DIPARTIMENTO DI
INFORMATICA, SISTEMISTICA E
COMUNICAZIONE

Viale Sarca, 336
20126 Milano

SEMINAR ANNOUNCEMENT

Monday February 5th, 2024

at 10:00 am
Room "Sala Seminari" - Abacus Building (U14)

Meeting with NVIDIA (NVAITC & Research)

Speakers
Andrea Pilzer Solutions Architect, NVIDIA
Iuri Frosio Principal Research Scientist, NVIDIA

Abstract

Introduction to NVAITC and High Performance Deep Learning.
This talk presents the NVIDIA AI Technology Center (NVAITC), a program to enable and accelerate AI research projects in Italy, with the mission of accelerating the pipelines of academic partners and advising them about the use of the most suitable NVIDIA technologies for each task. The talk will include examples and insights in dealing with large-scale model training with PyTorch. Topics covered in this PyTorch Multi-GPU approach to Deep learning Models include Data and Model Parallelism, Message Passing, Distributed training using Horovord, Mixed Precision, Memory Format, and Pipeline Parallelism.

The Best Defense is a Good Offense: Adversarial Augmentation against Adversarial Attacks.
Many defenses against adversarial attacks (e.g. robust classifiers, randomization, or image purification) use countermeasures put to work only after the attack has been crafted. We adopt a different perspective to introduce A5 (Adversarial Augmentation Against Adversarial Attacks), a novel framework including the first certified preemptive defense against adversarial attacks. The main idea is to craft a defensive perturbation to guarantee that any attack (up to a given magnitude) towards the input in hand will fail. To this aim, we leverage existing automatic perturbation analysis tools for neural networks. We study the conditions to apply A5 effectively, analyze the importance of the robustness of the to-be-defended classifier, and inspect the appearance of the robustified images. We show effective on-the-fly defensive augmentation with a robustifier network that ignores the ground truth label, and demonstrate the benefits of robustifier and classifier co-training. In our tests, A5 consistently beats state of the art certified defenses on MNIST, CIFAR10, FashionMNIST and Tinyimagenet. We also show how to apply A5 to create certifiably robust physical objects. Our code at https://github.com/NVlabs/ A5 allows experimenting on a wide range of scenarios beyond the man-in-the-middle attack tested here, including the case of physical attacks. Link: https://openaccess.thecvf.com/content/CVPR2023/papers/Frosio_The_Best_Defense_Is_a_Good_Offense_Adversarial_Augmentation_Against_CVPR_2023_paper.pdf

Andrea Pilzer is a Solution Architect at NVIDIA and works with NVIDIA AI Technology Center in Italy. He was a postdoc at Aalto University, working on uncertainty estimation for deep learning. He worked at Huawei Ireland and got his Ph.D. in CS from the University of Trento working with Nicu Sebe and Elisa Ricci.

Iuri Frosio got his PhD in biomedical engineering at the Politecnico of Milan in 2006. He was a research fellow at the Computer Science Department of the University of Milan from 2003 and an assistant professor in the same department from 2006 to 2013. In the same period, he worked as a consultant for various companies in Italy and in the US. He joined NVIDIA in 2014 as senior research scientist, and since 2021 he has the role of principal research scientist. His research interests include image processing, computer vision, robotics, parallel programming, machine learning, and reinforcement learning.

contact person for this Seminar: alessandro.raganato@unimib.it