1st Spanish Fusion HPC Workshop – November 27, 2020

EUROfusion

Use of HPC Infrastructures for Deep Learning in Fusion Research

Diogo R. Ferreira[1] and JET Contributors*

EUROfusion Consortium, JET, Culham Science Centre, Abingdon, OX14 3DB, UK
[1] Instituto de Plasmas e Fusão Nuclear (IPFN), Instituto Superior Técnico (IST), Universidade de Lisboa, Portugal
* See the author list of E. Joffrin et al. 2019 Nucl. Fusion 59 112021

Introduction

- First-principles plasma simulations
  - particle-in-cell, gyrokinetic codes, etc.
  - study of physical phenomena
  - use of synthetic data
  - CPU- and/or GPU-bound

- Deep learning
  - data analysis and processing
  - machine learning models
  - use of experimental data
  - GPU-bound

- Plasma simulations + Deep learning
  - models replace expensive simulations
  - study of physical phenomena
  - use of synthetic data
  - CPU- and GPU-bound

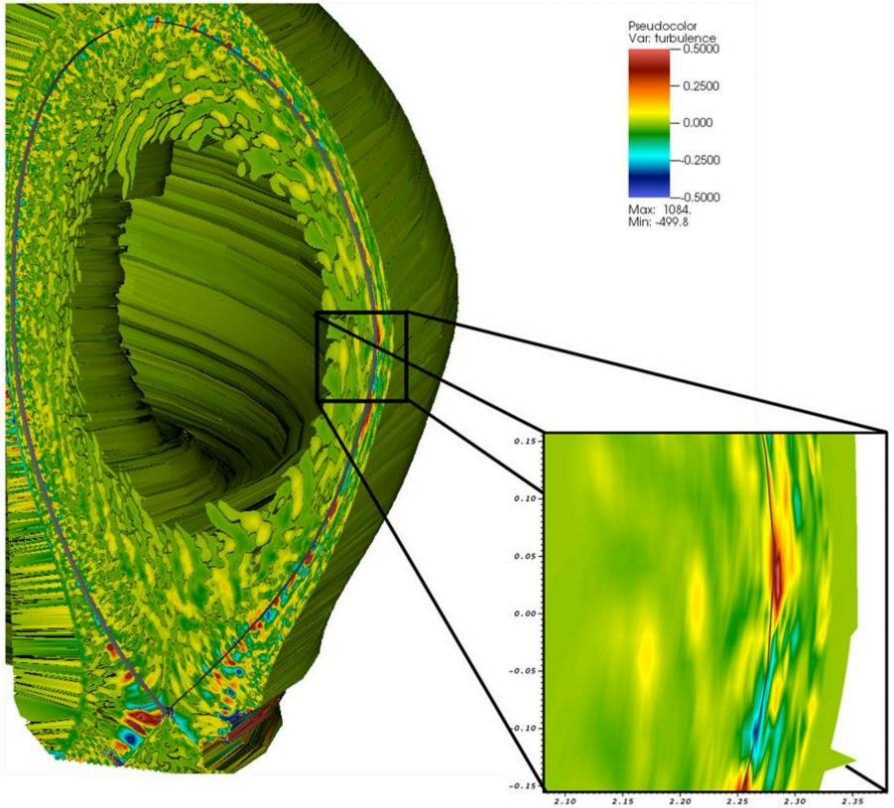

Computer simulation and visualization of edge turbulence in a fusion plasma
(Simulation: Seung-Hoe Ku/PPPL. Visualization: David Pugmire/ORNL)
PPPL News, June 1, 2015

Deep Learning

- Convolutional Neural Networks (CNNs)
  - image reconstruction
- Recurrent Neural Networks (RNNs)
  - disruption prediction
- Variational Autoencoders (VAEs)
  - anomaly detection

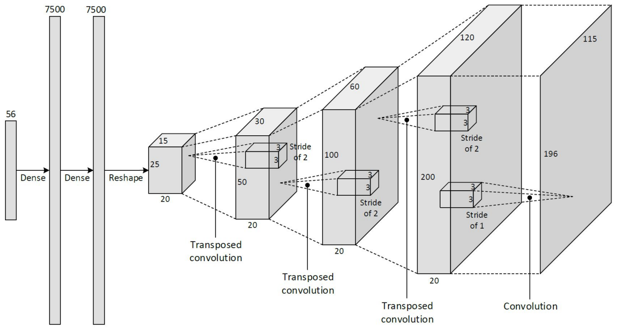

D. R. Ferreira et al, "Full-pulse tomographic reconstruction with deep neural networks", Fusion Sci Technol, vol. 74, no. 1-2, 2018

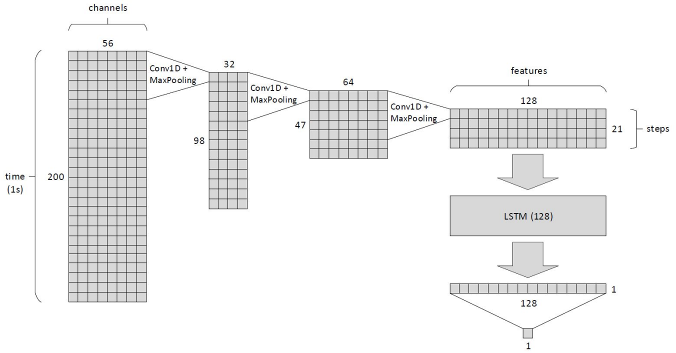

D. R. Ferreira et al, "Deep learning for plasma tomography and disruption prediction from bolometer data", IEEE T Plasma Sci, vol. 48, no. 1, 2020

