

MATH 126 Winter 18: Homework 2

posted on Jan. 19; due by Jan. 26

1. [10 points] Consider the following problem:

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^d} \quad & \mathbf{c}^\top \mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \preceq \mathbf{b} \text{ for all } \mathbf{A} \in \mathcal{A}, \end{aligned}$$

where $\mathcal{A} \subseteq \mathbb{R}^{p \times d}$ is the set

$$\mathcal{A} = \{\mathbf{A} \in \mathbb{R}^{p \times d} : \bar{A}_{ij} - V_{ij} \leq A_{ij} \leq \bar{A}_{ij} + V_{ij}, i = 1, \dots, p, j = 1, \dots, d\}$$

for given matrices $\bar{\mathbf{A}}$ and \mathbf{V} (with $V_{ij} \geq 0$ for all i, j). Express this problem as an LP.

2. [3 points each] Consider a gradient descent method with a constant step size s :

$$\mathbf{x}_{k+1} = \mathbf{x}_k - s \nabla f(\mathbf{x}_k), \quad k = 0, 1, \dots$$

for solving a quadratic problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left\{ f(\mathbf{x}) \equiv \frac{1}{2} \mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{p}^\top \mathbf{x} \right\}$$

where $\mathbf{Q} \in \mathbb{S}_{++}^d$, and we denote μ and L to be the smallest and largest eigenvalues of \mathbf{Q} for $0 < \mu < L$.

- (1) What is the closed-form optimal solution \mathbf{x}_* of the problem?
(2) Show that

$$\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2 \leq \|\mathbf{I} - s\mathbf{Q}\|_2 \|\mathbf{x}_k - \mathbf{x}_*\|_2. \quad (\text{a})$$

- (3) Considering (a), what is the condition of step size s that guarantees a sequence $\{\mathbf{x}_k\}$ to converge to \mathbf{x}_* (i.e., $\|\mathbf{x}_k - \mathbf{x}_*\|_2 \rightarrow 0$ as $k \rightarrow \infty$)?
(4) Considering (a), what is the optimal step size \hat{s} (i.e., the step size that has the fastest rate of convergence)?
(5) Specify a (two-dimensional) worst-case function f of a gradient descent method with the optimal step size \hat{s} (i.e., find f that satisfies $\|\mathbf{x}_{k+1} - \mathbf{x}_*\|_2 = \|\mathbf{I} - \hat{s}\mathbf{Q}\|_2 \|\mathbf{x}_k - \mathbf{x}_*\|_2$).
3. [5 points each] Consider the quadratic problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left\{ \frac{1}{2} \mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{p}^\top \mathbf{x} \right\} \quad (\text{P})$$

where \mathbf{Q} is the 50×50 Hilbert matrix defined by

$$Q_{i,j} = \frac{1}{i+j-1}, \quad i, j = 1, \dots, 50,$$

and $\mathbf{p} = (1, 1/2, 1/3, \dots, 1/50)^\top$. The matrix can be constructed via the MATLAB command `Q = hilb(50)`. Implement the (MATLAB) code for the following five methods efficiently (*i.e.*, avoid redundant computations for fair run time comparison), and compare their performance with the initial vector $\mathbf{x}_0 = (1, 2, 3, \dots, 50)^\top$ and a stopping criteria $\|\nabla f(\mathbf{x})\| \leq 10^{-3}$. (Hint: The optimal value is $p_* = -0.5$, so check whether your algorithm generates a sequence $f(\mathbf{x}_k)$ that reaches to p_* .)

- MATLAB linear system solver¹ (`linsolve` or backslash command),
 - Gradient descent methods with
 - Exact line search
 - Backtracking line search (initial $s = 1$, $\alpha = 0.4$, $\beta = 0.5$)
 - Constant step
 - Conjugate gradient (CG) method
(Its MATLAB code is available at https://en.wikipedia.org/wiki/Conjugate_gradient_method.)
- (1) Present the step size rule for the exact line search of the gradient method for solving (P). Describe your appropriate² choice of the constant step size for the gradient method for solving (P).
 - (2) Report the run time of MATLAB linear system solver. Report the total number of iterations and run time to reach given tolerance for the four iterative optimization algorithms. Also report the per-iteration run time³ for the four cases. (Hint: Use the MATLAB `tic/toc` commands to check the run time.)
 - (3) Compare the four iterative optimization algorithms (except the MATLAB solver) in terms of the total number of iterations and run time to reach given tolerance, and the per-iteration run time, with a brief reasoning for such results (*e.g.*, state a reason why one algorithm has shorter per-iteration run time compared to other algorithms, state why one algorithm took more iterations to reach same tolerance).

¹ Note that this will perform the best among the algorithms considered in Problem 3 for this toy (small-dimensional) problem, but not necessarily for the large-dimensional problems.

² Your *appropriate* step size should guarantee a descent update at every iteration.

³ For simplicity, let's denote the per-iteration run time to be $\frac{(\text{Run time})}{(\text{Total number of iterations})}$, ignoring the fact that some algorithms compute some operations before iterating (*i.e.*, before the first iteration).