Name:________________________________

PRINCETON UNIVERSITY

ORF 363/COS 323
Final Exam, Fall 2019

JANUARY 15, 2020

Instructor:
A.A. Ahmadi

AIs:
Dibek, El Khadir, Luo, Silin,
Tang, Yoo, Yu, Zhang

1. Please write out and sign the following pledge on top of the first page of your exam:
   "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. Don't forget to write your name on the exam. Make a copy of your solutions and keep it.

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

4. You are allowed to consult the lecture 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, but *nothing else*. Unless otherwise specified in the statement of a problem, you can only use the Internet in case you run into problems related to MATLAB or CVX.

5. You are allowed to refer to facts proven in the notes or problem sets without reproving them.

6. For all problems involving MATLAB or CVX, show your code. The MATLAB output that you present should come from your code.

7. Unless you have been granted an extension because of overlapping finals, the exam is to be turned in on Friday (January 17, 2020) at 10 AM in the instructor's office (Sherrerd 329). If you cannot make it on Friday and decide to turn in your exam sooner, or if your deadline is different under the rules of the exam, you have to drop your exam off in the ORF 363 box of the ORFE undergraduate lounge (Sherrerd 123). If you do that, you need to *write down the date and time on the first page of your exam and sign it.* You can also submit the exam electronically on Blackboard as a single PDF file.

8. Good luck!

Grading

Problem 1	30 pts	
Problem 2	30 pts	
Problem 3	20 pts	
Problem 4	20 pts	
TOTAL	100	

Problem 1: Optimal spying

At some universities like Princeton, professors have it easy. The students uphold such a high standard of academic integrity that there is no reason to worry about complications around cheating. As you may recall from Problem 1 of our Fall 2017 final however, this is sadly not the case at Cheatston University, where Professor Paranoid is (once again) holding a take-home final exam.

This time around, Professor Paranoid is tired of creating two different sets of exam questions. Instead, he decides to catch the cheaters by secretly monitoring electronic communication among his students. (After all, 2020 is far into the post-privacy era!) His plan is to covertly install surveillance software on his students' cell phones. This, as you can imagine, is a costly and risky thing to do. So Professor Paranoid wants to target as few students as possible while ensuring that any communication among friends is monitored.

More formally, suppose that Professor Paranoid has access to the "friendship network" of his students; this is a graph G with n students as nodes and an edge between two nodes if and only if the two students are friends. If we denote the adjacency matrix[1] of this graph by A, then the optimization problem that Professor Paranoid wants to solve is the following:

$$\begin{aligned} f^* := \min_{x \in \mathbb{R}^n} \quad & \sum_{i=1}^n x_i \\ \text{s.t.} \quad & x_i(x_i - 1) = 0, \ i = 1, \ldots, n, \\ & x_i + x_j \geq 1, \ \text{if } A_{ij} = 1. \end{aligned} \tag{1}$$

(a) We say that a friendship is *"monitored"* if the surveillance software is installed on the cell phone of at least one of the two students involved in this friendship. Argue why the optimal value f^* to problem (1) is equal to the minimum number of students who should be targeted by the surveillance software in order to ensure that all friendships in the class are monitored. Show that the above problem is not a convex optimization problem.[2]

It turns out that problem (1) is NP-hard to solve. Nevertheless, you know from ORF 363 that convex optimization can approximately solve NP-hard problems and find useful bounds on their optimal value. Let's see exactly how.

[1] Recall that the adjacency matrix is a symmetric matrix that has zeros on the diagonal and whose (i, j) entry equals 1 if students i and j are friends and 0 otherwise.

[2] For a more serious application of problem (1), think of replacing "Professor Paranoid" with the FBI and the "friendship network" with some criminal network.

