Alfvén eigenmode detection using Long-Short Term Memory Networks and CO_2 Interferometer data on the DIII-D National Fusion Facility

1st Alvin V. Garcia
Physics and Astronomy
University of California, Irvine
Irvine, CA, 92697, USA
alvin.garcia@uci.edu

2nd Azarakhsh Jalalvand
Mechanical and Aerospace Engineering
Princeton University
Princeton NJ 08544, USA
azarakhsh.jalalvand@princeton.edu

3rd Peter Steiner
Institute of Acoustics
and Speech Communication
Technische Universität Dresden
Dresden, Germany
peter.steiner@tu-dresden.de

4th Andrew Rothstein
Mechanical and Aerospace Engineering
Princeton University
Princeton NJ 08544, USA
arothstein@princeton.edu

5th Michael Van Zeeland
Experimental Science
General Atomics
San Diego, CA, 92121, USA
vanzeeland@fusion.gat.com

6th William W. Heidbrink
Physics and Astronomy
University of California, Irvine
Irvine, CA, 92697, USA
bill.heidbrink@uci.edu

7th Egemen Kolemen
Mechanical and Aerospace Engineering
Princeton University
Princeton NJ 08544, USA
ekolemen@princeton.edu

2023 International Joint Conference on Neural Networks (IJCNN) | 978-1-6654-8867-9/23/$31.00 ©2023 IEEE | DOI: 10.1109/IJCNN54540.2023.10191760

***Abstract*—The successful steady-state operation of burning fusion plasmas in planned future devices such as the ITER tokamak requires understanding of fast-ion physics. Alfvén eigenmodes are special cases of plasma waves driven by fast ions that are important to identify and control since they can lead to loss of confinement and potential damage to the inner walls of a plasma device. The goal of this work is to compare machine learning-based systems trained to classify Alfvén eigenmodes using CO_2 interferometer data from a labelled database on the DIII-D tokamak. A Long-Short Term Memory (LSTM) network is trained from scratch using simple spectrogram representations of the CO_2 phase data. The model is trained using a single chord (sequence) per training step. Results show a total true positive rate of = 90% and a false positive rate of = 18%. This paper demonstrates the potential of applying machine learning models to detect and identify different classes of Alfvén eigenmodes for real-time applications in steady-state plasma operations that could potentially drive actuators to mitigate Alfvén eigenmode impacts.**

***Index Terms*—Fusion Energy, Machine Learning Classification, Alfvén Eigenmodes, CO_2 Interferometry, DIII-D Tokamak**

I. Introduction

The success of nuclear fusion devices depends on well behaved populations of super-thermal particles that slow down in the bulk of the plasma [1]. Well confined alpha particles born from fusion reactions could be recycled and provide enough heat to sustain an ignited burning plasma. Unconfined fusion born alpha particles can carry away fusion power to the first walls and degrade the performance of the plasma [2]. External heating systems such as neutral beam injection (NBI) or radio frequency (RF) waves can produce populations of energetic particles (EP) useful for momentum transfer and current drive [3]. However, fusion born and externally driven EPs can resonate with special cases of plasma waves called Alfvén eigenmodes (AEs), exacerbate unstable plasma conditions, lead to a quench of reactions and degrade fusion performance.

Detection and control of AEs during real-time experiments is important for the realization of fusion energy. The International Thermonuclear Experimental Reactor (ITER) is the largest fusion device in the world and planned to generate more power than the heat required to produce fusion reactions [4]. AEs are commonly observed in currently operational fusion devices [5], [6], and are predicted to occur in ITER [7]–[12]. It is a mission critical goal for the Plasma Control System (PCS) at ITER to detect and control AEs in high performance plasmas that maintain fusion burn and mitigate damage to the inner walls of the machine vessel [13], [14]. A considerable effort in the fusion community is to determine the best set of external actuators that would control AE instabilities and alpha losses [15]. Currently, NBI, electron cyclotron resonance heating or current drive are three of the most promising techniques. Predictions from data-driven models can provide more insights in this area and pave the path towards ITER AE control.

Machine Learning (ML) models trained using experimental

data can be implemented into the PCS at fusion devices for the optimization of EP self-heating, current drive efficiency and mitigation of first wall impacts from losses. There is a need in the community for models with quick response times (milliseconds) that could accurately detect the presence of AEs in real-time experiments [16]. Large sets of experimental data are routinely collected at the DIII-D National Fusion Facility that are suitable for machine learning-based analysis techniques. Many operational regimes, scenarios and parameter spaces can be investigated with this ML-based assistance. Also, these models can be initially trained offline on large amounts of data and later get adapted to the new data offline or online. An example appears in [16].

Machine learning applied to challenges in fusion energy is an emerging area of research. One of the most developed fields using ML methods is disruption mitigation [17]–[24], where models are trained to predict disruptions to the plasma that could be damaging to the walls of the vessel. However, ML applications in EP physics offers exciting opportunities with only a few existing studies to date. For example, deep learning networks are trained to automatically detect magneto-hydrodynamic (MHD) and AE activity using manually-labeled targets and magnetics data from the TJ-II stellarator [25] and COMPASS tokamak [26]. Also, recent work detected AEs using supervised learning techniques [25]–[27], and data mining techniques plasma fluctuation extraction using clustering and time-series for event identification [28], [29].

