

SEMINAR ANNOUNCEMENT

Wednesday 6th May 2020, Time: 15.30-17.30

Microsoft Teams, Team code: f3co430



Giuseppe Fiameni, PhD

NVIDIA

Artificial Intelligence Technology Centre

Large Scale Training of Deep Neural Networks

Abstract: The computational requirements of deep neural networks used to enable AI applications like objects captioning or identification are enormous. A single training cycle can take weeks on a single GPU, or even months for the larger datasets like those used in computer vision research. Using multiple GPUs for deep learning can significantly shorten the time required to train lots of data, making solving complex problems with deep learning feasible.

This two-hour seminar will introduce you on how to use multiple GPUs to training neural networks. You'll learn:

- Approaches to multi-GPU training
- Algorithmic and engineering challenges to large-scale training

Upon completion, you'll be able to effectively parallelize training of deep neural networks.

Short bio: Giuseppe Fiameni, PhD, is a Solution Architect for AI and Accelerated Computing at NVIDIA, helping researchers in optimizing deep learning workloads on High Performance Computing systems. He is the technical lead of the Italian NVIDIA Artificial Intelligence Technology Centre.

Info: Prof. Carlo Sansone - tel. 081 7683606 – carlo.sansone@unina.it