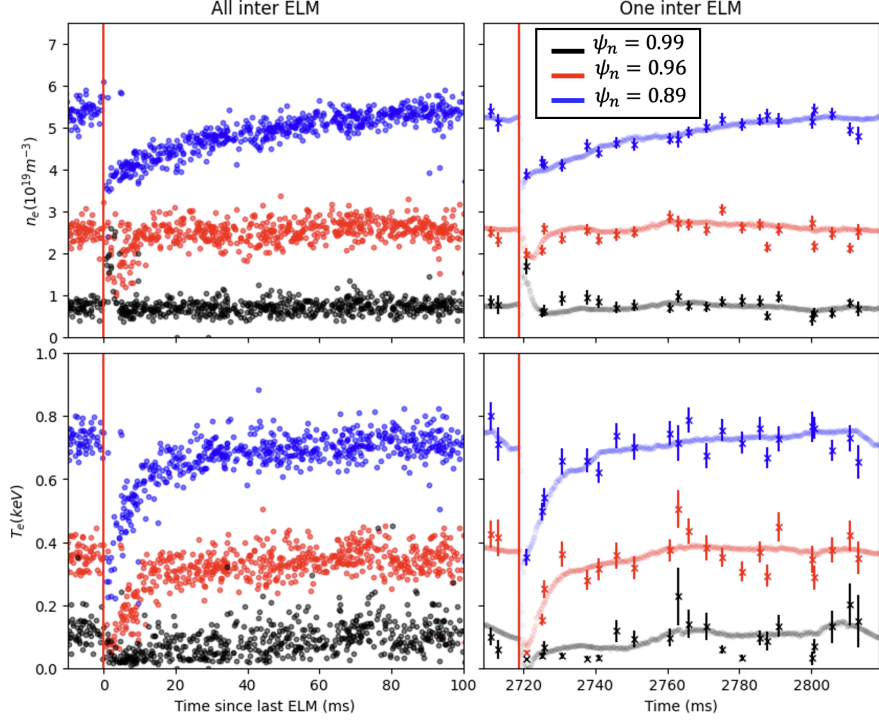
Fig. 6: This figure shows electron density n_e (top plots) and temperature T_e (bottom plots) between from $5ms$ before ELM to $100\,$ms after ELM. The red vertical line represents the time when ELM occurs. The blue points, red points, and black points represent $\psi_n = 0.89, 0.96, 0.99$ respectively. The left two plots show the TS signal accumulated through ELM cycles. The right two plots show the up-sampled TS from only one ELM period (dots). TS measurements (crosses) with error bars (vertical lines) are used to be compared with the up-sampled TS.

RMPs can create these magnetic islands on resonant magnetic surfaces. These islands promote strong radial transport and flatten the local gradients near the island, significantly degrading the density and temperature profiles. Recent modeling efforts [28–30] have shown that these magnetic islands can form both at the top and the foot of the pedestal, contributing to the observed degradation of temperature and density profiles. This model has successfully replicated the experimentally observed pedestal degradation, demonstrating good quantitative accuracy and reinforcing its role as a promising mechanism for RMP-induced pedestal degradation. However, this theory has not yet been fully substantiated, as there is no evidence to confirm the flattening of profiles at islands, which is central to the proposed mechanism. This is due to the spatially narrow island width (derived from theory), the oscillatory behavior of plasma boundary, and source evolution, which makes diagnostics analysis difficult. Here, measuring the perturbative time evolution of the profile by the modulation in RMP amplitude can be an effective way to overcome these difficulties to prove the island effect. However,

11

this was limited due to the insufficient time resolution of TS for achieving statistically reliable profile evolution for such a perturbative transport study.

