Name:________________________________

Princeton University

ORF 363/COS 323
Final Exam, Fall 2017

January 17, 2018

Instructor:
A.A. Ahmadi

AIs:
B. El Khadir, C. Dibek, G. Hall, J. Zhang, J. Ye, S. Uysal

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."

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*. 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 19, 2018) 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	25 pts	
Problem 2	25 pts	
Problem 3	25 pts	
Problem 4	25 pts	
TOTAL	100	

Problem 1: Minimizing the probability of cheating

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. This is sadly not the case at Cheatston University, where Professor Paranoid is about to hold a take-home final exam.

Because he is worried about cheating, Professor Paranoid goes through the trouble of designing two different sets of exam questions of equal difficulty level. He is now faced with the task of deciding which students should get Exam A and which should get Exam B. To make this decision, Professor Paranoid's strategy is to make sure that to the extent possible, students who are friends with each other get different exams.

More formally, suppose that Professor Paranoid has access to the "friendship network" of his students; this is a graph $G(V, E)$ 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:

$$\begin{aligned} f^* := \max_{x \in \mathbb{R}^n} \quad & \frac{1}{4} \sum_{i,j} A_{ij}(1 - x_i x_j) \\ \text{s.t.} \quad & x_i^2 = 1, \; i = 1, \ldots, n. \end{aligned} \tag{1}$$

(a) We say that a friendship is *"cheat-free"* if the two students involved in this friendship get different exams. Argue why the optimal value f^* to problem (1) is equal to the maximum possible number of cheat-free friendships. Show that the above problem is not a convex optimization problem.

It turns out that problem (1) is NP-hard to solve. Nevertheless, you know from ORF 363 that semidefinite programming is a powerful tool for approximately solving NP-hard problems. Let's see exactly how.

(b) *Upper bounding.* Consider the semidefinite program

$$\begin{aligned} f^{SDP} := \min_{X \in S^{n \times n}} \quad & Tr(AX) \\ \text{s.t.} \quad & X_{ii} = 1, \; i = 1, \ldots, n, \\ & X \succeq 0. \end{aligned} \tag{2}$$

[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.

Let $UB^{SDP} := \frac{1}{4}\sum_{i,j} A_{ij} - \frac{1}{4} f^{SDP}$. Show that

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

(c) *Lower bounding.* To find a suboptimal solution to (1), Professor Paranoid gives out the exams according to the following strategy[2]:

$$\hat{x}_i = \mathtt{sign}(X^*(i,1)), \quad i = 1, \ldots, n,$$

where X^* is the optimal solution to the SDP in (2) that the solver returns.[3] The objective value of (1) at $\hat{x}$ clearly gives a lower bound on f^*. Let's call this lower bound LB^{SDP}. So we have

