

COMPUTER SCIENCE
SEMINAR

Machine Learning Systems for the Data Tsunami

Christopher De Sa
Cornell University

Abstract: Much of the recent advancement in machine learning has been driven by the capability of machine learning systems to process and learn from very large data sets using very complicated models. Continuing to scale data up in this way to handle the present "data tsunami" presents a computational and algorithmic challenge, as power, memory, and time are all factors that limit performance. In this talk, I will discuss some recent advances from my lab that address these issues at every level of the systems stack, including algorithmic changes that make accurate statistical inference on large datasets feasible, numerical changes that increase our capabilities to train complicated models over unreliable networks, and principled approaches that ensure the accountability of large-scale learning systems.

Bio: Christopher De Sa is an Assistant Professor in the Cornell Department of Computer Science, with additional field membership in ECE and Statistics. His research covers algorithmic, software, and hardware techniques for high-performance machine learning, with a focus on relaxed-consistency variants of stochastic algorithms such as asynchronous and low-precision stochastic gradient descent (SGD) and Markov chain Monte Carlo.

Friday, October 30, 2020, 1:00 pm
<https://emory.zoom.us/j/92722816908>

COMPUTER SCIENCE
EMORY UNIVERSITY