
Deep convolutional neural networks for multi-scale time-series classification and application to disruption prediction in fusion devices

R.M. Churchill

Theory Department
Princeton Plasma Physics Laboratory
100 Stellarator Road, Princeton, NJ 08540, USA
rchurchi@pppl.gov

and the DIII-D team

General Atomics
P.O. Box 85608, San Diego, California 92186, USA

Abstract

The multi-scale, multi-physics nature of fusion plasmas makes predicting plasma events challenging. Recent advances in deep convolutional neural network architectures (CNN) utilizing dilated convolutions enable accurate predictions on sequences which have long-range, multi-scale characteristics, such as the time-series generated by diagnostic instruments observing fusion plasmas. Here we apply this neural network architecture to the popular problem of disruption prediction in fusion tokamaks, utilizing raw data from a single diagnostic, the Electron Cyclotron Emission imaging (ECEi) diagnostic from the DIII-D tokamak. ECEi measures a fundamental plasma quantity (electron temperature) with high temporal resolution over the entire plasma discharge, making it sensitive to a number of potential pre-disruptions markers with different temporal and spatial scales. Promising, initial disruption prediction results are obtained training a deep CNN with large receptive field ($\sim 30k$), achieving an F_1 -score of $\sim 91\%$ on individual time-slices using only the ECEi data.

1 Introduction

Plasma phenomena contain a wide range of temporal and spatial scales, often exhibiting multi-scale characteristics (see Figure 1). In fusion energy plasmas, many disparate diagnostic instruments are simultaneously used in order to capture these various spatiotemporal scales, and to cover the multiple physics present in these plasmas. In addition, fusion experiments are increasingly built to run longer pulses, with a goal of eventually running a reactor continuously. The confluence of these facts leads to large, complex datasets with phenomena manifest over long sequences. A key challenge is enabling scientists/engineers to utilize these long sequence datasets to, for example, automatically catalog events of interest or predict the onset of phenomena.

Many deep learning architectures have been created and successfully applied to sequence learning [12, 16, 17, 8] problems, in areas of time-series analysis or natural language processing. However, many of the typical architectures used for learning from sequences (e.g. recurrent neural networks (RNN) and its most popular variant Long Short Time Memory networks (LSTM)) suffer from memory loss; long-range dependencies in sequences are difficult for these architectures to track [2].




Figure 1: Example temporal and spatial scales of different broad physics phenomena in fusion plasmas, based on Ref. [1]

In this paper we discuss recent advances in neural networks, specifically an architecture that uses dilated convolutions in a deep convolutional neural network (CNN), which was designed to overcome these problems of learning on long sequences. We use this architecture to predict oncoming disruptions in fusion plasma discharges of the DIII-D tokamak utilizing only raw data from a single, high temporal resolution imaging diagnostic (the Electron Cyclotron Emission imaging diagnostic, or ECEi) [22]. Because the ECEi diagnostic is sensitive to a range of multi-scale dynamics in the plasma related to disruptions [3], it offers the potential to more accurately predict them. Avoiding disruptions is a grand challenge for tokamak fusion devices on the road to fusion energy [13]. While much research has gone into utilizing machine learning for disruption prediction [25, 21, 14], often global, reduced 0-D features are used in shallow machine learning methods. Recently work utilizing deep LSTM networks also added the use of low temporal resolution 1-D plasma profiles [14], and another work used a combination CNN/LSTM on resampled, low temporal resolution bolometer data [9]. The work we present here takes inspiration from these works in utilizing higher dimensional signals, and shows how to use newer deep learning architectures to learn on high-temporal resolution data with long-range dependencies due to multi-scale physics.

2 Deep convolutional neural networks with dilated convolutions

Recently there has been much research into deep learning architectures which can overcome the deficiencies of RNN/LSTM's, and handle long, multi-scale sequences [24, 5, 20, 26, 10]. A seminal paper presented one such architecture, WaveNET [23], which is a convolutional neural network (CNN) focused on generating realistic audio. One of the key insights of this paper was to use dilated convolutions to increase the receptive field of the network. This overcomes the dilemma faced with using normal convolutions in causal networks, where to be sensitive to long sequences you must increase the convolutional filter size and/or the number of layers in the network. Dilated convolutions have a dilation factor (d) which represents the number of input points skipped between filter parameters, e.g. the sequence output $y[n]$ from a dilated convolution with dilation d is:

$$y[n] = \sum_{i=0}^{k-1} w[i] x[n - d \cdot i]$$

where w represents the weights of the 1D dilated convolution filter of length k , and $x[n]$ is the input sequence. A normal convolution results by setting $d = 1$. By stacking layers of dilated convolutions, and increasing the dilation factor in each layer, the receptive field of the network can be increased while maintaining a tractable number of model parameters.