There has been significant progress in the area of classifying AEs using machine learning-based models and high resolution electron-cyclotron emission (ECE) data [30], [31]. Previous work used reservoir computing networks (RCN) to classify AEs observed at DIII-D using the large labeled database discussed in more detail in subsection II-B. This state-of-the-art (SOTA) model was trained using the AE-EP database to classify and predict AEs using time-series from the 40-channel ECE diagnostic on the DIII-D tokamak [30]. Following the RCN work, a manually labeled database of 26 discharges based on the original large DIII-D database was used to train simple neural-network-based models for spatiotemporal local identification of four AEs [31].

The goal here is to build on the prior work presented in [30], train two new machine learning architectures and compare their performance. Our prior work used ECE data and we instead use CO_2 interferometer data to train machine learning models. Also, we compute FOURIER transforms for the inputs unlike in our prior work, which used the 1D signals for training. In [30], Reservoir Computing Networks are trained to detect AEs and we instead train a Linear Regression (LR) model and a Long-Short Term Memory (LSTM) Network separately from scratch. We compare the results of these new models and aim to achieve performance values seen in our prior work. Both the LR and LSTM model use only one CO_2 chord (sequence) per step during the training process, which is also different in [30] since all 40 ECE channels are stacked and used as input for the RCN model.

This paper is organized as follows: The CO_2 interferometer system at DIII-D, large AE-EP database and interesting challenges are discussed in section II. The classification and performance of the LR and LSTM models are reported in section III. Conclusions are drawn in section IV.

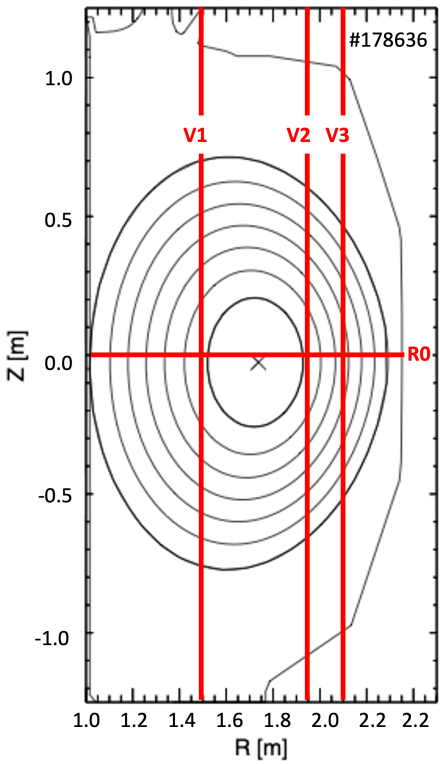

Fig. 1. A layout of the CO_2 interferometer at DIII-D. Vertical chords V1, V2 and V3 are located at major radii R_m of $1.48\,\mathrm{m}$, $1.94\,\mathrm{m}$ and $2.10\,\mathrm{m}$, respectively. The horizontal chord R0 is located at the midplane ($Z = 0\,\mathrm{m}$). The black "ovals" are contours of constant magnetic flux with the outermost representing the last closed magnetic flux surface, and the $\times$ symbol is the magnetic axis.

II. Experimental Setup

A. Multi-chord CO_2 Interferometer

The two-color vibration compensated CO_2 interferometer system installed on the DIII-D tokamak consists of three vertical and one horizontal chord. Figure 1 shows a layout of the diagnostic system. The full 3D toroidally shaped plasma is described by rotating this cross-section 360° about a vertical centerline axis located at $R_m = 0$ m (donut-shaped vessel). Each of these chords makes line integrated measurements of the electron density in the plasma, and are digitized for 9 seconds per discharge at a rate of 1.67 MS/s. Several signals are available and we use the CO_2 phase data in this work since AEs are well above typical mechanical vibration frequencies and it can be acquired by the Plasma Control System (PCS) in real time for future actuator driven mitigation of AE impacts.

Spectral analysis of the CO_2 interferometer diagnostic time series data is commonly done in the energetic particle community to detect AEs [32]. Computing spectrograms removes low frequency noise and machine vibrations observed in the data. Coherent activity for a given discharge can be studied

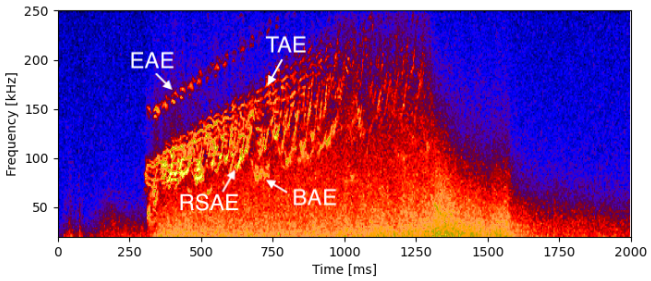

Fig. 2. An example crosspower spectrogram between chords V2 and R0 for shot 170677. Four plasma instabilities (EAE, TAE, RSAE and BAE) are labelled by visual inspection.

