Princeton University

ORF 363/COS 323
Final Exam, Fall 2020

December 9, 2020

Instructor:
A. A. Ahmadi
Co-Instructor:
S. Shen

AIs:
Chaudhry, Cheng, Dibek, Lee,
Luo, Underwood, Yoosefi

1. Please write your name on the first page of your solutions. Next to it, please write out and sign the following pledge: "I pledge my honor that I have not violated the Honor Code or the rules specified by the instructor during this examination. I have not spent more than 48 hours total on this exam."

2. The exam is not to be discussed with *anyone* except possibly the instructors and the AIs. You can only ask *clarification questions*, and only as *public* (and preferably non-anonymous) questions on Piazza. No emails.

3. You are allowed to consult the lecture notes, live notes, precept notes, your own notes, the reference books of the course as indicated on the syllabus, the problem sets and their solutions (yours and ours), the midterm and its solutions (yours and ours), the practice midterm and final exams and their solutions, all Piazza posts, all videos, but *nothing else*. You can only use the Internet in case you run into problems related to software (e.g. MATLAB or CVX).

4. You may refer to facts proven in the notes or problem sets without reproving them.

5. For all computational problems, include your code. The output that you present should come from your code.

6. Unless you have been granted an extension because of overlapping finals, the exam is to be turned in on Friday (December 11, 2020) at 10 AM EST on Gradescope as a single PDF file. You are free to write your solutions on paper or on a tablet, or to type them up. Only the latest version submitted before your deadline will be graded.

Problem 1 (20 pts): True or False? (Provide a proof or a counterexample.)

(a) Consider the problem of minimizing a quasiconvex function $f : \mathbb{R}^n \to \mathbb{R}$ over a convex set $\Omega \subseteq \mathbb{R}^n$. If a point is a strict local minimum for this problem, then it must be a strict global minimum.

(b) The set of quasiconvex univariate polynomials of degree less than or equal to 3 is convex.

Problem 2 (20 pts): How Martian Was My Valley

You are programming software for the Mars rover. The rover will need to minimize and maximize functions to locate the tops of hills and bottoms of valleys. There is only room for one optimization algorithm in the rover's memory and you must decide between gradient descent and Newton's method. Consider the following function (of a single variable $x \in \mathbb{R}$) which may resemble a Martian valley:

$$f(x) = \sqrt{x^2 + 1}.$$

(a) Show that f is strictly convex and has a strict global minimum.

(b) Consider the gradient descent algorithm with unit step size:

$$x_{k+1} = x_k - f'(x_k).$$

Classify the set of initial points $x_0 \in \mathbb{R}$ that converge to the global minimum of f under the iterations of this algorithm.

(c) Consider Newton's method with unit step size:

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k)}.$$

Classify the set of initial points $x_0 \in \mathbb{R}$ that converge to the global minimum of f under the iterations of this algorithm.

Hint: For parts (b) and (c), you must show that points in your set converge to the global minimum of f and that points outside of your set do not. You may run numerical experiments first to gain some intuition. You can potentially use Lyapunov's theorem when appropriate.

Problem 3 (20 pts): Headphones Were Never an Option

This holiday season you have decided to buy a new speaker, so that next semester you can enjoy listening to your online lectures. You want to use your newly-acquired optimization skills to determine the best location $(x, y) \in \mathbb{R}^2$ for the speaker in your room, and you come up with the following requirements:

- You want to be able to hear the speaker clearly while working, so it must be at most distance 4 from your desk, which is located at $(0, 0)$.
- The speaker has a short power cable, so it must be at most distance 6 from the power outlet, which is located at $(6, -4)$.
- To avoid disturbing your neighbor in the next room, you want to place the speaker as far as possible from the wall. This corresponds to maximizing $(-3x + 4y + 34)/5$.

In this problem we define distances using the usual Euclidean norm in $\mathbb{R}^2$.

(a) Formulate the above requirements as an optimization problem, with decision variables (x, y) corresponding to the location of the speaker. Call this problem (P1). Is (P1) a convex optimization problem? Why or why not?

(b) Formulate (P1) as a semidefinite program (not necessarily in standard form). You may want to consider matrices of the form

$$\begin{pmatrix} r-u & v \\ v & r+u \end{pmatrix},$$

for some appropriate choice of $u, v, r \in \mathbb{R}$.

(c) Use CVX to solve your semidefinite programming version of (P1), and report the best location for the speaker (optimal solution) and its distance from the wall (optimal value). Produce a plot showing the boundaries of the constraints, the feasible region, and the optimal solution (x^*, y^*). Also include on your plot the line $3x - 4y = 34$, representing the wall.

Hint: you can plot circles e.g. using the `viscircles` *function in* MATLAB.

(d) Is the optimal solution unique? Justify your answer.

Problem 4 (20 pts): The Dog Whisperer

Bored during winter break, you decide that you would like to make some extra money by starting your own dog-sitting business. Your premium business plan includes many different types of activities for the dogs, with the goal of maximizing the dogs' collective happiness (to be made more precise below). Each activity consumes different amounts of each dogs' energy per unit time. In addition, all the dogs you are dog-sitting will participate concurrently in each activity because you have to keep an eye on each dog at all times. Given this information, you would like to find out the optimal amount of time you should spend on each activity.

