

Beyond Convex Optimization

Stephen Boyd Maximilian Schaller Daniel Cederberg

Stanford University

Beyond convex optimization

- ▶ so far in this course, every problem we solved was *convex* (or convexified)
- ▶ convex optimization is great
 - we get the global solution (and an optimality certificate)
 - solvers are fast and reliable, with no initial point and no tuning
 - DCP and CVXPY make problems easy to specify
- ▶ but many problems are inherently nonconvex and cannot be convexified

What to do when your problem is not convex?

- ▶ sometimes possible to reformulate as convex using, *e.g.*, a change of variables
 - maximum likelihood estimation for exponential families
 - feedback control design
 - geometric programming
- ▶ implement simple heuristics based on convex optimization, such as the **convex-concave procedure** (CCP)
- ▶ use a nonlinear programming solver, via **disciplined nonlinear programming** (DNLP)

today's lecture: CCP and DNLP

What we are not covering

- ▶ training neural networks is also nonconvex optimization, and very successful
- ▶ but it is a different regime
 - huge number of variables, (mostly) no constraints
 - stochastic first-order methods (SGD, Adam), stopped early
- ▶ here we consider nonconvex problems
 - with structure and constraints
 - our focus: how to formulate and specify them, not algorithmic details
 - we'll specify them in CVXPY and (hopefully) solve them to high accuracy