Scalability Analysis of Super-Resolution Generative Adversarial Network Training for Turbulence Closure Modeling

L. Nista[1], S. Rocco[2,3], T. Grenga[1], A. Attili[4], G. Cavallaro[2], M. Riedel[2,3], H. Pitsch[1]

[1] *Institute for Combustion Technology, RWTH Aachen University, Aachen, Germany*
[2] *Jülich Supercomputing Center, Forschungszentrum Jülich, Germany*
[3] *School of Engineering and Natural Sciences, University of Iceland, Reykjavik, Iceland*
[4] *Institute for Multiscale Thermofluids, University of Edinburgh, Edinburgh, United Kingdom*

Main domain: Computational Fluid Dynamics, Data-driven numerical modeling of flows.
Possible subdomains: Deep Learning for Fluid Flows, Distributed Deep-Learning for Turbulence Closure Modeling.

Over the last years, deep-learning approaches have become increasingly popular in the context of Computational Fluid Dynamics (CFD), where there is growing interest in developing data-driven subfilter-scale (SFS) models for Large Eddy Simulations (LES). This has been promoted by the remarkable performance of recent Machine Learning (ML)-based architectures as well as to the exponential improvement in computer hardware performance. Deep Convolutional Neural Networks (CNN) can be used to super-resolve three-dimensional turbulent flow realizations from solely the coarse-grained data fields, enhancing subgrid physical structures. However, such architectures often lack generalization capabilities and cannot guarantee high-wavenumber details, fundamentals to represent dynamics, and transport below explicitly resolved scales [Ledig *et al.*, *IEEE CVPR*, 2017].

A recent viable solution is the usage of Generative Adversarial Networks (GAN), composed of two competing Deep CNNs, to increase information contents with similarity to genuine high-resolution images [Wang *et al.*, *ECCV*, 2018]. Despite the remarkable performance of GAN in single-image super-reconstruction, its application in turbulence modeling applications is relatively unexplored [Bode *et al.*, *Proc. Comb. Inst.*, 2021]. For example, available Super-Resolution GANs (SRGAN) are not yet extended to explicitly incorporate physics-based loss functions, such as symmetries and conserved quantities. Moreover, GANs also have the propensity for running into convergence oscillations due to the unbalanced training, for which an appropriate hyperparameter search is required. Many researchers are then still hesitant to explore such novel options as the adversarial training is usually computationally more expensive compared to more classical ML approaches and needs finer tuning.

Despite all the progress in parallelizing deep neural network models that have already revolutionized many application fields, including computer vision, speech recognition, and image processing, little has been done so far in turbulence closure modeling [Campos *et al.*, *IEEE/ACM*, 2017]. The typical Direct Numerical Simulation (DNS) datasets have significantly larger but fewer samples compared to image processing. Moreover, as the dynamic of turbulence is greatly influenced by three-dimensional vortex-stretching effects, deep-learning models must employ three-dimensional fully-connected CNNs with the consequence of increasing the number of trainable parameters. These characteristics may lead to very high pressure on Graphical Processing Units (GPU) memory and poor convergence capabilities [Keshar *et al.*, *ArXiv*, 2016].

The aim of this work is to show how to efficiently use High-Performance Computing (HPC) systems to speed up the training time of SRGAN architectures specifically for turbulence modeling applications. The objective is to provide some useful insights on how those challenges can be circumvented on distributed GPU clusters (e.g.: *JURECA-Booster* Julich) to allow the development of next-level data-driven turbulence modeling that learn and correctly reproduce the SFS statistics.