Let the number of different activities and the number of dogs be n and m respectively, and let us denote the time durations dedicated to each activity by $x_1, \ldots, x_n$, measured in hours. Let $A \in \mathbb{R}^{m \times n}$ denote the matrix of per hour energy consumption for the dogs, where A_{ij} denotes the per hour energy consumption for activity j for dog i. Hence, activity j consumes $A_{ij}x_j$ of dog i's energy. Because the total energy consumption is additive, the total energy e_i consumed by dog i is equal to $\sum_{j=1}^{n} A_{ij}x_j$. Further, since the energy level of dog i is limited by some amount $e_i^{max} > 0$, you must respect the requirement that $e_i \leq e_i^{max}$ for $i = 1, \ldots, m$. In addition, you have a maximum of *10 hours* to spend on all of the activities combined.

Being the dog whisperer that you are, you have discovered that for $j = 1, \ldots, n$, the total happiness level that all dogs collectively enjoy as a result of activity j can be approximated by the following piecewise linear function:

$$h_j(x_j) = \begin{cases} r_j x_j, & \text{if } 0 \leq x_j \leq t_j \\ r_j t_j + r_j^{disc}(x_j - t_j), & \text{if } x_j \geq t_j. \end{cases}$$

Here, for each activity j, $r_j > 0$ is the rate of happiness, $r_j^{disc} > 0$ is the discounted rate of happiness, and t_j is the time threshold for the change in happiness rates. We have that $0 < r_j^{disc} < r_j$ (hence modelling the "diminishing returns" phenomenon). The total happiness (TH) is then the sum of all the happiness associated with each activity, i.e., $TH = \sum_{j=1}^{n} h_j(x_j)$. Your goal is to find time durations for each activity to maximize TH while respecting the time budget and the individual dog energy constraints.

(a) Formulate the optimization problem described above as a linear program (LP). Your LP does not have to be in standard form.

(b) Suppose you are in charge of looking after five dogs with differing levels of energy, and the possible activities (in order) are: tug-of-war, fetch, walking, and nosework. The information on energy consumption and energy limits, and the parameters of the happiness functions are given below:

$$A = \begin{bmatrix} 0 & 2 & 3 & 0 \\ 1 & 3 & 4 & 2 \\ 0 & 2 & 3 & 1 \\ 1 & 2 & 3 & 2 \\ 0 & 2 & 3 & 1 \end{bmatrix}, \quad e^{max} = \begin{bmatrix} 25 \\ 30 \\ 20 \\ 20 \\ 15 \end{bmatrix}, \quad r = \begin{bmatrix} 3 \\ 6 \\ 5 \\ 4 \end{bmatrix}, \quad r^{disc} = \begin{bmatrix} 0 \\ 1 \\ 3 \\ 3 \end{bmatrix}, \quad t = \begin{bmatrix} 2 \\ 1 \\ 3 \\ 0.5 \end{bmatrix}.$$

Solve your LP for the given data using CVX. Report your optimal time spent on each activity, as well as the maximum TH achievable.

(c) Give a brief intuitive explanation of the solution you found in part (b), taking the context into account.

Problem 5 (20 pts): Balance-um Ambulances

Having observed that many COVID-19 cases come from community outbreaks, a city's local ambulance operation team would like to coordinate the fleet distribution such that more ambulances are near the neighborhoods with higher predicted demand.

There are five ambulance centers in the city; their coordinates are given as $(x_1, y_1), \ldots, (x_5, y_5)$. Each center currently has s_i ambulances, and is predicted to be in need of d_i ambulances for their designated service neighborhoods ($i = 1, \ldots, 5$). Denote the number of ambulances to be moved from center i to center j by z_{ij}. To simplify the relocation planning, assume that all ambulances are the same and the predicted demands can all be met, i.e. $\sum_{i=1}^{5} s_i \geq \sum_{i=1}^{5} d_i$. The cost associated with moving ambulances around is partially a transportation cost and partially a flat wage. The transportation cost is 10 dollars per unit distance travelled (where the distances between the ambulance centers are measured in the usual Euclidean norm), and the flat wage is 20 dollars per move (regardless of directions and the travel distance).

(a) Write down an optimization problem that solves the ambulance relocation problem; i.e., that finds a minimum cost relocation plan that meets the predicted demands. You are allowed to have constraints that require certain variables to take values among a finite set of integers.

(b) Given the following data, use CVX to solve a linear programming relaxation of the optimization problem you formulated in part (a). Examine the resulting optimal solution and from it retrieve and report an optimal relocation plan. What is the optimal relocation cost?

$$x = \begin{bmatrix} 36 \\ 23 \\ 24 \\ 10 \\ 15 \end{bmatrix}, \quad y = \begin{bmatrix} 20 \\ 30 \\ 56 \\ 15 \\ 5 \end{bmatrix}, \quad s = \begin{bmatrix} 7 \\ 2 \\ 3 \\ 3 \\ 5 \end{bmatrix}, \quad d = \begin{bmatrix} 8 \\ 3 \\ 2 \\ 4 \\ 3 \end{bmatrix}.$$