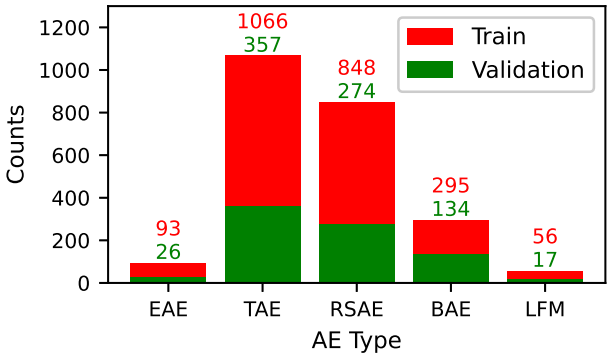

Fig. 3. Class distribution for expert-labeled AEs in the CO_2 interferometer dataset. The entire dataset is randomly shuffled into 801 training and 268 validation samples. There are clearly more RSAE and TAE labels and relatively few instances for EAE, BAE and LFM, making this classification problem highly imbalanced.

by calculating windowed crosspower spectrograms between two chords. An example is shown in Figure 2. Post-shot identification of AEs using crosspower spectra is possible, but is time intensive and requires extensive domain knowledge. In this work, we instead make identification instantaneous and simple by training machine learning models using simple magnitude spectrograms of CO_2 interferometer data.

B. The Large 2009–2017 DIII-D AE-EP Database

Recent work produced a database that is suitable for machine learning analysis. It includes the occurrences of six plasma instabilities: ellipticity (EAE), toroidal (TAE), reversed-shear (RSAE), beta-induced (BAE), low-frequency mode (LFM), and energetic particle-induced geodesic acoustic mode (EGAM) [33] . These flags sample a variety of plasma conditions or mode activity, and occur during the first $1.9\,\mathrm{s}$ due to the q profile steadily evolving during that phase of the discharge. Since flags are timestamps and originally motivated by physics analysis, there are several challenges for ML classification that are addressed in the remainder of this section.

After adapting the original nomenclature of Heidbrink et al., the database was discovered to have sparse temporal representation for LFM and EAE. In an effort to detect modes with flags having the highest degree of confidence, only stable (0) and unstable (2) flags are considered and made binary (one-hot encoded) as described in [30]. With these changes, labels 0 and 1 indicate the AE is *not present* or *present*, respectively. Due to the sparsity and campaign specific representation of LFM and EAE, random splitting is preferred over chronological splitting of the training and validation set. Figure 3 shows the CO_2 interferometer training and validation class distribution for 1069 available shots. There is a large imbalance favoring TAE and RSAE since they occur frequently in many experimental campaigns.

Since the labels are single timestamps for each AE at semi-arbitrary points throughout the database, temporal widening of the labels is necessary for accurate ML predictions. Windows of $250\,\mathrm{ms}$ centered about the timestamp approximates the time duration of each AE. Post-processing the expert marked flags by interpolating over a window of $250\,\mathrm{ms}$ adequately prepares the proper labels for ML training.

Reporting accuracy would be misleading since a model outputting all zeros would be higher than $94\,\%$ accurate and still miss the vast majority of AE activity. Thus, the metrics of success for this work are the following: True Positive Rate, $\mathrm{TPR} = \mathrm{TP}/(\mathrm{TP}+\mathrm{FN})$, and False Positive Rate, $\mathrm{FPR} = \mathrm{FP}/(\mathrm{FP}+\mathrm{TN})$. TP and TN are outcomes where the ML model *correctly* predicts the positive and negative label. Conversely, FP and FN are outcomes where the ML model *incorrectly* predicts the positive and negative label.

Given all of the challenges and initial inspection of the data, this project is framed as a class imbalanced multilabel multi-class classification problem. Although these challenges similarly applied to the prior work, excellent classification performance was achieved, with a true positive rate of $91\,\%$ and a false positive rate of $7\,\%$, see Table III of [30]. In this work, the aim is to match these results using new data (CO_2 interferometer) and new models (LR and LSTM) for the potential future application of Multi-Modal Machine Learning techniques in real-time control algorithms.

III. Classification and Performance

An exploratory approach is implemented and ML models are trained using CO_2 interferometer data to detect AEs. First, we train a Linear Regression model for a linear baseline technique. Then, we train a Recurrent Neural Network with Long-Short Term Memory cells for comparison. After computing spectrograms, we address two major goals in this work:

1) Compare different machine learning models, i.e., LR and LSTM architectures. The different models are introduced in subsection III-B and subsection III-C, respectively.
2) Determine if training using a single CO_2 chord per training step is sufficient to detect AEs, and which chord shows the best performance. This and the overall results are shown in subsection III-D.