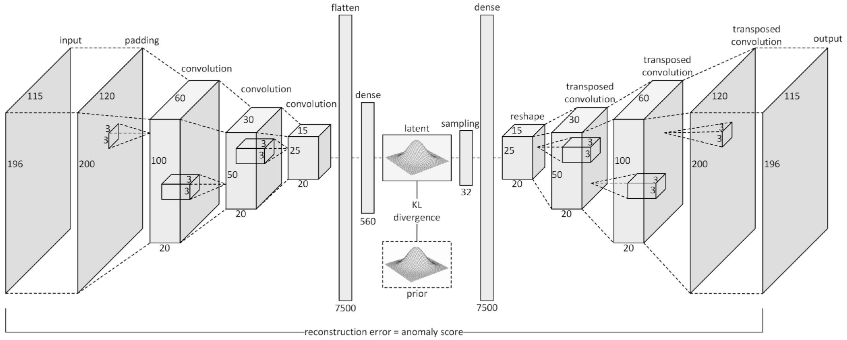

D. R. Ferreira et al, "Deep learning for the analysis of disruption precursors based on plasma tomography", Fusion Sci Technol, 2020 (to appear)

Deep Learning

- Convolutional Neural Networks (CNNs)
  - image reconstruction **from bolometer data**
- Recurrent Neural Networks (RNNs)
  - disruption prediction **from bolometer data**
- Variational Autoencoders (VAEs)
  - anomaly detection **from bolometer data**

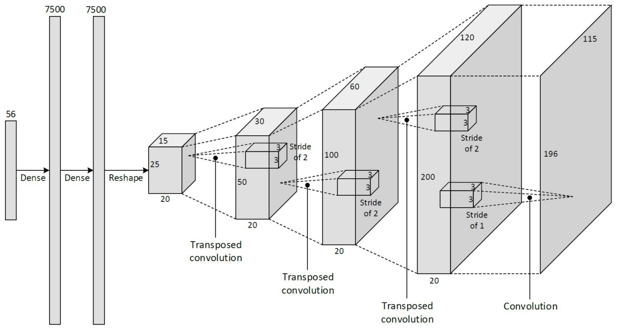

D. R. Ferreira et al, "Full-pulse tomographic reconstruction with deep neural networks", Fusion Sci Technol, vol. 74, no. 1-2, 2018

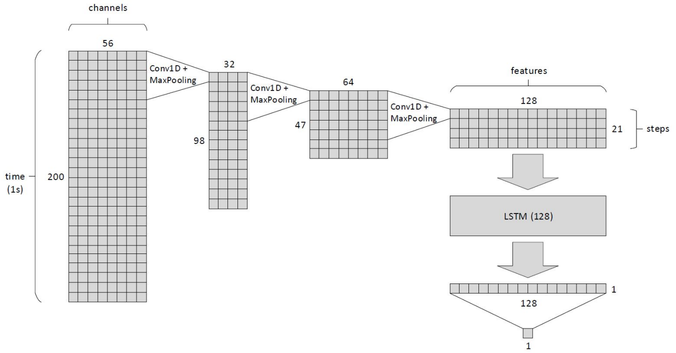

D. R. Ferreira et al, "Deep learning for plasma tomography and disruption prediction from bolometer data", IEEE T Plasma Sci, vol. 48, no. 1, 2020

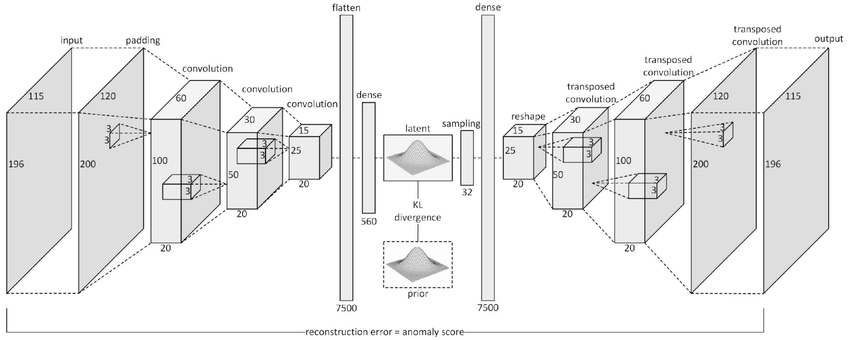

D. R. Ferreira et al, "Deep learning for the analysis of disruption precursors based on plasma tomography", Fusion Sci Technol, 2020 (to appear)

Plasma tomography at JET

- JET bolometer diagnostic
  - horizontal camera + vertical camera
  - 24 bolometers each + 8 reserve
  - 56 lines of sight over the plasma
  - line-integrated radiation
  - UV to soft X-ray range

- Bolometer tomography
  - reconstruct plasma radiation profile
  - several techniques available[1]
    - minimum Fisher, maximum likelihood, etc.
  - method developed at JET[2]
    - iterative constrained optimization algorithm
  - can be accelerated by deep learning[3]