(b) *Lower bounding by LP.* Consider the linear program

$$
\begin{aligned}
LB^{LP} := \min_{x \in \mathbb{R}^n} \quad & \sum_{i=1}^{n} x_i \\
\text{s.t.} \quad & 0 \leq x_i \leq 1, \; i = 1, \ldots, n, \\
& x_i + x_j \geq 1, \; \text{if } A_{ij} = 1.
\end{aligned}
\tag{2}
$$

Argue why we must have

$$LB^{LP} \leq f^*.$$

(c) *Lower bounding by SDP.* Consider the semidefinite program

$$
\begin{aligned}
f^{SDP} := \max_{X \in S^{n \times n}} \quad & \mathrm{Tr}(JX) \\
\text{s.t.} \quad & X_{ij} = 0, \; \text{if } A_{ij} = 1, \\
& \mathrm{Tr}(X) = 1, \\
& X \succeq 0,
\end{aligned}
\tag{3}
$$

where $S^{n \times n}$ denotes the set of real symmetric $n \times n$ matrices, J is the all-ones matrix, and Tr stands for the trace operation. Let $LB^{SDP} := n - f^{SDP}$. Show that

$$LB^{SDP} \leq f^*.$$

(d) *Upper bounding by rounding the LP solution.* Professor Paranoid installs the surveillance software according to the strategy[3]

$$
\hat{x}_i = \begin{cases} 1 & \text{if } x_i^{LP} \geq \frac{1}{2} \\ 0 & \text{if } x_i^{LP} < \frac{1}{2}, \end{cases} \quad \text{for } i = 1, \ldots, n,
$$

where x^{LP} is an optimal solution to the LP in (2). Let $UB^{LP} := \sum_{i=1}^{n} \hat{x}_i$. Show the following two inequalities:

$$f^* \leq UB^{LP} \leq 2f^*.$$

(e) *Numerical experiments.* To examine the quality of these bounds, we are going to compute three quantities on some random instances of this problem:

$$LB^{SDP} - LB^{LP}, \; \frac{UB^{LP}}{LB^{LP}}, \; \frac{UB^{LP}}{LB^{SDP}}.$$

[3] In other words, student i is a target for the surveillance software if and only if $\hat{x}_i$=1.

Suppose there are 20 students in the class and that any two randomly chosen students are friends with each other with probability 0.2. Generate 50 random instances of such a friendship network by running the following code (50 times):

```
n=20; p=.2; A=zeros(n);
for i=1:n-1
    for j=i+1:n
        h=rand; if h<p
            A(i,j)=1;A(j,i)=1;
        end
    end
end
```

In each run, compute[4] $LB^{SDP} - LB^{LP}$, $\frac{UB^{LP}}{LB^{LP}}$ and $\frac{UB^{LP}}{LB^{SDP}}$. Plot the histogram of each of these three quantities over your 50 runs. Briefly comment on any observations you may have.

[4]To compute $\hat{x}$, we suggest that you add 0.0001 to each entry of the solution x^{LP} that the solver returns before rounding it. This is because some entries of x^{LP} can in reality be equal to $\frac{1}{2}$, but be returned as 0.4999 by the numerical solver.

Problem 2: Convergence of Newton's method
Consider the univariate function[5] $f(x) = 20x \arctan(x) - 10\log(1+x^2) + x^2$.

(a) Is f convex? Strictly convex? Why or why not?

(b) What are the global minima of f? Show your reasoning.

(c) Is Newton's method for minimizing f globally convergent? In other words, do all initial points $x_0 \in \mathbb{R}$ converge to a global minimum of f under the iterations of the Newton method? Why or why not?

Hint: To answer this question, you may want to go through the following steps:

- On your computer, run Newton's method from various initial conditions to get a feel for its behavior. Comment on your numerical observations.
- Let $N(x) = x - \frac{f'(x)}{f''(x)}$ (this is the "Newton map"). Show that $N'(x) \leq 0, \forall x \in \mathbb{R}$.
- Show that $|x| \geq 2 \implies |N(x)| \geq 2$. You may want to start by showing that $x \geq 2 \implies N(x) \leq -2$.