A. Inputs

Simple magnitude spectrograms are prepared as input for the LR and LSTM models. Starting point for the feature set is

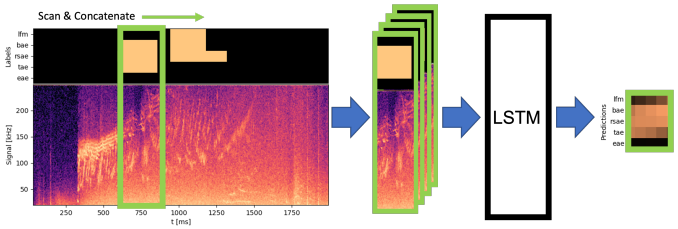

Fig. 4. Workflow of the "windowed" LSTM training process. Successive windows of $\sim 280\,\mathrm{ms}$ and $75\,\%$ overlap are concatenated and inserted into the LSTM during the training process. Since this overlapping method creates bins of $\sim 70\,\mathrm{ms}$ that will intersect several windows for the predictions of a given discharge, the predictions are collected such that the average AE score for all intersecting windows is deposited into this smaller bin.

the short-term FOURIER transform. Therefore, the input signal stream is divided into overlapping frames, and each frame has a spectrum computed. The final magnitude spectrogram is obtained by computing the absolute values of all spectra. Based on preliminary studies, the magnitude spectrograms are computed with a window length of $4.9\,\mathrm{ms}$ and an overlap of $80\,\%$. Also, the spectrograms use linear detrending and hanning windows. Finally, all spectrograms are downsampled by applying a maxpool function (8,4) to local frequency-time blocks in both directions making the final input shapes = (time, frequency) = (142, 508). Maxpooling is common in computer vision tasks and works well in this study.

Although both models "see" identical simple spectrograms, each model processes them differently. For a given discharge, the LSTM will process both spectrograms in windows of $\sim 280\,\mathrm{ms}$ with $75\,\%$ overlap per training step. These settings are close to the original label length of $250\,\mathrm{ms}$ and produce predictions over the same window. Since there are many overlapping windows for a given discharge due to the successive concatenation, averaging over smaller sub-windows $\sim 70\,\mathrm{ms}$ is necessary to collect the outputs. Figure 4 shows a workflow of this procedure. On the other hand, the LR model uses windowing in a slightly different manner. Although the model processes 1D vectors of frequencies per training step, frame stacking of nearby timestamps is implemented to provide the LR model temporal information about the past and future ($\pm 125\,\mathrm{ms}$) AE activity.

B. Linear Regression

Linear regression is the method chosen here since it provides a baseline for Reservoir Computing Networks, where the output weights are trained using linear regression. RCN classification was the method used in our prior work [30] to detect AEs using ECE data. An important advantage of using linear regression compared to other methods such as logistic regression is the online adaption in real-time control.

The Linear Regression model in this work is the conventional technique using Tikhonov regularization. The inputs are the spectrograms, and the model maps the spectral information to an AE score. All discharges are concatenated into a matrix $\mathbf{X}$ and a column of ones are added for the linear term. The labels are similarly collected into a matrix $\mathbf{y}$. Also, frame stacking over $250\,\mathrm{ms}$ windows is implemented at this point to provide the LR model with nearby temporal information during the training process. The mapping matrix is defined as follows:

$$\mathbf{W}_{\mathrm{out}} = \left(\mathbf{X}^T\mathbf{X} + \alpha\mathbf{I}\right)^{-1}\mathbf{X}^T\mathbf{y}, \tag{1}$$

where α is the regularization parameter. Since the sequences are linearized (discharge concatenation step), building $\mathbf{W}_{\mathrm{out}}$ becomes "1-shot" training and the process is very quick, particularly since the terms $\mathbf{X}^T\mathbf{X}$ and $\mathbf{X}^T\mathbf{y}$ in Equation (1) can be computed incrementally, as shown in [34].

TABLE I
TOTAL RATES FOR BOTH MODELS USING A THRESHOLD OF 0.1.

	Linear Regression	LSTM
TPR	0.86	0.90
FPR	0.19	0.18

C. LSTM

The base of the architecture consists of a LSTM block of three layers using 64 nodes per layer and relu activation function. The weights of each layer are initially set using uniform variance scaling [35]. The output of the LSTM is a vectorized feature map of length 64 and is input directly into an MLP block at the head of the network consisting of one 50% dropout layer followed by three layers with 128 nodes per layer using relu activation and a final classification layer with 5 nodes using sigmoid activation. The LSTM is trained using the Binary Crossentropy Loss Function and Adam Optimizer with a learning rate of 1×10^{-4}. The LSTM block considers memory in the inputs using recurrent connections, and the MLP functions as a fully connected feed-forward neural network. The hyperparameters were optimized by scanning values and observing performance on three selected discharges.

D. Results

The final classification results are summarized in Table I and show high performance. The LSTM rates have similar performance to the SOTA technique from the prior work [30]. The LSTM model has higher TPR and lower FPR since it can better identify the least common modes. Figure 5 shows the Linear Regression predictions over the entire validation set for LFM. Frame stacking was not enough to produce strong predictions and there is no discernible aggregate shape.

