Gersende Fort

CNRS, Institut de Mathématiques de Toulouse, France

A Variance Reduced Expectation Maximization algorithm for finite sum optimization

Abstract : The Expectation Maximization (EM) algorithm is of key importance for inference in latent variable models including mixture of regressors and experts, missing observations. Unfortunately, its computational cost is prohibitive in the large scale learning setting.

In this talk, we will introduce a novel EM algorithm, called SPIDER-EM, for inference from a large training set and designed for the so-called "finite sum" setting.

We will first show how this algorithm uses Variance Reduction techniques based on control variates, within a Stochastic Approximation scheme in order to provide an estimator of the full conditional expectation in the E-step. We will also outline a parallel with a variance reduced preconditioned stochastic gradient algorithm.

Then, we will derive finite-time complexity bounds for smooth non-convex likelihood optimization: we will discuss how some design parameters scale as a function of the number of examples and of the accuracy level epsilon, in order to control the convergence to an epsilon-approximate stationary point. Based on this criterion, SPIDER-EM improves over the state-of-the-art algorithms.

Numerical results will support our findings.

This is a joint work with Eric Moulines (CMAP, Ecole Polytechnique) and Hoi-To Wai (Chinese Univ. of Hong-Kong).