Chapter 1

Introduction

This course will be organized around algorithmic issues that arise in machine learning. The usual paradigm for algorithm design is to give an algorithm that succeeds on all possible inputs, but the difficulty is that almost all of the optimization problems that arise in modern machine learning are *computationally intractable*. Nevertheless, practitioners use a wide variety of heuristics that are successful in practice. However we often do not understand *when* and *why* these approaches work (an issue we would not have if our algorithms came with provable guarantees). The central questions in this course are:

Question 1 *Which* models *in machine learning lead to tractable algorithmic problems?*

Worst-case analysis is *comfortable* because if an algorithm works in this model, it certainly works in practice. But the optimization problems that machine learning systems "solve" everyday are indeed hard in the worst-case. However these lower bounds are not so frightening; many of the hard instances of machine learning problems are not ones we would want to solve in practice anyways! We will see a number of examples where choosing the right model will lead us to discover new algorithms with provable guarantees, where we really can understand when and why they work. In some cases, we will even be able to analyze approaches that practitioners already use and give new insights into their behavior.

Question 2 *Can new models – that better represent the instances we actually want to solve in practice – be the inspiration for developing fundamentally new algorithms for machine learning problems? Can we understand when and why widely used heuristics work?*

This course will focus on

(a) nonnegative matrix factorization

(b) topic modeling

(c) tensor decompositions

(d) sparse recovery

(e) dictionary learning

(f) learning mixtures models

(g) matrix completion

Hopefully more sections will be added to this course over time, since there are a vast number of topics at the intersection of algorithms and machine learning left to explore.

18.409 Algorithmic Aspects of Machine Learning

Spring 2015