Interestingly, the super-resolution TS has once again enabled the inspection of rapid profile evolution during perturbative RMP, providing the first experimental evidence of the RMP-induced island effect. Figure 7(b-i) illustrates the recovery of temperature and density pedestals within 10 ms after deactivating the RMP, as captured through numerical modeling Figure 7(b-e) and super-resolution TS Figure 7(f-i). The simulations reveal that the recovery of temperature and density pedestals begins at the top and foot, coinciding with the disappearance of islands at rational surfaces with safety factors q=11/3,13/3, and 14/3. As depicted in Figure 7(d) and Figure 7(e), the profile gradient recovers at these island locations, enhancing the overall profile. This recovery behavior is also mirrored in the super-resolution TS measurements. For instance, the measured electron temperature pedestal shows recovery at both the top and foot through an increasing gradient, displaying qualitative alignment with the simulation results. However, some discrepancies are noted, particularly, in the density evolution at the pedestal foot in the super-resolution TS, even though its gradient remains consistent with the modeling. These quantitative differences may stem from the TS's limited spatial resolution and the modeling assumptions, such as fixed boundary conditions [31]. Nevertheless, the gradient evolution directly indicates a change in transport due to the RMP-induced islands during this perturbative profile evolution, highlights that the super-resolution TS successfully reveals the experimental island effect. This provides the first diagnostic evidence of profile flattening at magnetic islands, a key mechanism of RMP-induced pedestal degradation. This successful application of super-resolution TS underscores its potential to reveal new physics beyond the limitations of conventional diagnostic techniques.

4 Discussion and conclusion

s The findings from our study on the Diag2Diag NN model highlight a significant advance in the field of plasma diagnostics, particularly in the context of temporal resolution enhancement. Our model's capability to utilize multi-modal inputs to reconstruct and enhance the TS diagnostic demonstrates a critical step beyond conventional unimodal interpolation methods, which are typically limited by the resolution of the input data alone.

The synthetic super-resolution TS diagnostic facilitated by our Diag2Diag model has been shown to increase the temporal resolution from a mere 200 Hz to an impressive 500 kHz. This improvement has unlocked new potentials in analyzing fast transient phenomena in plasma, such as the ELMs and the effects of RMPs on pedestal degradation, which were previously blurred or missed in lower resolution data. The ability to inspect these dynamics in greater detail provides new insights into plasma behavior, particularly in conditions where rapid changes occur within milliseconds.

Moreover, the use of high-resolution input diagnostics, like ECE and Interferometer, has not only supported the TS density and temperature profiles but also allowed us to investigate the micro-instabilities and fluctuations within the plasma, which are crucial for understanding the stability and efficiency of confinement in fusion reactors.
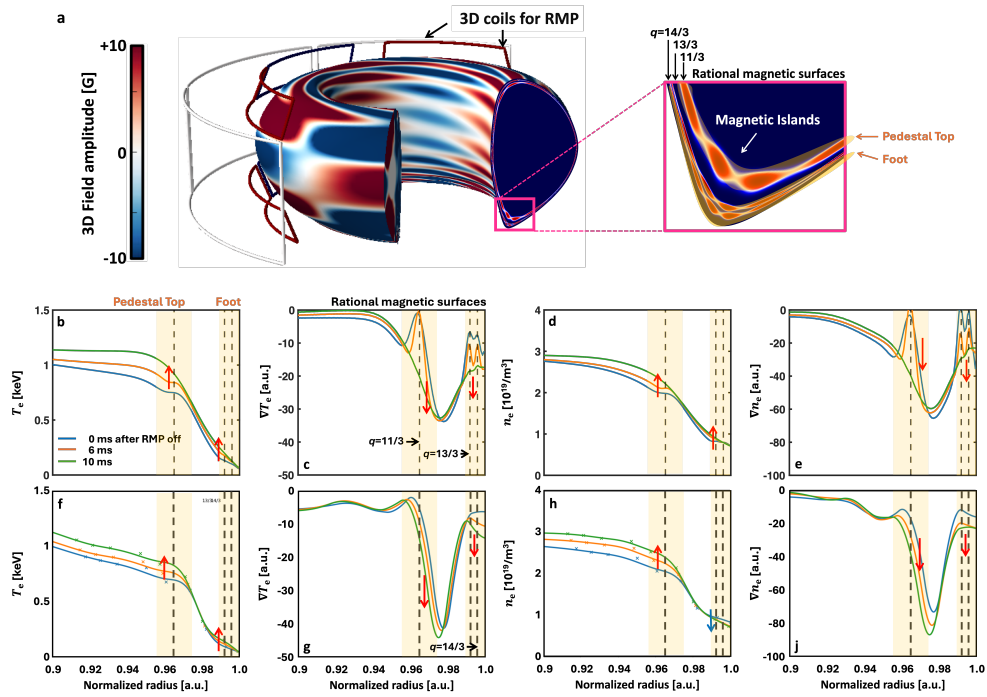
Fig. 7: Structure of 3D coils and islands by perturbed field (a), and the evidence in TM1 simulation (b-e) and super-resolution TS diagnostic (f-i) for RMP-induced island mechanism on the plasma boundary in DIII-D shot number 157545.