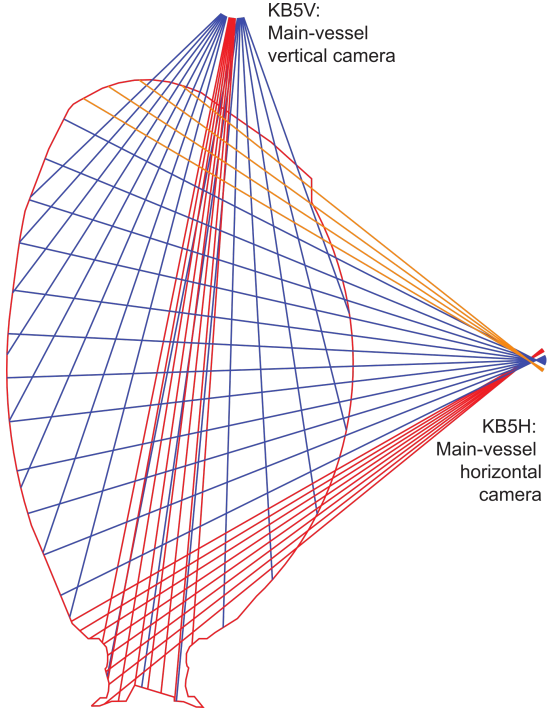

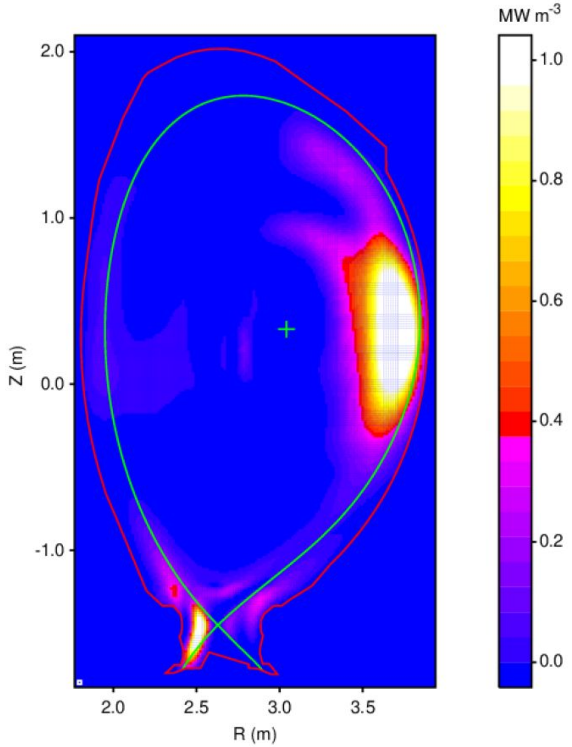

[1] J. Mlynar et al, "Current research into applications of tomography for fusion diagnostics", J. Fusion Energy, vol. 38, no. 3, 2019
[2] L. C. Ingesson et al, "Soft X ray tomography during ELMs and impurity injection in JET", Nucl. Fusion, vol. 38, no. 11, 1998
[3] F. A. Matos et al, "Deep learning for plasma tomography using the bolometer system at JET", Fusion Eng Des, vol. 114, 2017

Convolutional Neural Network (CNN)

- Deep learning for plasma tomography
  - input is **bolometer data**, output is **plasma profile**
  - trained on ~20k samples, fits memory of single GPU
  - day(s) on single GPU, hours on multi-GPU node