Dilated convolutions impose an inductive bias or specific structure to the architecture which guide the transformations learned by the neural network. Specifically, dilated convolutions have a natural connection with wavelet structures, which have been used for separating out structure in multi-scale data, including turbulent flows [7]. In a loose sense, these neural networks allow us to learn the wavelet coefficients needed to accomplish our classification task.

A simplified yet powerful architecture named temporal convolutional network (TCN) [2] built upon this WaveNET work, utilizing dilated convolutions and many modern neural network techniques, such as weight normalization and residual connections. Bai *et. al.* [2] showed the TCN could outperform LSTM and GRU architectures on many common sequence learning tasks, especially for long sequences with long-range dependencies. It is this TCN architecture that we will now apply to the problem of disruption prediction using ECEi data.

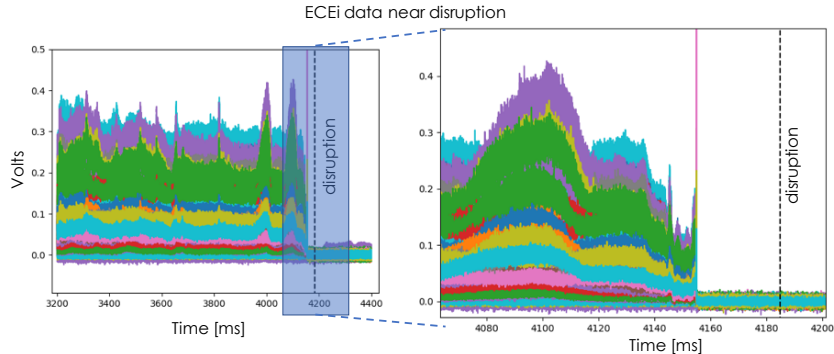


Figure 2: DIII-D ECEi diagnostic time-series data from each of the 160 channels of the LFS diagnostic, near a disruption event. The sudden drop in ECEi signal a few milliseconds before the disruption time (which is set at the current quench time) is due to the drop in temperature at the thermal quench.

3 Application to Disruption Prediction using Raw ECEi Imaging Data

Disruptions in tokamak plasmas are a sudden loss of control which cause a termination of the plasma and potentially large destructive forces and/or heating on the containment vessel and protective wall materials. Next-step devices such as ITER and beyond will have a low tolerance for disruptions [4]. We need to ensure disruptions can be avoided by accurate prediction of oncoming disruptions and mitigation techniques if necessary.

Here we apply the TCN architecture to high-temporal resolution, raw ECEi imaging data from the DIII-D tokamak for the purpose of predicting oncoming disruptions¹.

3.1 Data

The ECEi diagnostic [22] is used to measure electron temperature on very fast timescales, normally sampling at 1 MHz on the DIII-D tokamak. The diagnostic has 160 spatial channels, laid out in a rectangular grid with 20 vertical by 8 radial channels. Example time series of the DIII-D ECEi diagnostic near a disruption is shown in Figure 2. ECEi can capture a number of plasma phenomena such as turbulence fluctuations, tearing modes, sawteeth, and ELMs [22], which allow it to be sensitive to a number of pre-disruption markers. A dataset of good ECEi data ($\text{SNR} > 3$) from 2,747 DIII-D shots ($\sim 42\%$ disruptive, $\sim 58\%$ non-disruptive) was selected, measuring about 10 TB. Time length of each shot varies, typically between 5 to 10 seconds. Raw digitizer voltage output was corrected for digitizer drift, then z -normalized before inputting into the TCN. For ease of training the neural network, we decided as an initial step to temporally downsample the ECEi data to 100 kHz (i.e. factor of 10x less data).

3.2 Model and Training Setup

We treat the problem of disruption prediction as a binary classification problem, where we predict whether each time slice corresponds to a “non-disruptive” or “disruptive” class. We label all time slices within 300ms of a disruption as “disruptive” ($t_{\text{disrupt}} - t < 300\text{ms}$), and all other time slices as “non-disruptive” [21] (sequences from shots without disruptions are taken during established times of the discharge, i.e. during the plasma current flattop). Typical binary cross-entropy loss is used as the loss function for the neural network training.

We define our TCN model to have a receptive field of $N_{\text{recept}} \sim 30,000$. This is an order of magnitude larger than receptive fields in the original TCN [2] or WaveNET [23] papers. With the 100 kHz sampling rate, this means that each time slice prediction uses the previous $\sim 300\text{ms}$ in order to make the prediction. With our definition of disruptive time slices as within 300ms of the disruption, this implicitly assumes that 600ms before a disruption is sufficient to predict oncoming disruptions.

¹Code available at <https://github.com/rmchurch/disruptcnn>




Figure 3: Results from training the TCN on ECEi data.

