

Use of HPC Infrastructures for Deep Learning in Fusion Research

Diogo R. Ferreira¹ and JET Contributors*

EUROfusion Consortium, JET, Culham Science Centre, Abingdon, UK

¹ *IPFN / IST, University of Lisbon, Portugal*

In fusion research, the use of HPC infrastructures is mostly dedicated to plasma simulation and modeling tasks, using particle-in-cell, gyrokinetic codes, etc. However, more recently there have been several opportunities for applying deep learning to the analysis and processing of experimental data, in particular data collected from plasma diagnostics. In this talk we will show examples of such tasks, involving the use of Convolutional Neural Networks (CNNs) for image reconstruction [1], Recurrent Neural Networks (RNNs) for disruption prediction [2] and Variational Autoencoders (VAEs) [3] for anomaly detection. All of these examples are based on a single diagnostic – the bolometer system at JET. The use of HPC infrastructures for training such models is somewhat atypical. Rather than being CPU-intensive, deep learning tasks are GPU-centered, where the CPU has a subordinate role of dispatching work to the GPUs, in the form of batches of data. For tasks with small-sized models (or large models, but with small amounts of training data) it may be possible to reduce the batch size significantly in order to fit the task within the memory of a single GPU which, ideally, should be as large as possible. For large models with large amounts of training data, a multi-GPU setup will be required, but for convenience all GPUs should be in a single node, since multi-GPU support heavily depends on the features provided by the underlying deep learning framework. The use of multiple nodes, with multiple GPUs each, finds application in those scenarios which require hyperparameter tuning or, more recently, automated machine learning (AutoML) where each node trains a different variant of the same model, in a resource-exhausting search for the best-performing model or model architecture.

References

- [1] D. R. Ferreira, P. J. Carvalho, H. Fernandes, and JET Contributors. Full-pulse tomographic reconstruction with deep neural networks. *Fusion Science and Technology*, 74(1-2):47–56, 2018.
- [2] D. R. Ferreira, P. J. Carvalho, H. Fernandes, and JET Contributors. Deep learning for plasma tomography and disruption prediction from bolometer data. *IEEE Transactions on Plasma Science*, 48(1):36–45, 2020.
- [3] D. R. Ferreira, P. J. Carvalho, C. Sozzi, P. J. Lomas, and JET Contributors. Deep learning for the analysis of disruption precursors based on plasma tomography. *Fusion Science and Technology*, 2020. (to appear).

*See the author list of E. Joffrin et al. 2019 Nucl. Fusion 59 112021.