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

To examine the quality of these bounds, we are going to compute the ratio $\frac{LB^{SDP}}{UB^{SDP}}$ on some random instances of this problem. (If this ratio is 1, we have solved problem (1) exactly.) Suppose there are 70 students in the class and that any two randomly chosen students are friends with each other with probability 0.3. Generate 50 random instances of such a friendship network by running the following code (50 times):

```
n=70; p=.3; 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 $\frac{LB^{SDP}}{UB^{SDP}}$ and then show the histogram of this ratio over your 50 runs.

Problem 2: True or False? (Provide a proof or a counterexample.)

(a) For $k = 0, 1, \ldots,$ consider the iterations $x_{k+1} = x_k + \alpha_k d_k$ used to minimize a function $f : \mathbb{R}^n \to \mathbb{R}$. If the directions d_k are descent directions for all k, the step size is set to $\alpha_k = \frac{1}{2^k}$, and the function f is convex, then the iterations satisfy $f(x_{k+1}) \leq f(x_k), \forall k$.

[2] Here, the function $\mathtt{sign}(z)$ returns $+1$ if $z \geq 0$ and -1 otherwise.

[3] If you are curious, this heuristic is motivated by the fact that if the upper bound in part (b) is tight, then there is always a rank-1 optimal solution to the SDP whose first column (or any other column actually) is a ± 1 vector that is optimal for (1).

(b) For $k = 0, 1, \ldots,$ consider the iterations $x_{k+1} = x_k + \alpha_k d_k$ used to minimize a differentiable function $f : \mathbb{R}^n \to \mathbb{R}$. If the directions d_k are descent directions for all k, the function f is convex and nonnegative, and the step sizes are chosen to guarantee that $f(x_{k+1}) < f(x_k)$ for all k, then the sequence $\{x_k\}$ converges to a point $x^* \in \mathbb{R}^n$ that satisfies $\nabla f(x^*) = 0$.

(c) Consider the following SDP:

$$\begin{aligned} \min_{X \in S^{n\times n}} \quad & \mathrm{Tr}(CX) \\ \text{s.t.} \quad & \mathrm{Tr}(A_i X) = b_i,\ i = 1, \ldots, m, \\ & X \succeq 0. \end{aligned} \tag{3}$$

Let $b := (b_1, \ldots, b_m)^T$, and suppose there exists a vector $y \in \mathbb{R}^m$ such that $b^T y = 0$ and $\sum_{i=1}^m A_i y_i \preceq C$. Then, there exists a matrix $X^* \in S^{n\times n}$ that is feasible to (3) and such that

$$Tr(CX^*) \leq Tr(CX),$$

for any matrix X that is feasible to (3).

Problem 3: Nearest correlation matrix

You are the CEO of HoneyMoney Technologies LLC, a new hedge fund firm in NYC whose proprietary optimization algorithms has Wall Street raving. Your main competitor, RenaissancE Technologies[4], has sent in a spy, disguised as a summer intern, to interfere with your investments. The spy has gotten his hands on your *correlation matrix* C of n important stocks,[5] to which he has added some random noise, leaving you with a matrix $\hat{C}$. We remark that to be a valid correlation matrix, a matrix must be symmetric, positive semidefinite, and have all diagonal entries equal to one. The spy has been careful enough to make sure that the resulting matrix $\hat{C}$ is symmetric and has ones on the diagonal, but he hasn't noticed that his change has made $\hat{C}$ not positive semidefinite.

[4]Not to be confused with Renaissance Technologies that would never do such a thing.

[5]If you are curious, the correlation matrix is an $n \times n$ symmetric matrix used frequently in investment banking. Its (i, j)-th entry is a number between -1 and 1, with numbers close to 1 meaning that stocks i and j are likely to move up together, close to -1 meaning that the two stocks are likely to move in opposite directions, and close to zero meaning that they are likely uncorrelated. The problem of finding the closest correlation matrix to a given matrix is an important problem in financial engineering; see e.g. this article.

(a) Suppose we have

$$\hat{C} = \begin{pmatrix} 1.00 & -0.76 & 0.07 & -0.96 \\ -0.76 & 1.00 & 0.18 & 0.07 \\ 0.07 & 0.18 & 1.00 & 0.41 \\ -0.96 & 0.07 & 0.41 & 1.00 \end{pmatrix}.$$