Problem 3: SDPs with rational data but no rational feasible solution
In a previous problem set, we saw that unlike linear programs, semidefinite programs with rational data may have irrational optimal solutions. In this problem, you are asked to show that the same phenomenon can occur for feasible solutions.
Give an example of symmetric $n \times n$ matrices $A_1, \ldots, A_m$ with rational entries and rational numbers $b_1, \ldots, b_m$ such that the set

$$S := \{X \in S^{n\times n} | \operatorname{Tr}(A_i X) = b_i, i = 1, \ldots, m, X \succeq 0\}$$

is non-empty, but only contains matrices that have at least one irrational entry. Here, $S^{n\times n}$ denotes the set of symmetric $n \times n$ matrices with real entries and Tr stands for the trace operation. You can pick any value for n and m that you like as long as the above requirements are met.
(*Hint*: You may want to first impose positive semidefiniteness constraints on two 2×2 decision matrices and then convert them into a positive semidefiniteness constraint on a single 4×4 decision matrix.)

[5]Here, $\log(x)$ is the natural logarithm and $\arctan(x)$ is the inverse of the $\tan(x) = \frac{\sin(x)}{cos(x)}$ function. You are allowed to look up the derivatives of these functions if you do not remember them.

Problem 4: Putting the F in ORFE

Last year, your friend Finn Anshall graduated from Princeton with \$10,000 in his pocket and a desire to use the skills he learned from his optimization courses to make some money. He recalled learning about the following optimization problem,[6] which determines in what fractions one's money should be invested in n given stocks so the resulting portfolio gives maximum return while having limited risk (or variance):

$$\begin{aligned} \max_{x \in \mathbb{R}^n} \quad & \mu^T x \\ \text{s.t.} \quad & x^T \Sigma x \leq \delta, \\ & x_i \geq 0, i = 1, \ldots, n, \\ & \sum_{i=1}^{n} x_i = 1. \end{aligned} \tag{4}$$

Here, $\delta \in \mathbb{R}$ is a risk-aversion parameter. To describe what μ and Σ are, let us first talk about returns. The return of a particular stock in a particular month is computed as

$$\frac{\text{Stock price at end of month} - \text{Stock price at end of previous month}}{\text{Stock price at end of previous month}}.$$

Let $r_j \in \mathbb{R}^n$ denote the vector of returns of all stocks in a particular month j. If we have access to past return vectors over a period of m months, then $\mu \in \mathbb{R}^n$ is the vector of average monthly returns; i.e., $\mu = \frac{1}{m}\sum_{j=1}^{m} r_j$. The symmetric $n \times n$ matrix Σ is the so-called empirical covariance matrix[7] of the monthly returns. It is computed as

$$\Sigma = \frac{1}{m-1}\sum_{j=1}^{m}(r_j - \mu)(r_j - \mu)^T.$$

(a) Is (4) a convex optimization problem? Why or why not?

(b) Finn looked up on Yahoo! Finance the monthly closing prices of five stocks—Apple (AAPL), Coca Cola (COKE), Google (GOOG), Alibaba (BABA), and CMC (CMC)—from September 2014 to August 2018. He then followed the procedure described above to compute μ and Σ, which can be found in the file `StockData.mat` (using the command

[6] This model was proposed by Harry Markowitz in 1952 and earned him a Nobel Prize in Economics in 1990. For simplicity, we are ignoring transaction costs or the possibility of shorting a stock, though the model can be easily adjusted to allow for these.

[7] The off-diagonal entries of this matrix indicate whether two stocks are likely to move up or down in price together.

`load StockData`). Solve problem (4) using this μ and Σ and with $\delta = .003$. Report Finn's optimal solution and optimal value.

(c) On the last Friday of September 2018, Finn invested his \$10,000 according to the optimal solution of (4) you just obtained. Find out how much his portfolio is worth now (i.e., on January 14, 2020). The stock prices on these two dates are available in `StockData.mat` and are copied below:

Stock	September 28, 2018	January 14, 2020
Apple	\$221.63	\$312.68
Coca Cola	\$181.33	\$277.43
Google	\$1193.47	\$1430.88
Alibaba	\$164.76	\$226.49
CMC	\$19.86	\$22.22

Compare this to the portfolio of Finn's roommate, who did not take any optimization courses and split his \$10,000 evenly across the five stocks on September 28, 2018. How much is the roommate's portfolio worth on January 14, 2020?