The F_2 score is a harmonic mean of the precision and recall metrics, where $\beta = 2$ in the following equation:

$$\mathrm{F}_\beta = \frac{1+\beta^2}{\frac{\beta^2}{Recall} + \frac{1}{Precision}}. \tag{2}$$

The F_2 score gives more weighting to TPR, and this metric is computed for both models across each chord in Figure 6. The LSTM model performs better than the LR model and

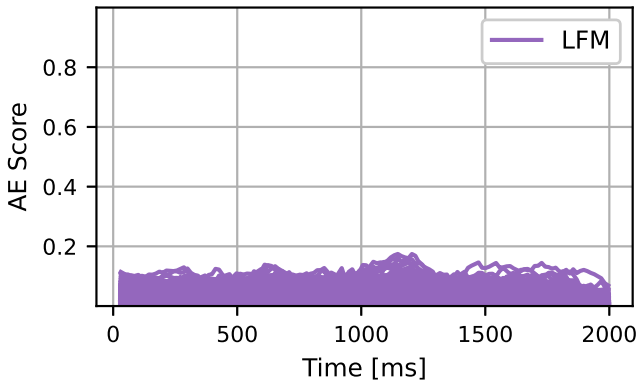

Fig. 5. Predictions for LFM using the Linear Regression model are low and don't provide much information. Recurrent connections are necessary for detecting this mode.

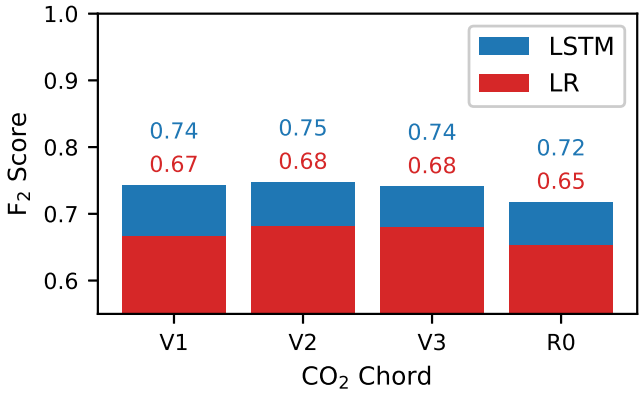

Fig. 6. F2 scores for the Recurrent Neural Network with Long-Short Term Memory cells and Linear Regression Model show the former performs consistently better. Also, predictions on chord V2 are slightly better than the other three chords.

chord V2 has slightly higher performance than the other three chords. Deep layers with recurrent connections are necessary to capture the patterns of the least common modes. Therefore, the remainder of this section will focus on the LSTM model and its predictions.

Only a few hyperparameters had a noticeable effect on the performance of the LSTM model. Adding a third LSTM layer and increasing the number of LSTM nodes to 64 improved the predictions for LFM. With less model capacity, the network was incapable of triggering and output mostly 0 for this mode. Improved performance for LFM with the addition of layers and nodes is expected since LFMs have a unique "christmas light" pattern [33] that would be more sensitive to networks with increased temporal memory. Changes for hyperparameters in the MLP block (layers, nodes or dropout) or further increasing the capacity of the LSTM block had little effect on the performance.

The TPR and FPR curves for the LSTM model are displayed in Figure 7. TAE and RSAE have the best classification performance since these modes occur most frequently throughout the database and are relatively easier to detect. TAE and RSAE often occur in the middle of the frequency range and for a long time in many discharges. Since these modes occur frequently, they will have the largest impact on the total TPR and FPR. On the other hand, predictions for EAE, BAE and LFM are expected to be lower since they do not appear as much in the database. Nevertheless, the LSTM is able to capture their patterns by taking advantage of memory in the recurrent connections.

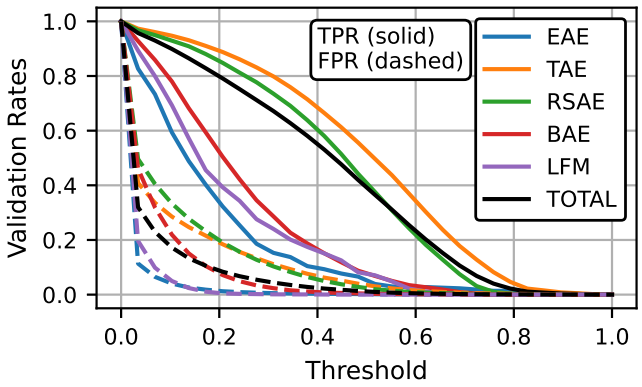

Fig. 7. TPR (solid) and FPR (dashed) curves for the LSTM model.

The relatively larger FPR for RSAE and TAE indicate that the LSTM model slightly overestimate detection for these modes. Although effects from the quality of the training data and overfitting are likely to contribute, the partially flagged AE activity in the original database are likely to have a substantial effect. Each model will likely identify AEs that are marked *not present*, but will really have AE activity in the selected discharges, see Figure 8. Also, unflagged regions could be areas where the interpolated training window of 250 ms was not large enough to capture the full domain of the AE activity. Figure 9 shows an example of this behavior. Although these issues influence the overall performance, clearly the LSTM model is still capable of learning the patterns associated with the AE activity.

Shot 178636 and 175985 are interesting examples, where the LSTM model performs very well for the least common modes. Although shot 178636 is difficult to classify since there is a lot of AE activity, the LSTM is capable of identifying all four modes, see Figure 10. Predictions align well with the provided labels. Shot 175985 is another example of a difficult to predict discharge, only this one contains EAE. Although predictions are light for this mode, the model identifies the *presence* of EAE, see Figure 11. For both of these shots, it was important to build models with sufficient memory capabilities to capture temporal information.

