



Deconstructing Distributed Deep Learning (and other problems)

Dr. Leana Golubchik
Stephen and Etta Varra Professor
Department of Computer Science
USC Viterbi School of Engineering

Thursday, September 8, 2022
10:00am – 11:00am
EEB 248

<https://usc.zoom.us/j/92960680288?pwd=QndJKzUrSIVsaXZCSkJKRTlIdTNOUT09>
Meeting ID: 929 6068 0288 **Passcode:** 610848

Abstract: Deep learning has made substantial strides in many applications. New training techniques, larger datasets, increased computing power, and easy-to-use machine learning frameworks all contribute to this success. An important missing piece is that deep learning frameworks do not assist users with provisioning cloud resources; most users need to try different job configurations to determine the resulting training performance. When resources are shared among hundreds of jobs, this approach quickly becomes infeasible. At a larger scale, when multiple datacenters need to manage deep learning (and other) workloads, different degrees of affinity for their resources create incentives to collaborate, e.g., as in cloud federations. In this talk, I will focus on our approach to predicting performance metrics and scheduling algorithms that use these metrics to guide resource allocation. Our goal is to broaden the population of users capable of developing deep learning models and applying them to novel applications. I will also discuss resource sharing for such workloads, mainly in the context of hybrid cloud federations, and conclude with some future directions.

Bio: Leana Golubchik is the Stephen and Etta Varra Professor at USC. She also serves as the Director of the Women in Science and Engineering (WiSE) program. Prior to that, she was on the faculty at the University of Maryland and Columbia University. Leana received her PhD from UCLA. Her research interests are broadly in the design and evaluation of large scale distributed systems, including hybrid clouds and data centers and their applications in data analytics, machine learning, and privacy. Leana received several awards, including the IBM Faculty Award, the NSF CAREER Award, the Okawa Foundation Award, the WTS-LA Diversity Leadership Award, the USC Remarkable Women Award, and the USC Mellon Culture of Mentoring Award. She is the Editor-in-Chief of the ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS) and on the Editorial Board of the Performance Evaluation journal as well as a member of the IFIP WG 7.3 (elected in 2000). She is a Fellow of AAAS.

Host: Dr. Richard M. Leahy (leahy@sipi.usc.edu)