Recover the original matrix by finding the nearest correlation matrix to $\hat{C}$ in Frobenius norm[6] (i.e., the correlation matrix C that minimizes $||C - \hat{C}||_F$). Give your optimal solution.

(b) Show that for any symmetric matrix $\hat{C}$, the problem of finding the closest correlation matrix to $\hat{C}$ in Frobenius norm has a unique solution.

(c) Suppose $\hat{C}$ is symmetric, has ones on the diagonal, but is not positive semidefinite. Show that the optimal solution to the problem of finding the closest correlation matrix in Frobenius norm to $\hat{C}$ must have at least one zero eigenvalue. (*Hint:* You may use the fact that the minimum eigenvalue of a symmetric matrix is a continuous function of its entries.)

Problem 4: Nash equilibria in 2-player games

A *bimatrix game* is a game between two players each having a finite number of strategies. If Player 1 has m strategies and Player 2 has n strategies, then the game is fully defined by two real $m \times n$ *payoff matrices* A and B. When Player 1 plays strategy i and Player 2 plays strategy j, then the payoffs that they get are respectively A_{ij} and B_{ij}.

In the case where Player 1 chooses to only play strategy i, we represent her choice by a vector $x \in \mathbb{R}^m$ which consists of all zeros except for a 1 in the i^{th} position. This is called a *"pure strategy"*. However, Player 1 can also choose to play a *"mixed strategy"*, where she chooses to play a *convex combination* of her m strategies. In either case, Player 1's strategy is always an element of the "unit simplex", which by definition is the following set:

$$\Delta_m = \left\{ x \in \mathbb{R}^m | \sum_{i=1}^{m} x_i = 1, x_i \geq 0 \right\}.$$

Likewise, Player 2's strategy will be an element of Δ_n. Under a pair of mixed strategies $(x, y) \in \Delta_m \times \Delta_n$, the payoff of Player 1 is $x^T A y$ and the payoff of Player 2 is $x^T B y$.

[6] Recall what the Frobenius norm is from a previous problem set.

A pair of strategies (x^*, y^*) constitutes a Nash equilibrium if no player can improve his/her payoff by a unilateral deviation:

$$\begin{aligned} &x^* \in \Delta_m, \\ &y^* \in \Delta_n, \\ &{x^*}^T A y^* \geq x^T A y^*, \ \forall x \in \Delta_m, \\ &{x^*}^T B y^* \geq {x^*}^T B y, \ \forall y \in \Delta_n. \end{aligned} \tag{4}$$

This means that if Player 2 sticks to his strategy y^*, then there is no incentive for Player 1 to change her strategy from x^* to some other strategy, and, conversely, if Player 1 sticks to her strategy x^*, then there is no incentive for Player 2 to change his strategy from y^* to some other strategy. Hence, we are at some sort of an "equilibrium" (Trump will not bomb Kim Jong-un, Kim Jong-un will not bomb Trump, and life will remain good).
Nash's big result was to prove that independent of A and B, a (Nash) equilibrium always exists.[7] His paper however, did not give an algorithm for finding such equilibria.

(a) Show that a pair of strategies (x^*, y^*) is a Nash equilibrium if and only if it satisfies

$$\begin{aligned} &x^* \in \Delta_m, \\ &y^* \in \Delta_n, \\ &{x^*}^T A y^* \geq e_i^T A y^*, \ i = 1, \ldots, m, \\ &{x^*}^T B y^* \geq {x^*}^T B e_i, \ i = 1, \ldots, n, \end{aligned} \tag{5}$$

where e_i is a vector (in $\mathbb{R}^m$ or $\mathbb{R}^n$) with only zeros, except for a 1 in the i^{th} position.

(b) An important class of bimatrix games is the so-called *zero-sum* games. These are games whose payoff matrices satisfy $B = -A$. This means that the two players are in direct competition (any win for Player 1 is a loss for Player 2 and vice versa). For example, the classic "rock-paper-scissors" game is a zero-sum game. Show that in a zero-sum game, a pair of strategies (x^*, y^*) is a Nash equilibrium if and only if it satisfies

$$x^* \in \Delta_m, y^* \in \Delta_n, \qquad {x^*}^T A e_j \geq e_i^T A y^*, \forall \{i, j\} \in \{1, \ldots, m\} \times \{1, \ldots, n\}.$$

(c) Consider the zero-sum game given by

$$A = \begin{pmatrix} 1 & 3 & -3 \\ 0 & 1 & 2 \\ 3 & -1 & -1 \end{pmatrix}, B = -A.$$

Find a Nash equilibrium of this game. Is the Nash equilibrium unique? Justify.

[7] In fact, he proved the result more generally for games among a finite number of players.