The model validation results, with an average R^2 score of 0.92 on the validation set, underscore the model's robustness in accurately reconstructing and enhancing the TS diagnostic.

Despite these advancements, there are limitations to the current approach. The accuracy of the reconstructed signals is dependent on the quality and resolution of the input diagnostics. There is also a need to explore the extrapolation capabilities of the model further to ensure that it remains reliable under different plasma conditions and configurations.

Future work should focus on refining the model to enhance its generalizability across different diagnostic types and experimental setups. Expanding the model's application to other forms of diagnostics could significantly impact the broader field of fusion research, potentially leading to more efficient and stable fusion reactors.

This research not only confirms the efficacy of using AI and neural networks in enhancing diagnostic resolutions but also opens up new avenues for the application of machine learning techniques in plasma physics and other areas requiring high-resolution temporal analysis.

13

5 Methods

5.1 Data Acquisition

For this experiment, we used discharges from the DIII-D tokamak that include all data from the key diagnostics of interest (CER, Interferometer, ECE, MSE, and TS). We randomly selected 4000 discharges recorded between the years 2017 and 2022 to ensure a diverse and representative dataset. The diagnostic data was collected using the DIII-D MDSplus [32] and PTDATA [33] systems. These diagnostics are generally provided as time-series data streams with varying sampling frequencies, ranging from 200 Hz for TS up to 1.66 MHz for Interferometer. The specific pre-processing steps applied to the data for the different experiments conducted in this study are detailed in the following sections.

5.2 Feature extraction

For the spectrogram experiments, we consider the Interferometer and ECE diagnostics. We compute logarithmic magnitude spectrogram from raw series using hamming windows of 1ms with 0.5ms overlap. Given the noisy nature of the ECE signals and after rescaling the spectrograms to the range of [0,1], the spectrograms are enhanced using a pipeline of image processing filters that includes

- Quantile Filtering with a threshold of 0.9,
- Gaussian Blur Filtering on patches of size 31x3,
- Subtracting average per frequency bin
- Morphological filtering on windows of 4x4

We used the ECE spectrograms as inputs to our model. Since we treated every ECE channel independently during feature extraction, we obtained one spectrogram per channel, resulting in 40 input spectrograms (one per ECE channel). Since our model is designed to estimate the Interferometer spectrograms, it predicts four output spectrogram channels corresponding to the four Interferometer interferometer channels.

For the time-series models, the different diagnostic measurements have varying sampling rates, and some are even non-uniformly sampled in time. Since the aim of time-series data analysis was to increase the resolution of TS, we used its timestamps as a reference and aligned all diagnostic modalities to TS by matching their most recent measured samples in time. This resulted in an amount of 135 233 training, 22 084 validation, and 18 721 test samples.

For Interferometer and ECE, we also included the first and second temporal derivatives. Therefore, we smoothed the signals with a moving average window of 1 ms (1660 Interferometer samples and 500 ECE samples), and then computed the first and second temporal derivatives of the smoothed signal also with a window of 1 ms. In this way, we can consider a temporal context of 4 ms.

The diagnostics CER and MSE have a low temporal resolution, i.e., sampling frequencies of 200 Hz and 4 kHz, respectively. In this paper, we assume that they evolve only slowly in time. For the upsampling experiments, we thus pad these diagnostics after a measured sample with constant values until the next measured sample arrived.

The diagnostics (CER, Interferometer, ECE, and MSE) together with the derivatives of Interferometer (4 channels → 12 dimensions including derivatives) and ECE (42 channels → 126 dimensions) lead to an input size of 192. From there, we map to TS with 288 dimensions for plasma density and temperature.

5.3 Spectrogram model development

We developed a CNN model for the task of predicting spectrograms. This model was trained to estimate Interferometer interferometer spectrograms based on input from ECE spectrograms. The optimization process of the model involved several key steps:

- The model underwent training for up to 500 epochs.
- We implemented early stopping with a patience threshold of 20 epochs, during which we monitored the validation loss for any improvements.
- The AdamW optimizer [34], known for decoupling weight decay from the learning rate, was utilized to minimize the $\mathcal{L}1$ loss function.
- We conducted a comprehensive hyper-parameter optimization through a randomized search across 1000 iterations for all hyper-parameters listed in Table 1.

The exact search space of the hyper-parameters and their optimized values obtained from the randomized search are summarized in Table 1.

