UNCLASSIFIED

Information Science and Technology Seminar Speaker Series

Joseph K. Bradley
Carnegie Mellon University

Optimizing Trade-offs for Scalable Machine Learning

Wednesday, February 13, 2013

3:30 - 4:30 PM

TA-3, Bldg. 1690, Room 102 (CNLS Conference Room)

Abstract: Modern machine learning applications require large models, lots of data, and complicated optimization. I will discuss scaling machine learning by decomposing learning problems into simpler sub-problems. This decomposition allows us to trade off accuracy, computational complexity, and potential for parallelization, where a small sacrifice in one can mean a big gain in another. Moreover, we can tailor our decomposition to our model and data in order to optimize these trade-offs.

I will present two examples. First, I will discuss parallel optimization for regression, where the goal is to model or predict a label given many other measurements. Our Shotgun algorithm parallelizes coordinate descent, a seemingly sequential method. Shotgun theoretically achieves near-linear speedups and empirically is one of the fastest methods for multicore sparse regression. Second, I will discuss parameter learning for Probabilistic Graphical Models, a powerful class of models of probability distributions. In both examples, our analysis provides strong theoretical guarantees which guide our very practical implementations.

Biography: Joseph Bradley is a Ph.D. candidate in Machine Learning at Carnegie Mellon University, advised by Carlos Guestrin. His thesis is on learning large-scale Probabilistic Graphical Models, focusing on methods which decompose problems to take advantage of parallel computation. Previously, he received a B.S.E. in Computer Science from Princeton University.

For more information contact the technical host Reid Porter, rporter@lanl.gov, 665-7508.

Los Alamos NATIONAL LABORATORY EST.1943

Hosted by the Information Science and Technology Institute (ISTI)

ISTI INFORMATION SCIENCE & TECHNOLOGY INSTITUTE