



Machine Learning Training: Research Challenges and Opportunities for the Distributed Systems Community

Dr. Giovanni Neglia

INRIA,
Université Côte D'Azur, Sophia Antipolis, FR

Monday, February 10th 2020, h:15.30
Aula 5.6, Building: "Aule Nuove"

About the speaker

Giovanni Neglia received the master's degree in electronic engineering and the PhD degree in telecommunications from the University of Palermo, Italy, in 2001 and 2005, respectively. He has been a researcher at Inria, Sophia Antipolis, France, since September 2008. In 2005, he was a research scholar with the University of Massachusetts, Amherst, visiting the Computer Networks Research Group. Before joining Inria, he was a post-doctorate with the University of Palermo and an external scientific advisor with the Maestro Team at Inria. His research focus on modeling and performance evaluation of networks. He is area editor for the Elsevier Computer Communication journal (COMCOM), and associate editor for IEEE Trans. on Mobile Computing.

Abstract

In this talk, I will support the thesis that the distributed systems community is not meant to simply apply machine learning (ML) tools to their classic research problems, but can also contribute to design faster and more efficient distributed ML systems both for training and inference. I will first introduce machine learning training and show that computational speedups directly translate into better ML models. I will then explain why design choices for ML systems are inevitably entangled with optimization and statistical considerations. Finally, I will provide two examples from my recent research activity: dynamic (TCP-like) adaptation of the number of ML workers, and topology design.

Contact

Prof. Giuseppe Notarstefano giuseppe.notarstefano@unibo.it