We use a 4 hidden layer TCN with dilations $[1, 10, 100, 961]$ (i.e. increasing by a factor of about 10 each layer), with a filter kernel size of 15. The number of filters per hidden layer was held constant at 80 (varying number of filters per hidden layer was not attempted).

The TCN architecture allows parallelization of the sequence prediction by inputting sequences of length N_{seq} , which are longer than N_{recept} , resulting in $N_{seq} - N_{recept} + 1$ predictions per sequence. Empirically it was found that sequence lengths of $N_{seq} = 78, 125$ allowed model computations that fit inside the GPU memory constraints, while allowing a batch size of 12 (per GPU) to ensure sufficient variety within each batch for training with stochastic gradient descent (the total batch size with data parallelism was 192. Larger batch size can be achieved reducing the sequence length, though at an increased computational cost due to more data reads). The set of sequences with timeslices consisting of only the majority class (“non-disruptive”) was undersampled such that there were balanced disruptive and non-disruptive sequences.

Stochastic Gradient Descent (SGD) with Nesterov momentum 0.9 was used to train the model, with an initial learning rate of 0.5 that was decreased automatically upon plateau (ReduceLRonPlateau). A warmup period was used for the first 5 epochs, increasing the learning rate from 0.0625 to 0.5 to enable larger batch training [11]. Multi-node, multi-GPU setup was used to parallelize the training. The Pytorch built-in synchronous data parallel training routine `DistributedDataParallel` was used [19], training on 16 GPUs over 2 days.

3.3 Results

The results of training this TCN model on ECEi data for disruption prediction on DIII-D are shown in Figure 3. Results are plotted over 1000 training epochs. The training binary cross-entropy loss continually decreases over the training, showing our model has the capacity to learn the task from this dataset. The validation loss also continually decreases, slightly flattening towards the end, indicating the model is reaching the limit of its generalizability after 1000 epochs. Two validation metrics are also shown: accuracy (how many time slices were predicted correctly as disruptive or non-disruptive), and F1-score (a geometric mean between precision and recall). Because the time slice classes are imbalanced (even though the sequence sets are balanced), the F1-score gives a better indication of how well our classifier does on the minority class (disruptive).

The metric of accuracy reaches $\sim 94\%$, but more importantly the metric of F1-score reaches $\sim 91\%$, showing the neural network has learned to predict individual time slices of both disruptive and non-disruptive time slices very well. Current machine learning disruption predictors typically achieve a true-positive rate in the low 90% on shots [25, 21, 14, 9], with the goal of $>95\%$ with a false-positive rate of $<5\%$ [4]. The results presented here offer a promising path to overcome this gap. Consolidation of the time slice predictions to make shot predictions is left for future work, including not training on the last 30 ms before a disruption since this is a minimum amount of time needed to trigger mitigation systems. Most likely techniques like the hysteresis threshold algorithm will be needed [18], as the fast, noisy ECEi data could easily trigger occasional false predictions.

4 Discussion and Future Work

These results show the usefulness of deep convolutional neural networks with dilated convolutions for fusion problems where the multi-scale, multi-physics nature mandates capturing long-range dependencies in time-series. They have shown that it is possible to apply deep learning directly on the raw data from a single diagnostic with high temporal resolution in order to make useful disruption predictions, a topic critical to the success of magnetic confinement fusion. They also show that training TCN networks with large receptive fields on the order of $\sim 30k$ is possible, allowing learning on long sequences with long-range dependencies.

Future work in various areas is planned. At the base level, using the full dataset at full temporal resolution could give further improvement, though may require model parallelism to train. Further, combining multiple modalities (including more diagnostics) [6] can allow the disruption predictions to be sensitive to the various physics which can trigger disruptions [4]. Also, interpretability of the network decisions is highly desired, especially to understand the physics and extend to future machines [15].

Acknowledgments

The main author would like to thank Ben Tobias, Yilun Zhu, Dave Schissel, C.S. Chang, Bill Tang, Julien Kates-Harbeck, Raffi Nazikian, Cristina Rea, Bob Granetz, Neville Luhmann, Sean Flanagan, Ahmed Diallo, and Ken Silber for various contributions to this work. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Fusion Energy Sciences, under AC02-09CH11466, DE-FC02-04ER54698, and FG02-99ER54531. We also recognize the Princeton Research Computing center for the computational resources used in this paper.

