



Convex and Nonlinear Optimization for Computer Vision

Andrew Fitzgibbon
Microsoft Research, Cambridge, UK

Abstract

Many tasks in vision and machine learning are best expressed as large-scale nonlinear optimization problems. Although linear models such as PCA may be easily tractable, real world problems often involve nonlinearities (e.g. perspective projection, manifold learning), and lead to nonlinear optimization problems. Although it is widely known that "one shouldn't just use gradient descent", and sometimes known that "alternation works well for a while but then flatlines", it appears to be less well known how to do better, and often, trying a second order method gives worse results. This tutorial will cover several aspects of this domain, from computing matrix derivatives, to classes of optimization algorithms, to reasons why your second-order method may be giving you problems.

Syllabus: Convex and Nonlinear Optimization