

Scalable Deep Learning over Parallel and Distributed Infrastructures (ScaDL)

Scope of the Workshop

Recently, Deep Learning (DL) has received tremendous attention in the research community because of the impressive results obtained for a large number of machine learning problems. The success of state-of-the-art deep learning systems relies on training deep neural networks over a massive amount of training data, which typically requires a large-scale distributed computing infrastructure to run. In order to run these jobs in a scalable and efficient manner, on cloud infrastructure or dedicated HPC systems, several interesting research topics have emerged which are specific to DL. The sheer size and complexity of deep learning models when trained over a large amount of data makes them harder to converge in a reasonable amount of time. It demands advancement along multiple research directions such as, model/data parallelism, model/data compression, distributed optimization algorithms for DL convergence, synchronization strategies, efficient communication and specific hardware acceleration.

This intersection of distributed/parallel computing and deep learning is becoming critical and demands specific attention to address the above topics which some of the broader forums may not be able to provide. The aim of this workshop is to foster collaboration among researchers from distributed/parallel computing and deep learning communities to share the relevant topics as well as results of the current approaches lying at the intersection of these areas.

General Chairs

Gauri Joshi, Carnegie Mellon University (gaurij@andrew.cmu.edu)

Ashish Verma, IBM Research AI (ashish.verma1@us.ibm.com)

Program Chairs

Yogish Sabharwal, IBM Research AI

Parijat Dube, IBM Research AI

Local Chair

Eduardo Rodrigues, IBM Research

Steering Committee

Vijay K. Garg, University of Texas at Austin

Vinod Muthuswamy, IBM Research AI

Technical Program Committee

Alvaro Coutinho - Federal University of Rio de Janeiro

Dimitris Papailiopoulos, University of Wisconsin-Madison

Esteban Meneses, Costa Rica Institute of Technology

Kangwook Lee, KAIST

Li Zhang, IBM Research

Lydia Chen, TU Delft

Philippe Navaux, University of Rio Grande do Sul

Rahul Garg, Indian Institute of Technology Delhi

Vikas Sindhwani, Google Brain

Wei Zhang, IBM Research

Xiangru Lian, University of Rochester