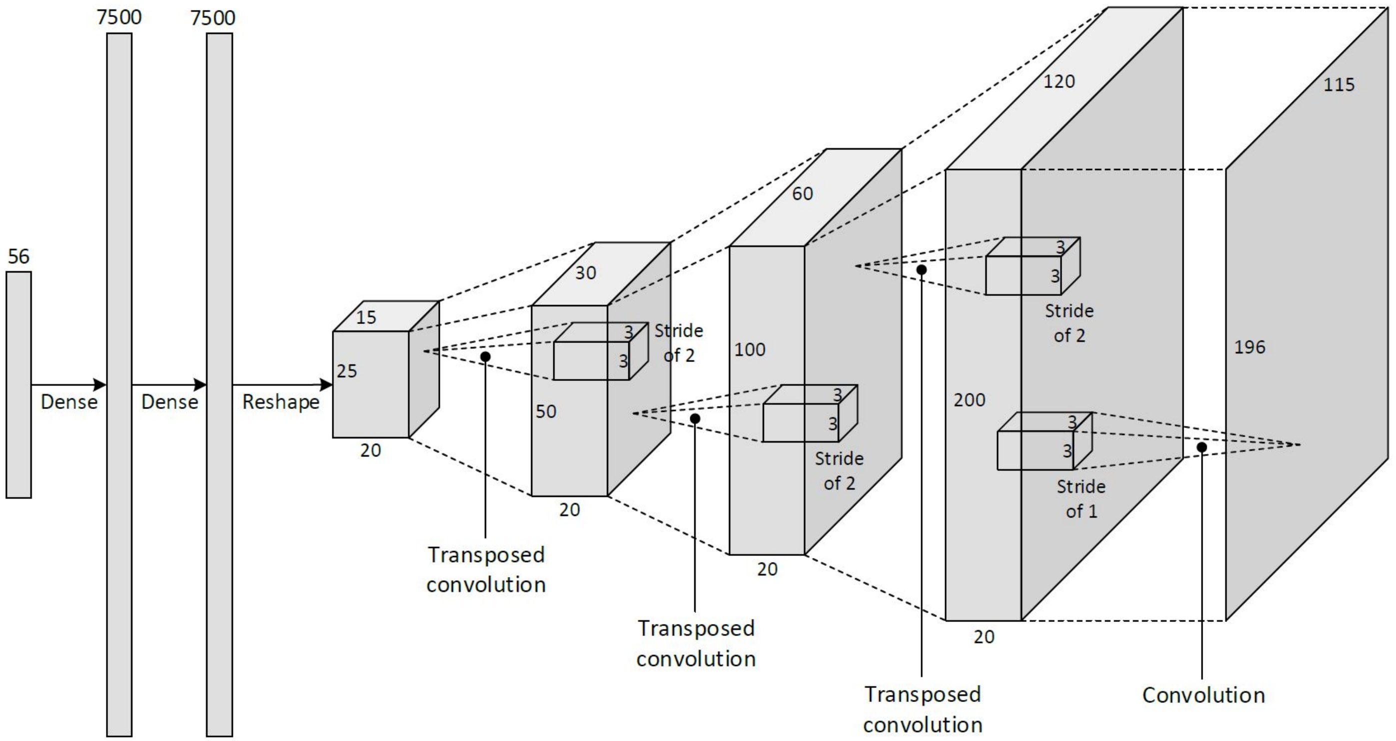

D. R. Ferreira et al, “Full-pulse tomographic reconstruction with deep neural networks”, Fusion Sci Technol, vol. 74, no. 1-2, 2018

Convolutional Neural Network (CNN)

- Deep learning for plasma tomography
  - input is **bolometer data**, output is **plasma profile**
  - trained on ~20k samples, fits memory of single GPU
  - day(s) on single GPU, hours on multi-GPU node