IV. Conclusion

A Linear Regression Model and a Recurrent Neural Network with Long-Short Term Memory cells are trained us-

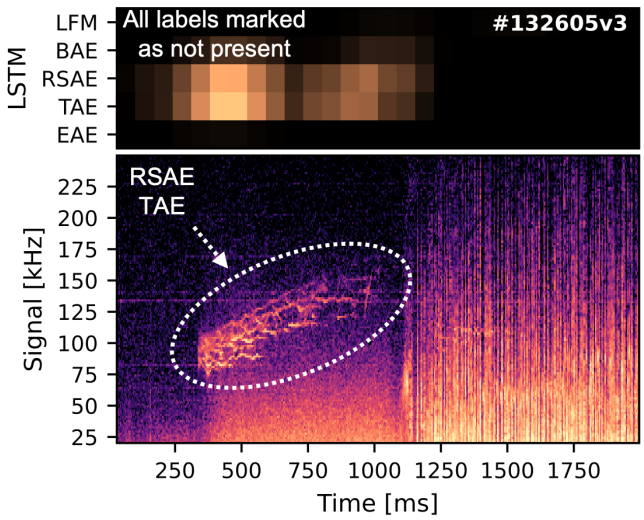

Fig. 8. An interesting example where the LSTM model is robust enough to capture activity for TAE and RSAE in shot #132605, where labels indicate all modes are not present. Orange pixels show that the predictions are able to detect the AE activity observed by human inspection in the dotted oval.

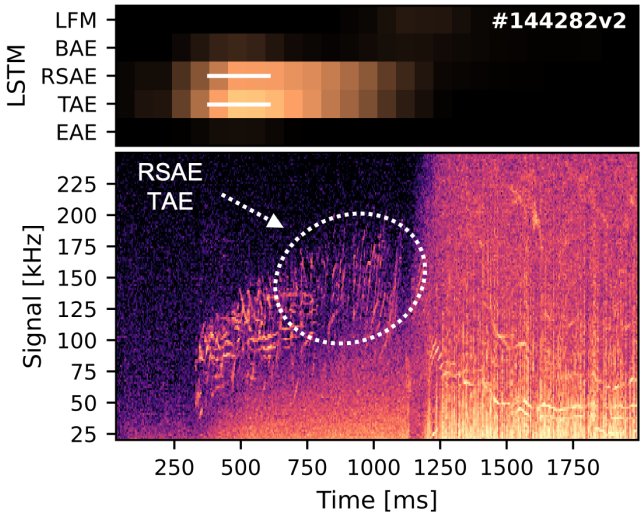

Fig. 9. LSTM predictions using simple spectrograms as input for discharge #144282. White lines indicate the "ground truth" areas, which is a 250 ms window centered about the original expert made timestamp (in this case the timestamp is located at 500 ms). The LSTM is able to predict AE activity in regions where the labels are marked *not present*.

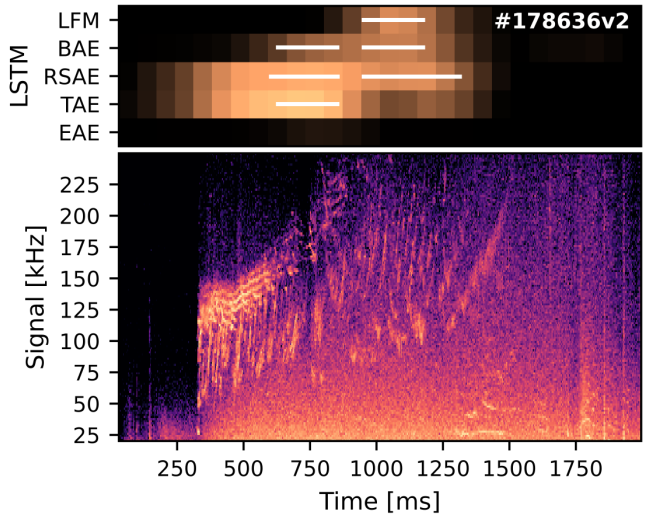

Fig. 10. LSTM predictions using simple spectrograms as input for shot #178636. "Ground-truth" labels are denoted by the white strikethroughs. The model is able to correctly identify the least common mode, LFM, and the rest of the AEs. Also, the model is robust enough to make good predictions for times before 620 ms, where the curated database doesn't say anything about the AE activity.

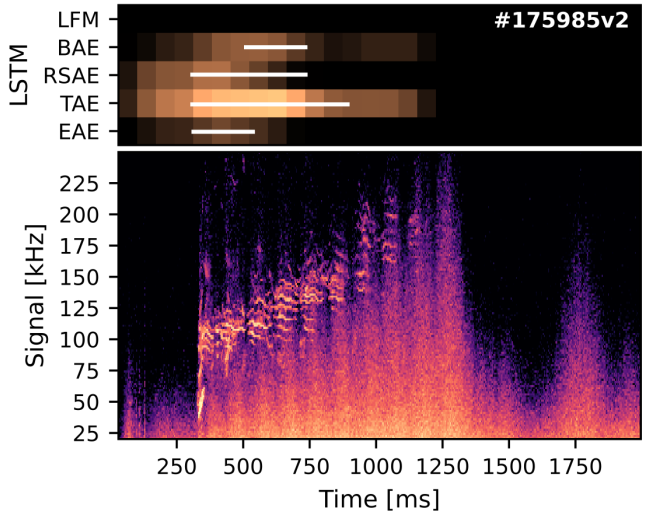

Fig. 11. Identical input and model setup as Figure 10, only predictions are for shot #175985. Although predictions are light, this example shows good agreement for the second least common mode, EAE, and all other modes.