References

- [1] FESAC ISOFS Subcommittee Final Report. Technical report, 2002.
- [2] S. Bai, J. Z. Kolter, and V. Koltun. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv e-prints*, page arXiv:1803.01271, mar 2018.
- [3] M. Choi, H. Park, G. Yun, W. Lee, N. Luhmann, K. Lee, W.-H. Ko, Y.-S. Park, B. Park, and Y. In. 2D/3D electron temperature fluctuations near explosive MHD instabilities accompanied by minor and major disruptions. *Nucl. Fusion*, 56(6):066013, jun 2016.
- [4] P. C. de Vries, G. Pautasso, D. Humphreys, M. Lehnen, S. Maruyama, J. A. Snipes, A. Vergara, and L. Zabeo. Requirements for Triggering the ITER Disruption Mitigation System. *Fusion Sci. Technol.*, 69(2):471–484, apr 2016.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. oct 2018.
- [6] V. Dumoulin, E. Perez, N. Schucher, F. Strub, H. Vries, A. Courville, and Y. Bengio. Feature-wise transformations. *Distill*, 3(7):e11, jul 2018.
- [7] M. Farge. Wavelet Transforms and their Applications to Turbulence. *Annu. Rev. Fluid Mech.*, 24(1):395–458, jan 1992.
- [8] H. I. Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller. Deep learning for time series classification: a review. sep 2018.
- [9] D. R. Ferreira. Applications of Deep Learning to Nuclear Fusion Research. nov 2018.
- [10] J. Gehring, M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin. Convolutional Sequence to Sequence Learning. may 2017.
- [11] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He. Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour. *arXiv e-prints*, page arXiv:1706.02677, jun 2017.

- [12] A. Graves. Generating Sequences With Recurrent Neural Networks. aug 2013.
- [13] T. Hender, J. Wesley, J. Bialek, A. Bondeson, A. Boozer, R. Buttery, A. Garofalo, T. Goodman, R. Granetz, Y. Gribov, O. Gruber, M. Gryaznevich, G. Giruzzi, S. Günter, N. Hayashi, P. Helander, C. Hegna, D. Howell, D. Humphreys, G. Huysmans, A. Hyatt, A. Isayama, S. Jardin, Y. Kawano, A. Kellman, C. Kessel, H. Koslowski, R. L. Haye, E. Lazzaro, Y. Liu, V. Lukash, J. Manickam, S. Medvedev, V. Mertens, S. Mirnov, Y. Nakamura, G. Navratil, M. Okabayashi, T. Ozeki, R. Paccagnella, G. Pautasso, F. Porcelli, V. Pustovitov, V. Riccardo, M. Sato, O. Sauter, M. Schaffer, M. Shimada, P. Sonato, E. Strait, M. Sugihara, M. Takechi, A. Turnbull, E. Westerhof, D. Whyte, R. Yoshino, H. Zohm, D. Group, the ITPA MHD, and Magnet. Chapter 3: MHD stability, operational limits and disruptions. *Nucl. Fusion*, 47(6):S128–S202, jun 2007.
- [14] J. Kates-Harbeck, A. Svyatkovskiy, and W. Tang. Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature*, 568(7753):526–531, apr 2019.
- [15] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, and R. Sayres. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). nov 2017.
- [16] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, may 2015.
- [17] Z. C. Lipton, J. Berkowitz, and C. Elkan. A Critical Review of Recurrent Neural Networks for Sequence Learning. may 2015.
- [18] K. Montes, C. Rea, R. Granetz, R. Tinguely, N. Eidietis, O. Meneghini, D. Chen, B. Shen, B. Xiao, K. Erickson, and M. Boyer. Machine learning for disruption warnings on Alcator C-Mod, DIII-D, and EAST. *Nucl. Fusion*, 59(9):096015, sep 2019.
- [19] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *NIPS-W*, oct 2017.
- [20] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language Models are Unsupervised Multitask Learners. Technical report, 2019.
- [21] C. Rea and R. S. Granetz. Exploratory Machine Learning Studies for Disruption Prediction Using Large Databases on DIII-D. *Fusion Sci. Technol.*, pages 1–12, mar 2018.
- [22] B. Tobias, C. W. Domier, T. Liang, X. Kong, L. Yu, G. S. Yun, H. K. Park, I. G. J. Classen, J. E. Boom, A. J. H. Donné, T. Munsat, R. Nazikian, M. Van Zeeland, R. L. Boivin, and N. C. Luhmann. Commissioning of electron cyclotron emission imaging instrument on the DIII-D tokamak and first data. *Rev. Sci. Instrum.*, 81(10):10D928, oct 2010.
- [23] A. Van Den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. WaveNet: A Generative Model for Raw Audio. *arXiv e-prints*, page arXiv:1609.03499, 2016.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention Is All You Need. jun 2017.
- [25] J. Vega, S. Dormido-Canto, J. M. López, A. Murari, J. M. Ramírez, R. Moreno, M. Ruiz, D. Alves, and R. Felton. Results of the JET real-time disruption predictor in the ITER-like wall campaigns. *Fusion Eng. Des.*, 88(6-8):1228–1231, oct 2013.
- [26] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local Neural Networks. nov 2017.