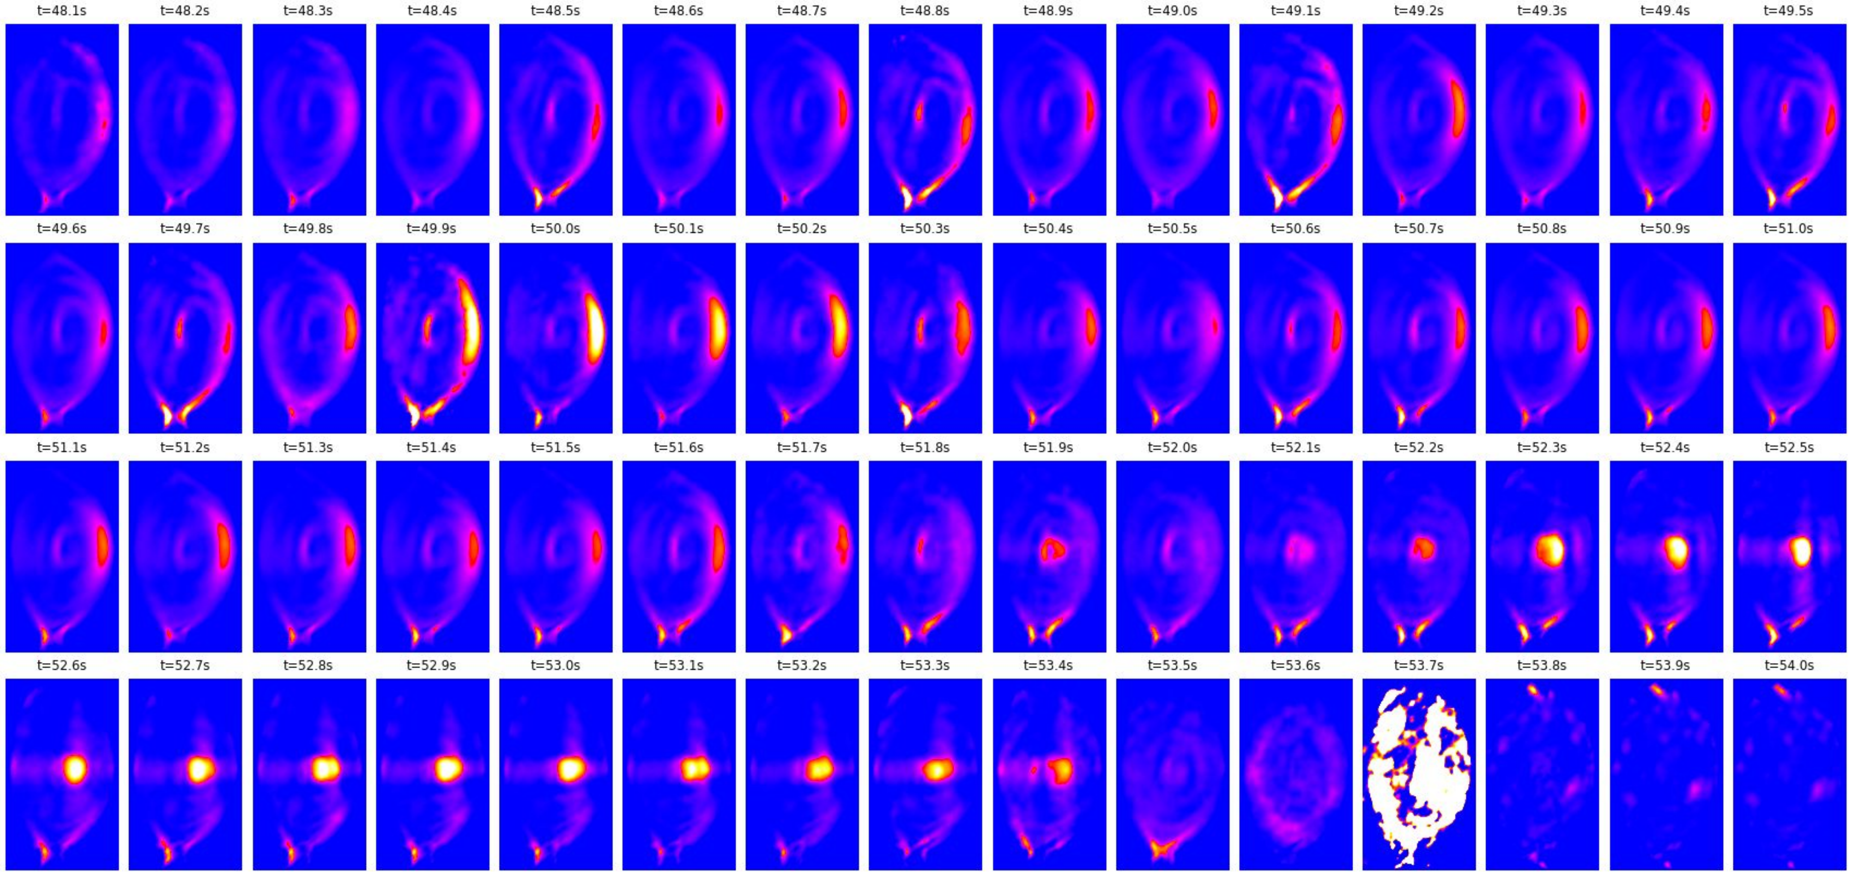

Recurrent Neural Network (RNN)

- Deep learning for disruption prediction
  - input is **bolometer data**, output is **time to disruption** or **probability of disruption** (two models, same architecture)
  - trained on samples drawn at random from ~10k pulses
  - the two models can be trained simultaneously on separate nodes/GPUs

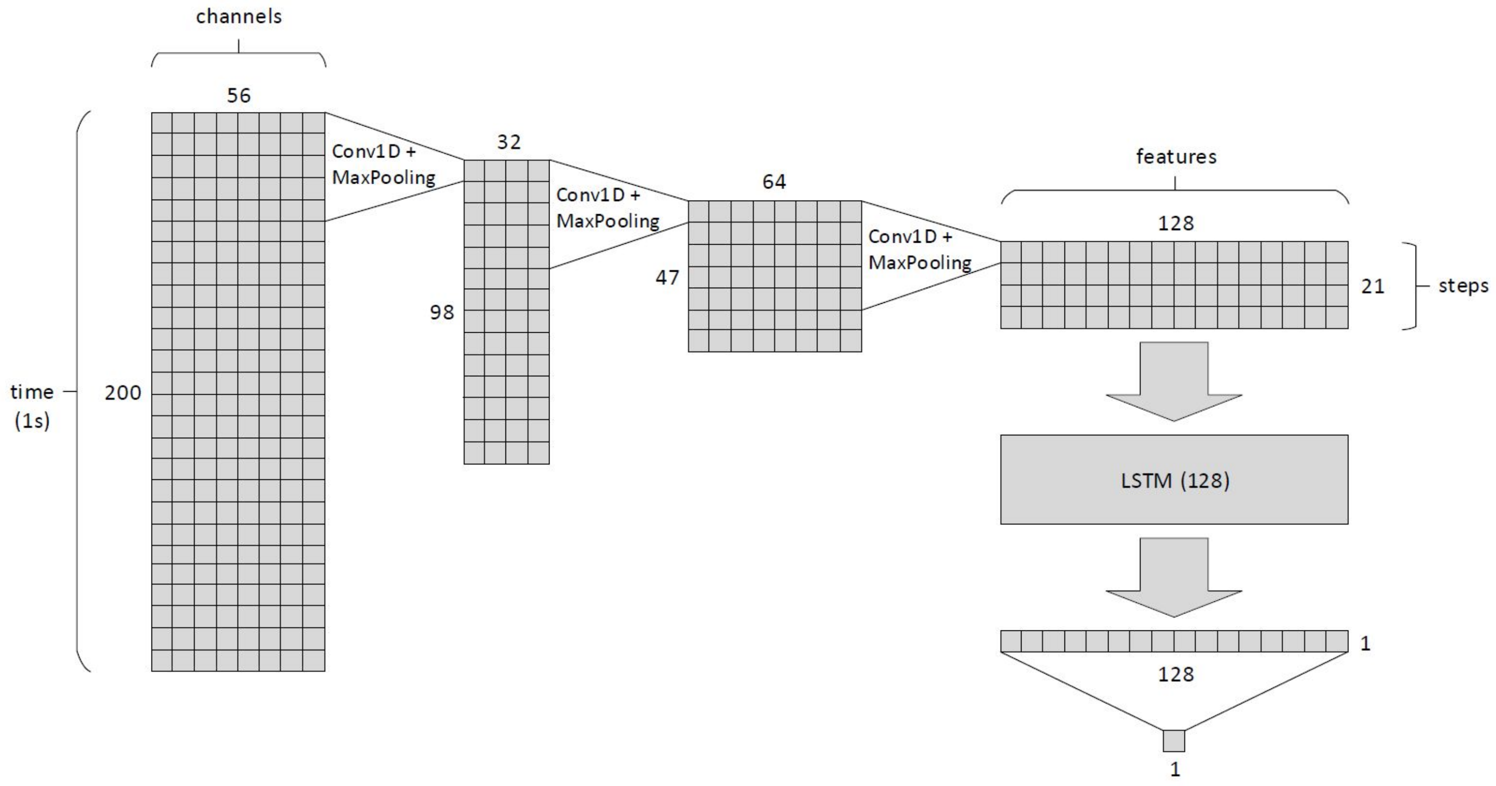

D. R. Ferreira et al, “Deep learning for plasma tomography and disruption prediction from bolometer data”, IEEE T Plasma Sci, vol. 48, no. 1, 2020