ing CO_2 interferometer data and labels from the large AE-EP database on the DIII-D tokamak. The inputs are simple magnitude spectrograms, and both models use one CO_2 interferometer chord per step during training. Predictions are made binary by thresholding the outputs for each mode. For the LSTM model, the total TPR $= 90\,\%$ and FPR $= 18\,\%$ using a threshold value of 0.1 (Linear Regression also has high performance). Predictions are most accurate for CO_2 interferometer chord V2. These results are in close agreement of the state-of-the-art technique in prior work, and show detection using only one chord is possible. This work is useful for the implementation of ML systems into real-time algorithms on the Plasma Control System for the detection and control of AEs in upcoming experiments at DIII-D. It would also be interesting to perform cross-machine classification using spectrograms from a different device in future work.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, Office of Workforce Development for Teachers

and Scientists, Office of Science Graduate Student Research (SCGSR) program, using the DIII-D National Fusion Facility, a DOE Office of Science user facility, under Award(s) DE-FC02-04ER54698, DE-SC0021275, DE-SC0020337, DE-SC0014664, Army Research Office (ARO W911NF-19-1-0045), National Science Foundation under 1633631 and Ghent University Special Research Award No. BOF19/PDO/134. The SCGSR program is administered by the Oak Ridge Institute for Science and Education (ORISE) and managed by ORAU for the DOE.

REFERENCES

[1] A. Fasoli, C. Gormenzano, H. Berk, *et al.*, "Physics of energetic ions," *Nuclear Fusion*, vol. 47, no. 6, S264, 2007.

[2] H. Duong, W. Heidbrink, E. Strait, *et al.*, "Loss of energetic beam ions during TAE instabilities," *Nuclear Fusion*, vol. 33, no. 5, p. 749, 1993.

[3] N. N. Gorelenkov, S. Pinches, and K. Toi, "Energetic particle physics in fusion research in preparation for burning plasma experiments," *Nuclear Fusion*, vol. 54, no. 12, p. 125 001, 2014.

[4] *ITER*, https://www.iter.org/sci/goals, 2023.

[5] K.-L. Wong, "A review of Alfvén eigenmode observations in toroidal plasmas," *Plasma physics and controlled fusion*, vol. 41, no. 1, R1, 1999.

[6] W. Heidbrink, "Basic physics of Alfvén instabilities driven by energetic particles in toroidally confined plasmas," *Physics of Plasmas*, vol. 15, no. 5, p. 055 501, 2008.

[7] H. Yamada, S. Murakami, K. Yamazaki, *et al.*, "Impact of heat deposition profile on global confinement of NBI heated plasmas in the LHD," *Nuclear fusion*, vol. 43, no. 8, p. 749, 2003.

[8] G. Vlad, S. Briguglio, G. Fogaccia, F. Zonca, and M. Schneider, "Alfvénic instabilities driven by fusion generated alpha particles in ITER scenarios," *Nuclear Fusion*, vol. 46, no. 1, p. 1, 2005.

[9] M. Van Zeeland, N. Gorelenkov, W. Heidbrink, *et al.*, "Alfvén eigenmode stability and fast ion loss in DIII-D and ITER reversed magnetic shear plasmas," *Nuclear Fusion*, vol. 52, no. 9, p. 094 023, 2012.

[10] P. Rodrigues, A. Figueiredo, J. Ferreira, *et al.*, "Systematic linear-stability assessment of Alfvén eigenmodes in the presence of fusion α-particles for ITER-like equilibria," *Nuclear Fusion*, vol. 55, no. 8, p. 083 003, 2015.

[11] T. Hayward-Schneider, P. Lauber, A. Bottino, and Z. Lu, "Global linear and nonlinear gyrokinetic modelling of Alfvén eigenmodes in ITER," *Nuclear Fusion*, vol. 61, no. 3, p. 036 045, 2021.

[12] E. Bass and R. Waltz, "Prediction of Alfvén eigenmode energetic particle transport in ITER scenarios with a critical gradient model," *Nuclear Fusion*, vol. 60, no. 1, p. 016 032, 2019.

[13] J. Snipes, D. Campbell, T. Casper, *et al.*, "MHD and plasma control in ITER," *Fusion Science and Technology*, vol. 59, no. 3, pp. 427–439, 2011.

[14] J. Snipes, S. Bremond, D. Campbell, *et al.*, "Physics of the conceptual design of the ITER plasma control system," *Fusion Engineering and Design*, vol. 89, no. 5, pp. 507–511, 2014.

[15] M. Garcia-Munoz, S. Sharapov, M. Van Zeeland, *et al.*, "Active control of Alfvén eigenmodes in magnetically confined toroidal plasmas," *Plasma Physics and Controlled Fusion*, vol. 61, no. 5, p. 054 007, 2019.

