

Master Internship proposal – April to september 2012

Project: **Large-scale convex optimization for supervised learning**

Many applications of machine learning now take place in situations where both the size p of each observations and the number n of these observations are large (i.e., $n \sim 10^9$, $p \sim 10^6$). In this context, the running time of classical supervised learning techniques (such as the SVM or logistic regression) may become prohibitive, and current research focuses on algorithms based on simple updates after each observation is made [1].

The goal of the internship is to extend recent work [2] which is dedicated to finite amount of data arriving as a batch, to data that may arrive as a stream. The work is at the intersection of algorithms, statistics and optimization, and may focus primarily on any these three aspects depending on the candidate.

Students interested in this project should contact Francis Bach to discuss further.

This internship may be extended to a PhD within the SIERRA team of Département d'Informatique de l'Ecole Normale Supérieure.

Pre-requisites : Machine learning classes (theory, kernel methods)
Convex optimization

Length: 4-6 months

Compensation: Approximately 1200 per month

Laboratory: Département d'Informatique de l'ENS - INRIA
23, avenue d'Italie, Paris 13^e

Contact : Francis Bach (francis.bach@ens.fr)

[1] L. Bottou, O. Bousquet. The trade-off of large scale learning. Advances in Neural Information Processing Systems (NIPS), 2008.

[2] N. Le Roux, M. Schmidt, F. Bach. A Stochastic Gradient Method with an Exponential Convergence Rate for Strongly-Convex Optimization with Finite Training Sets. Advances in Neural Information Processing Systems (NIPS), 2012.
http://hal.inria.fr/docs/00/71/50/15/PDF/sag_arxiv.pdf