Table 1: Optimized hyperparameters for the spectrogram prediction CNN model.

Hyper-parameter	Search space	Optimized value
Batch size	1 to 8, random integers	2
Kernel size	3 to 15 odd integers	7
Learning rate	1×10^{-5} to 1 log uniform	0.482×10^{-3}
Final R^2 score	–	0.87

To reduce the amount of training time, we randomly selected 518 discharges from the entire dataset to conduct the hyperparameter optimization. The model with the best performing hyperparameter setting (achieving an R^2 score of 0.87 on the validation set) was then re-trained on all 5000 available discharges.

5.4 Time-series model development

For the time-series prediction task, we employed a MLP model. The input data to the MLP comprised the CER, Interferometer, ECE, MSE, and magnetic diagnostics, along with the first and second temporal derivatives of the Interferometer and ECE signals, resulting in a total input size of 236 dimensions. The target output was the TS diagnostic data, which had 80 dimensions representing electron temperature and density across various spatial locations. The target data were augmented by factor 2 by using the upper and lower intervals of each sample as additional targets.

15

The [MLP](#) model was trained for a maximum of 500 epochs, with an early stopping mechanism implemented to halt the training process if the validation loss did not improve for 20 consecutive epochs. The AdamW optimizer [34] was employed to minimize the $\mathcal{L}1$ loss function during training.

As for the spectrogram model, a comprehensive hyperparameter optimization was undertaken using a randomized search approach spanning 2000 iterations. The hyperparameters jointly optimized included the batch size, hidden layer size, dropout rate, and learning rate.

Table 2 summarizes the optimized hyperparameter values obtained from the randomized search process.

Table 2: Optimized hyperparameters for the time-series [MLP](#) model.

Hyper-parameter	Search space	Optimized value
Batch size	1 to 2048, powers of 2	1024
Hidden layer size	192 to 2048 integers	952
Dropout	0 to 1 uniform	0.076
Learning rate	1×10^{-5} to 1 log uniform	1.998×10^{-3}
Final R^2 score	–	0.92

Disclaimer: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

[1] Seo, J., Kim, S., Jalalvand, A., Conlin, R., Rothstein, A., Abbate, J., Erickson, K., Wai, J., Shousha, R., Kolemen, E.: Avoiding fusion plasma tearing instability

with deep reinforcement learning. Nature **626**(8000), 746–751 (2024) https://doi.org/10.1038/s41586-024-07024-9

[2] Boivin, R.L., Luxon, J.L., Austin, M.E., Brooks, N.H., Burrell, K.H., Doyle, E.J., Fenstermacher, M.E., Gray, D.S., Groth, M., Hsieh, C.-L., Jayakumar, R.J., Lasnier, C.J., Leonard, A.W., McKee, G.R., Moyer, R.A., Rhodes, T.L., Rost, J.C., Rudakov, D.L., Schaffer, M.J., Strait, E.J., Thomas, D.M., Van Zeeland, M., Watkins, J.G., Watson, G.W., Wong, C.P.C.: Diii-d diagnostic systems. Fusion Science and Technology **48**(2), 834–851 (2005) https://doi.org/10.13182/FST05-A1043

[3] Haskey, S.R., Grierson, B.A., Stagner, L., Chrystal, C., Bortolon, A., Laggner, F.M.: Details of the neutral energy distribution and ionization source using spectrally resolved Balmer-alpha measurements on DIII-D. Review of Scientific Instruments **93**(10), 103504 (2022) https://doi.org/10.1063/5.0101854 https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/5.0101854/16594766/103504_1_online.pdf

[4] Strait, E.J.: Magnetic diagnostic system of the DIII-D tokamak. Review of Scientific Instruments **77**(2), 023502 (2006) https://doi.org/10.1063/1.2166493 https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/1.2166493/16053320/023502_1_online.pdf

[5] Holcomb, C.T., Makowski, M.A., Jayakumar, R.J., Allen, S.A., Ellis, R.M., Geer, R., Behne, D., Morris, K.L., Seppala, L.G., Moller, J.M.: Motional Stark effect diagnostic expansion on DIII-D for enhanced current and Er profile measurements. Review of Scientific Instruments **77**(10), 10–506 (2006) https://doi.org/10.1063/1.2235812 https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/1.2235812/11170995/10e506_1_online.pdf