Recurrent Neural Network (RNN)

- Deep learning for disruption prediction
  - input is **bolometer data**, output is **time to disruption** or **probability of disruption** (two models, same architecture)
  - trained on samples drawn at random from ~10k pulses
  - the two models can be trained simultaneously on separate nodes/GPUs

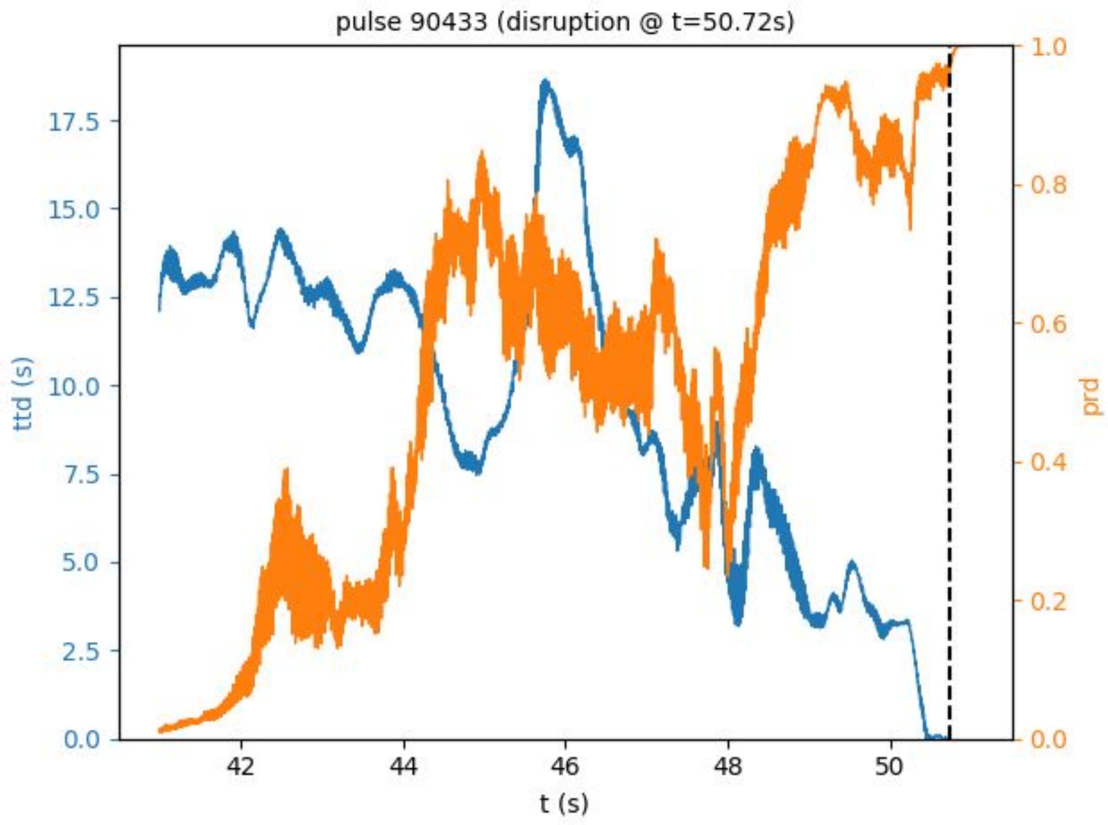

Variational Autoencoder (VAE)

- Deep learning for anomaly detection
  - input is **plasma profile**, output is **plasma profile** (reconstruction error = anomaly score)
  - trained on ~1.4 million profiles from ~250 non-disruptive pulses, tested on disruptive pulses
  - does not fit memory of single GPU, takes day(s) on multi-GPU node

