

## Final Exam

This is a 24 hour take-home final. Please turn it in at Bytes Cafe in the Packard building, 24 hours after you pick it up.

You may use any books, notes, or computer programs (*e.g.*, Matlab, CVX), but you may not discuss the exam with anyone until March 18, after everyone has taken the exam. The only exception is that you can ask us for clarification, via the course staff email address. We've tried pretty hard to make the exam unambiguous and clear, so we're unlikely to say much.

Please make a copy of your exam before handing it in.

**Please attach the cover page to the front of your exam.** Assemble your solutions in order (problem 1, problem 2, problem 3, ...), starting a new page for each problem. Put everything associated with each problem (*e.g.*, text, code, plots) together; do not attach code or plots at the end of the final.

**We will deduct points from long needlessly complex solutions, even if they are correct.** Our solutions are not long, so if you find that your solution to a problem goes on and on for many pages, you should try to figure out a simpler one. We expect neat, legible exams from everyone, including those enrolled Cr/N.

When a problem involves computation you must give all of the following: a clear discussion and justification of exactly what you did, the Matlab source code that produces the result, and the final numerical results or plots.

To download Matlab files containing problem data, you'll have to type the whole URL given in the problem into your browser. To get a file called `filename.m`, for example, you would retrieve

```
http://www.stanford.edu/class/ee364a/data_for_final/filename.m
```

with your browser.

All problems have equal weight. Some are easy. Others, not so much.

Be sure you are using the most recent version of CVX, which is Version 2.1, build 1074. You can check this using the command `cvx_version`. Check your email often during the exam, just in case we need to send out an important announcement.

Some problems involve applications. But you do not need to know *anything* about the problem area to solve the problem; the problem statement contains everything you need.

1. *Lightest structure that resists a set of loads.* We consider a mechanical structure in 2D (for simplicity) which consists of a set of  $m$  nodes, with known positions  $p_1, \dots, p_m \in \mathbf{R}^2$ , connected by a set of  $n$  bars (also called struts or elements), with cross-sectional areas  $a_1, \dots, a_n \in \mathbf{R}_+$ , and internal tensions  $t_1, \dots, t_n \in \mathbf{R}$ .

Bar  $j$  is connected between nodes  $r_j$  and  $s_j$ . (The indices  $r_1, \dots, r_n$  and  $s_1, \dots, s_n$  give the structure topology.) The length of bar  $j$  is  $L_j = \|p_{r_j} - p_{s_j}\|_2$ , and the total volume of the bars is  $V = \sum_{j=1}^n a_j L_j$ . (The total weight is proportional to the total volume.)

Bar  $j$  applies a force  $(t_j/L_j)(p_{r_j} - p_{s_j}) \in \mathbf{R}^2$  to node  $s_j$  and the negative of this force to node  $r_j$ . Thus, positive tension in a bar pulls its two adjacent nodes towards each other; negative tension (also called compression) pushes them apart. The ratio of the tension in a bar to its cross-sectional area is limited by its yield strength, which is symmetric in tension and compression:  $|t_j| \leq \sigma a_j$ , where  $\sigma > 0$  is a known constant that depends on the material.

The nodes are divided into two groups: free and fixed. We will take nodes  $1, \dots, k$  to be free, and nodes  $k+1, \dots, m$  to be fixed. Roughly speaking, the fixed nodes are firmly attached to the ground, or a rigid structure connected to the ground; the free ones are not.

A *loading* consists of a set of external forces,  $f_1, \dots, f_k \in \mathbf{R}^2$  applied to the free nodes. Each free node must be in equilibrium, which means that the sum of the forces applied to it by the bars and the external force is zero. The structure can *resist* a loading (without collapsing) if there exists a set of bar tensions that satisfy the tension bounds and force equilibrium constraints. (For those with knowledge of statics, these conditions correspond to a structure made entirely with pin joints.)

Finally, we get to the problem. You are given a set of  $M$  loadings, *i.e.*,  $f_1^{(i)}, \dots, f_k^{(i)} \in \mathbf{R}^2$ ,  $i = 1, \dots, M$ . The goal is to find the bar cross-sectional areas that minimize the structure volume  $V$  while resisting all of the given loadings. (Thus, you are to find *one* set of bar cross-sectional areas, and  $M$  sets of tensions.) Using the problem data provided in `lightest_struct_data.m`, report  $V^*$  and  $V^{\text{unif}}$ , the smallest feasible structure volume when all bars have the same cross-sectional area. The node positions are given as a  $2 \times m$  matrix  $\mathbf{P}$ , and the loadings as a  $2 \times k \times M$  array  $\mathbf{F}$ . Use the code included in the data file to visualize the structure with the bar cross-sectional areas that you find, and provide the plot in your solution.

*Hint.* You might find the graph incidence matrix  $A \in \mathbf{R}^{m \times n}$  useful. It is defined as

$$A_{ij} = \begin{cases} +1 & i = r_j \\ -1 & i = s_j \\ 0 & \text{otherwise.} \end{cases}$$

*Remark.* You could reasonably ask, ‘Does a mechanical structure really solve a convex optimization problem to determine whether it should collapse?’. It sounds odd, but the answer is, yes it does.

2. *Theory-applications split in a course.* A professor teaches an advanced course with 20 lectures, labeled  $i = 1, \dots, 20$ . The course involves some interesting theoretical topics, and many practical applications of the theory. The professor must decide how to split each lecture between theory and applications. Let  $T_i$  and  $A_i$  denote the fraction of the  $i$ th lecture devoted to theory and applications, for  $i = 1, \dots, 20$ . (We have  $T_i \geq 0$ ,  $A_i \geq 0$ , and  $T_i + A_i = 1$ .)