[16] A. Jalalvand, J. Abbate, R. Conlin, G. Verdoolaege, and E. Kolemen, "Real-time and adaptive reservoir computing with application to profile prediction in fusion plasma," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–12, 2021. DOI: 10 . 1109 / TNNLS.2021.3085504.

[17] C. Rea and R. S. Granetz, "Exploratory machine learning studies for disruption prediction using large databases on DIII-D," *Fusion Science and Technology*, vol. 74, no. 1-2, pp. 89–100, 2018.

[18] Y. Fu, D. Eldon, K. Erickson, *et al.*, "Machine learning control for disruption and tearing mode avoidance," *Physics of Plasmas*, vol. 27, no. 2, p. 022 501, 2020.

[19] B. Cannas, P. C. de Vries, A. Fanni, A. Murari, A. Pau, and G. Sias, "Automatic disruption classification in JET with the ITER-like wall," *Plasma Physics and Controlled Fusion*, vol. 57, no. 12, p. 125 003, 2015.

[20] C. Rea, R. S. Granetz, K. Montes, *et al.*, "Disruption prediction investigations using machine learning tools on DIII-D and Alcator C-Mod," *Plasma Physics and Controlled Fusion*, vol. 60, no. 8, p. 084 004, 2018.

[21] A. Murari, M. Lungaroni, E. Peluso, *et al.*, "Adaptive predictors based on probabilistic SVM for real time disruption mitigation on JET," *Nuclear Fusion*, vol. 58, no. 5, p. 056 002, 2018.

[22] J. Kates-Harbeck, A. Svyatkovskiy, and W. Tang, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning," *Nature*, vol. 568, no. 7753, pp. 526–531, 2019.

[23] E. Aymerich, G. Sias, F. Pisano, *et al.*, "Disruption prediction at JET through deep convolutional neural networks using spatiotemporal information from plasma profiles," *Nuclear Fusion*, vol. 62, no. 6, p. 066 005, 2022.

[24] K. J. Montes, C. Rea, R. S. Granetz, *et al.*, "Machine learning for disruption warnings on Alcator C-Mod, DIII-D, and EAST," *Nuclear Fusion*, vol. 59, no. 9, p. 096 015, 2019.

[25] A. Bustos, E. Ascasıbar, A. Cappa, and R. Mayo-Garcıa, "Automatic identification of MHD modes in magnetic fluctuation spectrograms using deep learning techniques," *Plasma Physics and Controlled Fusion*, vol. 63, no. 9, p. 095 001, 2021.

[26] V. Škvára, V. Šmıdl, T. Pevnỳ, J. Seidl, A. Havránek, and D. Tskhakaya, "Detection of Alfvén eigenmodes on

COMPASS with generative neural networks," *Fusion Science and Technology*, vol. 76, no. 8, pp. 962–971, 2020.

[27] B. J. Woods, V. N. Duarte, E. D. Fredrickson, N. N. Gorelenkov, M. Podestà, and R. G. Vann, "Machine learning characterization of Alfvénic and sub-Alfvénic chirping and correlation with fast-ion loss at NSTX," *IEEE Transactions on Plasma Science*, vol. 48, no. 1, pp. 71–81, 2020.

[28] S. Haskey, B. Blackwell, and D. Pretty, "Clustering of periodic multichannel timeseries data with application to plasma fluctuations," *Computer Physics Communications*, vol. 185, no. 6, pp. 1669–1680, Jun. 2014, ISSN: 00104655. DOI: 10.1016/j.cpc.2014.03.008. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0010465514000885 (visited on 12/21/2021).

[29] S. Haskey, B. Blackwell, C. Nührenberg, *et al.*, "Experiment-theory comparison for low frequency BAE modes in the strongly shaped H-1NF stellarator," *Plasma Physics and Controlled Fusion*, vol. 57, no. 9, p. 095 011, 2015.

[30] A. Jalalvand, A. A. Kaptanoglu, A. V. Garcia, *et al.*, "Alfvén eigenmode classification based on ECE diagnostics at DIII-D using deep recurrent neural networks," *Nuclear Fusion*, vol. 62, no. 2, p. 026 007, 2021.

[31] A. A. Kaptanoglu, A. Jalalvand, A. Garcia, *et al.*, "Exploring data-driven models for spatiotemporally local classification of Alfvén eigenmodes," *Nuclear Fusion*, 2022.

[32] M. Van Zeeland, G. Kramer, R. Nazikian, H. Berk, T. Carlstrom, and W. Solomon, "Alfvén eigenmode observations on DIII-D via two-colour CO2 interferometry," *Plasma physics and controlled fusion*, vol. 47, no. 9, p. L31, 2005.

[33] W. W. Heidbrink, M. A. Van Zeeland, M. E. Austin, *et al.*, "'BAAE' instabilities observed without fast ion drive," *Nuclear Fusion*, vol. 61, no. 1, p. 016 029, 2020.

[34] P. Steiner, A. Jalalvand, S. Stone, and P. Birkholz, "PyRCN: A toolbox for exploration and application of Reservoir Computing Networks," *Engineering Applications of Artificial Intelligence*, vol. 113, p. 104 964, 2022.

[35] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.