

Large-scale machine learning and stochastic algorithms

Leon BOTTOU

Abstract:

During the last decade, data sizes have outgrown processor speed. We are now frequently facing statistical machine learning problems for which datasets are virtually infinite. Computing time is then the bottleneck.

The first part of the presentation shows distinct tradeoffs for the case of small-scale and large-scale learning problems. Small-scale learning problems are subject to the usual approximation--estimation tradeoff.

Large-scale learning problems are subject to qualitatively different tradeoffs involving the computational complexity of the underlying optimization algorithms in non-trivial ways.

Unlikely optimization algorithm such as stochastic gradient shows amazing performance for large-scale machine learning problems.

The second part makes a detailed exploration of stochastic learning algorithms and of their implementation, with both simple and complex examples.

Keywords:

Large scale machine learning, approximation, estimation, optimization, tradeoffs, stochastic gradient