

Computer Science Seminar Series

National Capital Region

Scalable Methods for Large-scale Machine Learning and Sparse Signal Recovery

Speaker: Dr. Thomas Goldstein
Department of Computer Science
University of Maryland
Friday, April 8, 2016
1:00PM - 2:00PM, NVC 325

Abstract

The abundance of large, distributed web-based data sets and the recent popularity of cloud computing platforms has opened many doors in machine learning and statistical modeling. However, these resources present a host of new algorithmic challenges. Practical algorithms for large-scale data analysis must scale well across many machines, have low communication requirements, and have low (nearly linear) runtime complexity to handle extremely large problems. In this talk, we discuss several practical and general tools for solving a wide range of model-fitting problems in a distributed framework. We first discuss new "transpose reduction" strategies that allow extremely large regression problems to be solved quickly on a single node. We also discuss new variants of stochastic gradient descent that achieve fast convergence rates with low memory overhead. We will study the performance of these algorithms for fitting linear classifiers and sparse regression models on tera-scale datasets using thousands of cores.

Biography



Tom Goldstein (<http://www.cs.umd.edu/~tomg/>) recently joined the University of Maryland as an assistant professor in the Department of Computer Science, and also serves as a member of the University of Maryland Institute for Advanced Computer Studies (UMIACS). Before joining the faculty at UMD, Tom held research positions at Stanford University and Rice University. Tom's research focuses on efficient, low complexity methods for model fitting and data analysis. His work ranges from large-scale computing on distributed architectures, to inexpensive power-aware algorithms for small-scale embedded systems. Applications of his work include scalable machine learning algorithms, efficient model-fitting methods for computer vision and imaging, and signal processing methods for wireless communications. Tom's research is supported by grants from the National Science Foundation and Office of Naval Research, in addition to resources from Google and the United States Naval Academy. He also conducts research in collaboration with the Department of Defense Super Computing Resource Center.