[6] Ponce-Marquez, D.M., Bray, B.D., Deterly, T.M., Liu, C., Eldon, D.: Thomson scattering diagnostic upgrade on DIII-Da). Review of Scientific Instruments **81**(10), 10–525 (2010) https://doi.org/10.1063/1.3495759 https://pubs.aip.org/aip/rsi/article-pdf/doi/10.1063/1.3495759/13981451/10d525_1_online.pdf

[7] HE, Z., SMITH, C., ZHANG, Z., BIEWER, T.M., JIANG, N., HSU, P.S., ROY, S.: Pulse-burst laser-based 10 khz thomson scattering measurements. Plasma Science and Technology **21**(10), 105603 (2019) https://doi.org/10.1088/2058-6272/ab2e30

[8] Jiang, H., Sun, D., Jampani, V., Yang, M.-H., Learned-Miller, E., Kautz, J.: Super slomo: High quality estimation of multiple intermediate frames for video interpolation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)

[9] Niklaus, S., Liu, F.: Context-aware synthesis for video frame interpolation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)

[10] Bao, W., Lai, W.-S., Ma, C., Zhang, X., Gao, Z., Yang, M.-H.: Depth-aware video frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)

[11] Xue, T., Chen, B., Wu, J., Wei, D., Freeman, W.T.: Video enhancement with task-oriented flow. International Journal of Computer Vision **127**(8), 1106–1125 (2019) https://doi.org/10.1007/s11263-018-01144-2

[12] Che, Z., Purushotham, S., Li, G., Jiang, B., Liu, Y.: Hierarchical deep generative models for multi-rate multivariate time series. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 784–793. PMLR, ??? (2018). https://proceedings.mlr.press/v80/che18a.html

[13] Recla, M., Schmitt, M.: Deep-learning-based single-image height reconstruction from very-high-resolution sar intensity data. ISPRS Journal of Photogrammetry and Remote Sensing **183**, 496–509 (2022) https://doi.org/10.1016/j.isprsjprs.2021.11.012

[14] Recla, M., Schmitt, M.: Improving deep learning-based height estimation from single sar images by injecting sensor parameters. In: IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium, pp. 1806–1809 (2023). https://doi.org/10.1109/IGARSS52108.2023.10282228

[15] Roßberg, T., Schmitt, M.: Temporal upsampling of ndvi time series by rnn-based fusion of sparse optical and dense sar-derived ndvi data. In: IGARSS 2023 - 2023 IEEE International Geoscience and Remote Sensing Symposium, pp. 5990–5993 (2023). https://doi.org/10.1109/IGARSS52108.2023.10282861

[16] Bellos, D., Basham, M., Pridmore, T., French, A.P.: A convolutional neural network for fast upsampling of undersampled tomograms in x-ray ct time-series using a representative highly sampled tomogram. Journal of Synchrotron Radiation **26**(3), 839–853 (2019) https://doi.org/10.1107/S1600577519003448

[17] Pascual, S., Bonafonte, A., Serrà, J.: Segan: Speech enhancement generative adversarial network. In: Proceedings of the Interspeech 2017, pp. 3642–3646 (2017). https://doi.org/10.21437/Interspeech.2017-1428

[18] Donahue, C., McAuley, J., Puckette, M.: Adversarial audio synthesis. In: International Conference on Learning Representations (2019). https://openreview.net/forum?id=ByMVTsR5KQ

[19] Pandey, A., Wang, D.: Densely connected neural network with dilated convolutions for real-time speech enhancement in the time domain. In: ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6629–6633 (2020). https://doi.org/10.1109/ICASSP40776.2020.9054536

[20] Kumar, K., Kumar, R., Boissiere, T., Gestin, L., Teoh, W.Z., Sotelo, J., Brébisson, A., Bengio, Y., Courville, A.C.: Melgan: Generative adversarial networks for conditional waveform synthesis. In: Wallach, H., Larochelle, H., Beygelzimer, A., Alché-Buc, F., Fox, E., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 32. Curran Associates, Inc., ??? (2019). https://proceedings.neurips.cc/paper_files/paper/2019/file/6804c9bca0a615bdb9374d00a9fcba59-Paper.pdf

[21] Yoon, J., Zame, W.R., Schaar, M.: Estimating missing data in temporal data streams using multi-directional recurrent neural networks. IEEE Transactions on Biomedical Engineering **66**(5), 1477–1490 (2019) https://doi.org/10.1109/TBME.2018.2874712