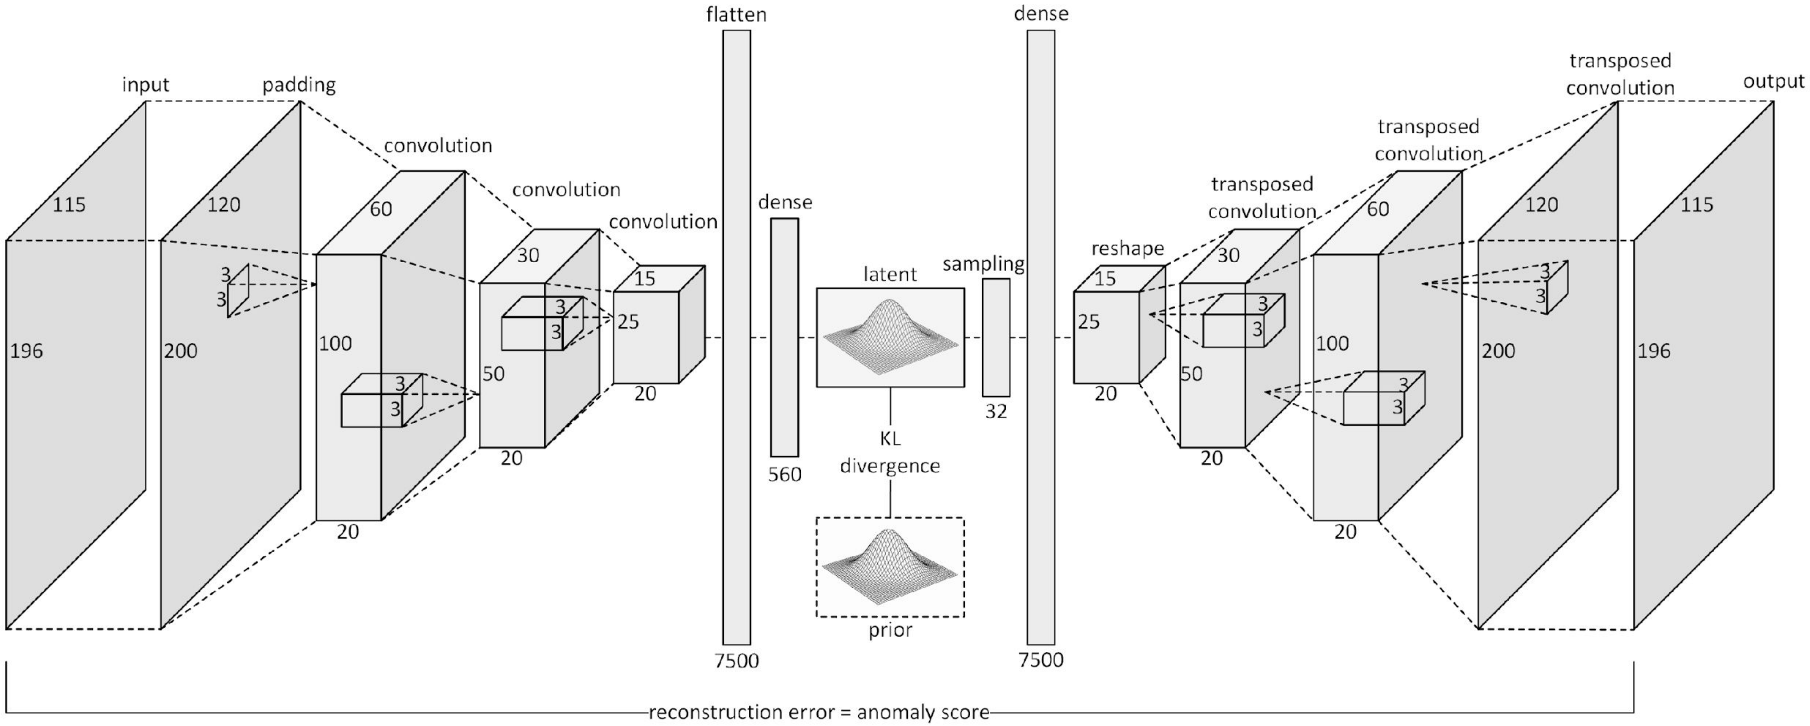

D. R. Ferreira et al, "Deep learning for the analysis of disruption precursors based on plasma tomography", Fusion Sci Technol, 2020 (to appear)

Variational Autoencoder (VAE)

- Deep learning for anomaly detection
  - input is **plasma profile**, output is **plasma profile** (reconstruction error = anomaly score)
  - trained on ~1.4 million profiles from ~250 non-disruptive pulses, tested on disruptive pulses
  - does not fit memory of single GPU, takes day(s) on multi-GPU node

