

Workshop 19: ScaDL Scalable Deep Learning over Parallel and Distributed Infrastructures

Preface

It is our great pleasure to welcome you to the second edition of the workshop on Scalable Deep Learning over Parallel and Distributed Infrastructure (ScaDL)!

Recently, Deep Learning (DL) has received tremendous attention in the research community because of 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 massive amounts of training data, which typically requires large-scale distributed computing infrastructure to run. 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 thus becoming critical; this workshop aims to bring these two communities together to foster collaboration, discuss relevant topics and share results.

In addition to five peer-reviewed research papers, ScaDL 2020 also features four invited presentations, from Dr. Manish Gupta (Google Research, India), Prof. Geoffrey Fox (Indiana University, USA), Prof. Wen-mei Hwu (UIUC, USA) and Dr. Minsik Cho (IBM, USA). We sincerely thank our invited speakers for their time and efforts.

We would like to extend special thanks to the ScaDL 2020 program committee members for their substantial effort in writing high-quality reviews within the short reviewing timeframe (Feb 14 - Mar 6, 2020). Each submission received constructive feedback that we are confident was of considerable value to its authors, whether or not it was accepted. We would also like to thank the authors of all the submissions for taking the time to document their research findings, thereby advancing the science of scalable deep learning platforms and infrastructure.

Due to the unprecedented situation surrounding COVID-19, ScaDL 2020 is being held virtually (live). We hope to make it as close as possible to an in-person event.

May 8, 2020

Ashish Verma (IBM Research, USA) & Christopher Carothers (RPI, USA)
General Chairs

K. R. Jayaram (IBM Research, USA) & Parijat Dube (IBM Research, USA)
Program Chairs

Organizing Committee

General Chairs

Ashish Verma, IBM Research, USA
Christopher Carothers, Rensselaer Polytechnic Institute, USA

Program Chairs

K. R. Jayaram, IBM Research, USA
Parijat Dube, IBM Research, USA

Publicity Chair

Danilo Ardagna, Politecnico di Milano, Italy

Program Committee

Kangwook Lee, University of Wisconsin, Madison, USA
Li Zhang, IBM Research, USA
Xiangru Lian, University of Rochester, USA
Eduardo Rocha Rodrigues, IBM, Brazil
Wagner Meira Jr., Universidade Federal de Minas Gerais, Brazil
Stacy Patterson, Rensselaer Polytechnic Institute, USA
Alex Gittens, Rensselaer Polytechnic Institute, USA
Catherine Schuman, Oak Ridge National Laboratory, USA
Ignacio Blanquer, Universitat Politècnica de Valencia, Spain
Leandro Balby Marinho, Universidade Federal de Campina Grande, Brazil
Chen Wang, IBM Research, USA

Steering Committee

Vijay K. Garg, University of Texas at Austin, USA
Vinod Muthusamy, IBM Research, USA
Yogish Sabharwal, IBM Research, USA
Danilo Ardagna, Politecnico di Milano, Italy