[22] Li, Y., Wang, Y., Meng, C., Duan, Y., Ji, J., Zhang, Y., Zhang, Y.: Farfusion: A practical roadside radar-camera fusion system for far-range perception. IEEE Robotics and Automation Letters, 1–8 (2024) https://doi.org/10.1109/LRA.2024.3387700

[23] Melis, J.M., Siwanowicz, I., Dickinson, M.H.: Machine learning reveals the control mechanics of an insect wing hinge. Nature **628**(8009), 795–803 (2024) https://doi.org/10.1038/s41586-024-07293-4

[24] Garcia, A.V., Jalalvand, A., Steiner, P., Rothstein, A., Van Zeeland, M., Heidbrink, W.W., Kolemen, E.: Comparison of machine learning systems trained to detect alfvén eigenmodes using the co2 interferometer on diii-d. Nuclear Fusion **63**(12), 126039 (2023) https://doi.org/10.1088/1741-4326/acfe8b

[25] Nelson, A.O., Xing, Z.A., Izacard, O., Laggner, F.M., Kolemen, E.: Interpretative sol modeling throughout multiple elm cycles in diii-d. Nuclear Materials and Energy **26**, 100883 (2021) https://doi.org/10.1016/j.nme.2020.100883

[26] Nazikian, R., Hu, Q., Ashourvan, A., Eldon, D., Evans, T.E., Grierson, B.A., Logan, N.C., Orlov, D.M., Park, J.-K., Paz-Soldan, C., Poli, F.M., Yu, Q.: Pedestal collapse by resonant magnetic perturbations. Nuclear Fusion **61**(4), 044001 (2021) https://doi.org/10.1088/1741-4326/abd863

[27] Evans, T.E., Moyer, R.A., Burrell, K.H., Fenstermacher, M.E., Joseph, I., Leonard, A.W., Osborne, T.H., Porter, G.D., Schaffer, M.J., Snyder, P.B., Thomas, P.R., Watkins, J.G., West, W.P.: Edge stability and transport control with resonant magnetic perturbations in collisionless tokamak plasmas. Nature Physics, vol. 2, no. 6, June 1, 2006, pp. 419-423 **2**(6) (2006) https://doi.org/10.1038/nphys312

[28] Hu, Q.M., Nazikian, R., Grierson, B.A., Logan, N.C., Park, J.-K., Paz-Soldan, C., Yu, Q.: The density dependence of edge-localized-mode suppression and pump-out by resonant magnetic perturbations in the DIII-D tokamak. Physics of Plasmas **26**(12), 120702 (2019) https://doi.org/10.1063/1.5134767 https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/1.5134767/13931866/120702_1_online.pdf

[29] Hu, Q.M., Nazikian, R., Grierson, B.A., Logan, N.C., Orlov, D.M., Paz-Soldan, C., Yu, Q.: Wide operational windows of edge-localized mode suppression by resonant magnetic perturbations in the diii-d tokamak. Phys. Rev. Lett. **125**, 045001 (2020) https://doi.org/10.1103/PhysRevLett.125.045001

[30] Hu, Q.M., Park, J.-K., Logan, N.C., Yang, S.M., Grierson, B.A., Nazikian, R., Yu, Q.: Nonlinear two-fluid modeling of plasma response to rmps for the elm control in the iter baseline. Nuclear Fusion **61**(10), 106006 (2021) https://doi.org/10.1088/1741-4326/ac1d85

[31] Yu, Q., Strumberger, E., Igochine, V., Lackner, K., Laqua, H.P., Zanini, M., Braune, H., Hirsch, M., Höfel, U., Marsen, S., Stange, T., Wolf, R.C., Günter, S., Wendelstein 7-X Team: Numerical modeling of the electron temperature crashes observed in wendelstein 7-x stellarator experiments. Nuclear Fusion **60**(7), 076024 (2020) https://doi.org/10.1088/1741-4326/ab9258

[32] Fredian, T.W., Stillerman, J.A.: Mdsplus. current developments and future directions. Fusion Engineering and Design **60**(3), 229–233 (2002) https://doi.org/10.1016/S0920-3796(02)00013-3

[33] Schissel, D., McHarg Jr, B.: Data analysis infrastructure at the diii-d national fusion facility. General Atomics report No. GA-A23474 (2000)

[34] Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2019). https://openreview.net/forum?id=Bkg6RiCqY7