ORF 363/COS 323	PROBLEM SET 3	Fall 2025, Princeton University
Instructor: A.A. Ahmadi		
TAs: Budway, Hua, Shi, Tang, Yang		
Due on Tuesday, October 7, 2025, at 1:30pm EST, on Gradescope		

For all problems that involve coding, please include your code.

Problem 1: Radiation treatment planning[1]

In radiation treatment, radiation is delivered to a patient, with the goal of killing or damaging the cells in a tumor, while carrying out minimal damage to other tissue. The radiation is delivered in beams, each of which has a known pattern; the level of each beam can be adjusted. (In most cases multiple beams are delivered at the same time, in one 'shot', with the treatment organized as a sequence of 'shots'.) We let b_j denote the level of beam j, for $j = 1, \dots, n$. These must satisfy $0 \le b_j \le B^{\max}$, where $B^{\max}$ is the maximum possible beam level. The exposure area is divided into m voxels, labeled $i = 1, \dots, m$. The dose d_i delivered to voxel i is linear in the beam levels, i.e., $d_i = \sum_{j=1}^n A_{ij} b_j$. Here $A \in \mathbb{R}_+^{m\times n}$ is a (known) matrix that characterizes the beam patterns. We now describe a simple radiation treatment planning problem.

A (known) subset of the voxels, $\mathcal{T} \subset \{1, \dots, m\}$, corresponds to the tumor or target region. We require that a minimum radiation dose D^{target} be administered to each tumor voxel, i.e., $d_i \ge D^{\text{target}}$ for $i \in \mathcal{T}$. For all other voxels, we would like to have $d_i \le D^{\text{other}}$, where D^{other} is a desired maximum dose for non-target voxels. This is generally not feasible, so instead we settle for minimizing the penalty

$$E = \sum_{i \notin \mathcal{T}} (d_i - D^{\text{other}})_+,$$

where $(\cdot)_+$ denotes the nonnegative part of its argument (i.e., $(z)_+ = \max\{0, z\}$). We can interpret E as the total nontarget excess dose.

1. Show that the treatment planning problem is convex. The optimization variable is $b \in \mathbb{R}^n$; the problem data are $B^{\max}$, A, $\mathcal{T}$, D^{target}, and D^{other}.

2. Solve the problem instance with data generated by the file `treatment_planning_data.m`. If you are using Python, you can download `Atumor.csv` and `Aother.csv`, then use `treatment_planning_data.py` to load them. Here we have split the matrix A into

[1] Courtesy of S. Boyd and L. Vandenberghe.

`Atumor`, which contains the rows corresponding to the target voxels, and `Aother`, which contains the rows corresponding to other voxels. Plot the dose histogram for the target voxels, and also for the other voxels. You can use the function `hist` in both MATLAB and Python (Matplotlib package) to plot histograms. Make a brief comment on what you see. *Remark:* The beam pattern matrix in this problem instance is randomly generated, but similar results would be obtained with realistic data.

Problem 2: Would your GPA be higher at Yale?

An article appeared recently in the New York Times with the title “*Nearly Everyone Gets A’s at Yale. Does That Cheapen the Grade?*”[2] After reading the article, you may wonder whether it is easier to get an A at Yale than at Princeton and, if so, how one could adjust GPAs to account for course difficulty. In this problem, we approach this question using an optimization-based idea proposed by Professor Vanderbei and his collaborators.

We begin by asking four familiar Princeton/Yale students to take some courses at their own institution and a few similar ones at the other institution. The letter grades of these students are summarized in Table 1. Their GPAs are calculated using the grade points in Table 2.

[2] `https://www.nytimes.com/2023/12/05/nyregion/yale-grade-inflation.html`

	Princeton ORF 363	Princeton ENG 351	Princeton ORF 309	Yale CPSC 365	Yale ENGL 305	Yale STAT 241	GPA	Aptitude
M. Obama	A-	A	B+		A+		3.825	?
J. Bezos	A+	A-	B-	A			3.675	?
M. Streep	B-			A	A	A	3.675	?
R. DeSantis			B-	A+	A+	A+	3.9	?
Inflatedness	?	?	?	?	?	?		

Table 1: Performance of four students in six courses

Letter Grade	A+	A	A-	B+	B	B-
Grade Point	4.3	4.0	3.7	3.3	3.0	2.7

Table 2: Converting letter grades to grade points

We assume that the grade point g_{ij} that student i receives in course j should nearly be equal to $a_i + b_j$, where a_i is the "*aptitude*" of student i and b_j is the "*inflatedness*" of course j. In our example, $i \in \{1, \ldots, 4\}$, and $j \in \{1, \ldots, 6\}$. We normalize the inflatedness scores with the constraint $\sum_{j=1}^{6} b_j = 0$ (negative inflatedness scores correspond to more difficult courses). This leads us to the following optimization problem which simultaneously computes student aptitudes and course inflatedness scores:

$$\begin{array}{ll} \min\limits_{a \in \mathbb{R}^4, b \in \mathbb{R}^6} & \sum_{(i,j) \in \mathcal{G}} (g_{ij} - a_i - b_j)^2 \\ \text{s.t.} & \sum\limits_{j=1}^{6} b_j = 0. \end{array} \tag{1}$$

Here, the index set $\mathcal{G}$ denotes the student-course pairs for which a grade is available.

(a) Is problem (1) a convex optimization problem? Why or why not?

(b) Use Tables 1 and 2 to solve problem (1) via `cvx` or `cvxpy`. Fill in the question marks in Table 1 with your optimal solution.

(c) How do the four students rank based on their aptitude (which can be thought of as an "adjusted GPA")? Compare this to the GPA-based ranking. Which courses have the lowest/highest inflatedness score?

Problem 3: Minimizers of convex problems

Let $f : \mathbb{R}^n \to \mathbb{R}$ be a convex function and $\Omega \subseteq \mathbb{R}^n$ be a convex set. Show that the set of minimizers of f over Ω is convex (i.e., a convex set).

Problem 4: True or False? (Provide a proof or a counterexample.)

(a) A quadratic function $f(x) = x^T Q x + b^T x + c$ is convex if and only if it is quasiconvex.

(b) A convex homogeneous polynomial $p(x)$ of degree $d \geq 2$ is nonnegative. (Recall that a polynomial is homogeneous of degree d if all of its monomials have degree exactly d, and that $p(x)$ is nonnegative if $p(x) \geq 0$ for all $x \in \mathbb{R}^n$.)

(c) Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is a norm, then f^2 must be convex.

(d) Suppose $f : \mathbb{R}^n \to \mathbb{R}$ is a norm, then f^2 must be strictly convex.