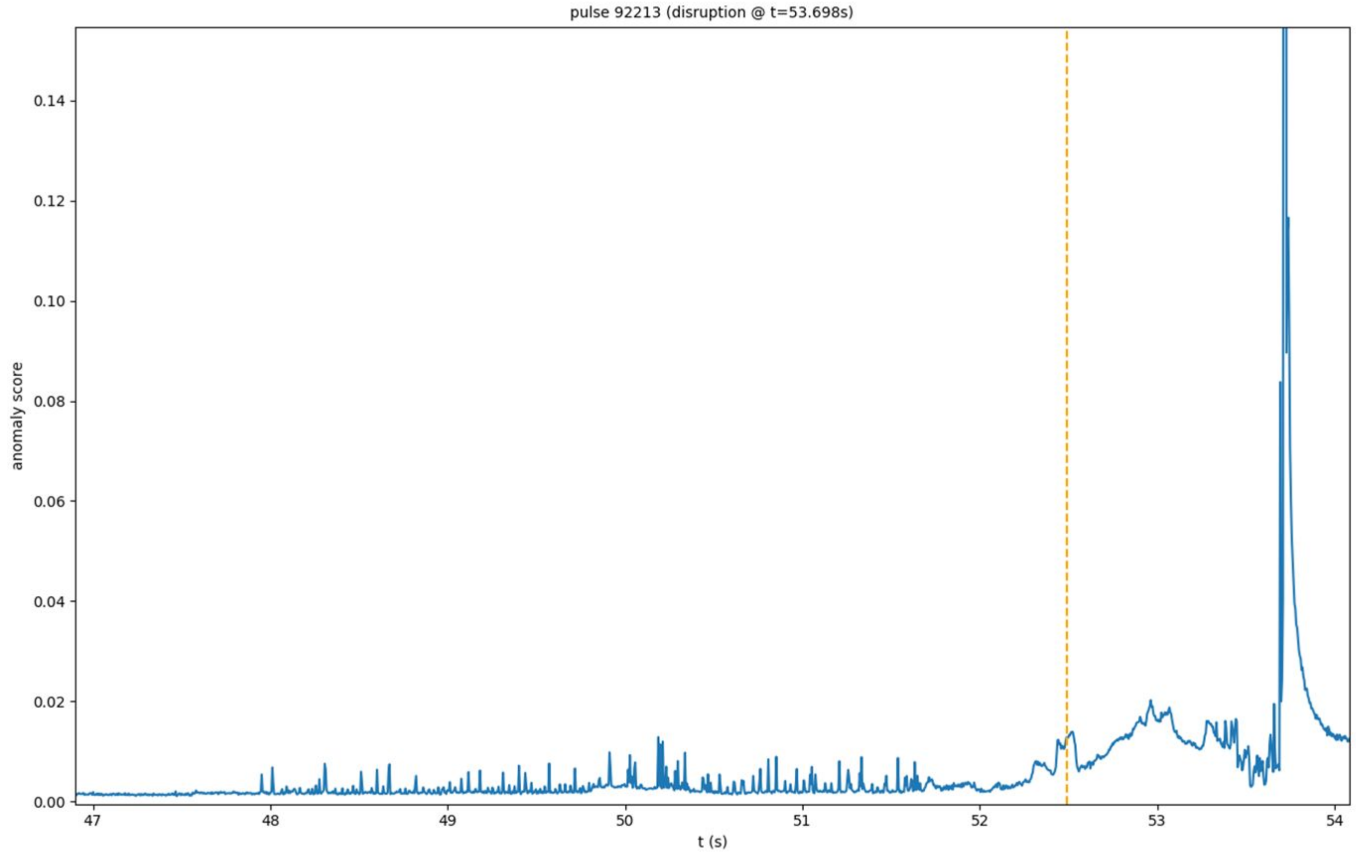

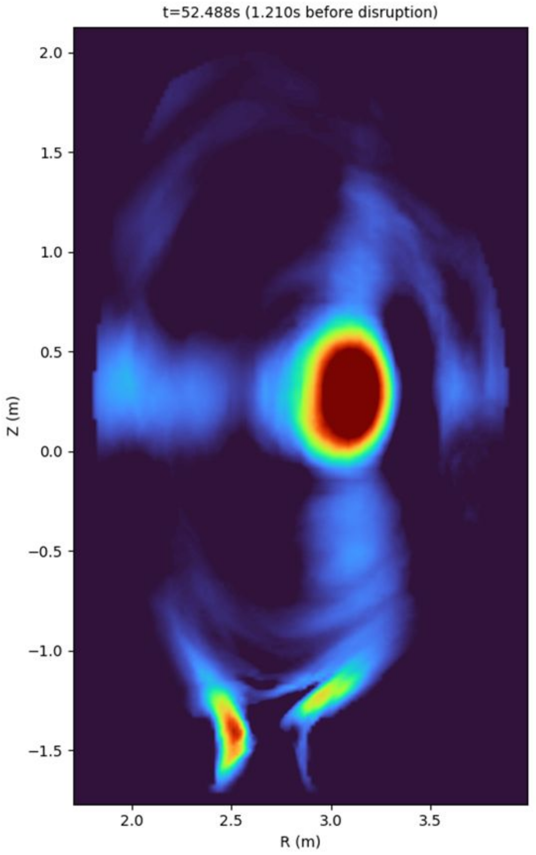

Scaling it up

- Single GPU
  - small model
  - large model, small training data

- Multiple GPUs, single node
  - large model, large training data
  - small models trained separately

- Multiple GPUs, multiple nodes
  - large models trained separately
  - hyperparameter tuning
  - automated machine learning (autoML)

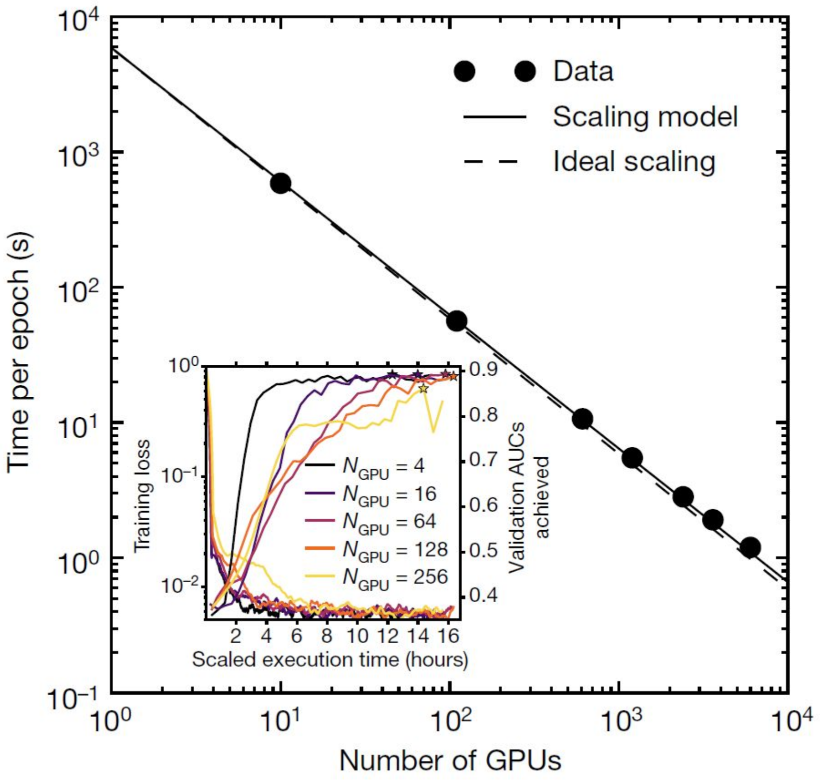

J. Kates-Harbeck et al, "Predicting disruptive instabilities in controlled fusion plasmas through deep learning", Nature, vol. 506, 2019

Conclusion

- A lot of opportunities for deep learning in fusion research
- Availability of GPU partitions in most HPC clusters
- Fundamentally different levels of computation

NVIDIA CEO Jensen Huang introduces the NVIDIA A100 data center GPU
(May 2020)