A certain amount of theory has to be covered before the applications can be taught. We model this in a crude way as

$$A_1 + \dots + A_i \leq \phi(T_1 + \dots + T_i), \quad i = 1, \dots, 20,$$

where  $\phi : \mathbf{R} \rightarrow \mathbf{R}$  is a given nondecreasing function. We interpret  $\phi(u)$  as the cumulative amount of applications that can be covered, when the cumulative amount of theory covered is  $u$ . We will use the simple form  $\phi(u) = a(u - b)_+$ , with  $a, b > 0$ , which means that no applications can be covered until  $b$  lectures of the theory is covered; after that, each lecture of theory covered opens the possibility of covering  $a$  lectures on applications.

The theory-applications split affects the emotional state of students differently. We let  $s_i$  denote the emotional state of a student after lecture  $i$ , with  $s_i = 0$  meaning neutral,  $s_i > 0$  meaning happy, and  $s_i < 0$  meaning unhappy. Careful studies have shown that  $s_i$  evolves via a linear recursion (dynamics)

$$s_i = (1 - \theta)s_{i-1} + \theta(\alpha T_i + \beta A_i), \quad i = 1, \dots, 20,$$

with  $s_0 = 0$ . Here  $\alpha$  and  $\beta$  are parameters (naturally interpreted as how much the student likes or dislikes theory and applications, respectively), and  $\theta \in [0, 1]$  gives the emotional volatility of the student (*i.e.*, how quickly he or she reacts to the content of recent lectures). The student's terminal emotional state is  $s_{20}$ .

Now consider a specific instance of the problem, with course material parameters  $a = 2$ ,  $b = 3$ , and three groups of students, with emotional dynamics parameters given as follows.

	Group 1	Group 2	Group 3
$\theta$	0.05	0.1	0.3
$\alpha$	-0.1	0.8	-0.3
$\beta$	1.4	-0.3	0.7

Find (four different) theory-applications splits that maximize the terminal emotional state of the first group, the terminal emotional state of the second group, the terminal emotional state of the third group, and, finally, the minimum of the terminal emotional states of all three groups.

For each case, plot  $T_i$  and the emotional state  $s_i$  for the three groups, versus  $i$ . Report the numerical values of the terminal emotional states for each group, for each of the four theory-applications splits.

3. *Optimal electric motor drive currents.* In this problem you will design the drive current waveforms for an AC (alternating current) electric motor. The motor has a magnetic rotor which spins with constant angular velocity  $\omega \geq 0$  inside the stationary stator. The stator contains three circuits (called *phase windings*) with (vector) current waveform  $i : \mathbf{R} \rightarrow \mathbf{R}^3$  and (vector) voltage waveform  $v : \mathbf{R} \rightarrow \mathbf{R}^3$ , which are  $2\pi$ -periodic functions of the angular position  $\theta$  of the rotor. The circuit dynamics are

$$v(\theta) = Ri(\theta) + \omega L \frac{d}{d\theta} i(\theta) + \omega k(\theta),$$

where  $R \in \mathbf{S}_{++}^3$  is the resistance matrix,  $L \in \mathbf{S}_{++}^3$  is the inductance matrix, and  $k : \mathbf{R} \rightarrow \mathbf{R}^3$ , a  $2\pi$ -periodic function of  $\theta$ , is the back-EMF waveform (which encodes the electromagnetic coupling between the rotor permanent magnets and the phase windings). The angular velocity  $\omega$ , the matrices  $R$  and  $L$ , and the back-EMF waveform  $k$ , are known.

We must have  $|v_i(\theta)| \leq v^{\text{supply}}$ ,  $i = 1, 2, 3$ , where  $v^{\text{supply}}$  is the (given) supply voltage. The output torque of the motor at rotor position  $\theta$  is  $\tau(\theta) = k(\theta)^T i(\theta)$ . We will require the torque to have a given constant nonnegative value:  $\tau(\theta) = \tau^{\text{des}}$  for all  $\theta$ .

The average power loss in the motor is

$$P^{\text{loss}} = \frac{1}{2\pi} \int_0^{2\pi} i(\theta)^T Ri(\theta) d\theta.$$

The mechanical output power is  $P^{\text{out}} = \omega \tau^{\text{des}}$ , and the motor efficiency is

$$\eta = P^{\text{out}} / (P^{\text{out}} + P^{\text{loss}}).$$

The objective is to choose the current and voltage waveforms to maximize  $\eta$ .

*Discretization.* To solve this problem we consider a discretized version in which  $\theta$  takes on the  $N$  values  $\theta = h, 2h, \dots, Nh$ , where  $h = 2\pi/N$ . We impose the voltage and torque constraints for these values of  $\theta$ . We approximate the power loss as

$$P^{\text{loss}} = (1/N) \sum_{j=1}^N i(jh)^T Ri(jh).$$

The circuit dynamics are approximated as

$$v(jh) = Ri(jh) + \omega L \frac{i((j+1)h) - i(jh)}{h} + \omega k(jh), \quad j = 1, \dots, N,$$

where here we take  $i((N+1)h) = i(h)$  (by periodicity).

Find optimal (discretized) current and voltage waveforms for the problem instance with data given in `ac_motor_data.m`. The back-EMF waveform is given as a  $3 \times N$  matrix  $K$ . Plot the three current waveform components on one plot, and the three voltage waveforms on another. Give the efficiency obtained.

4. *Multi-label support vector machine.* The basic SVM described in the book is used for classification of data with two labels. In this problem we explore an extension of SVM that can be used to carry out classification of data with more than two labels. Our data consists of pairs  $(x_i, y_i) \in \mathbf{R}^n \times \{1, \dots, K\}$ ,  $i = 1, \dots, m$ , where  $x_i$  is the feature vector and  $y_i$  is the label of the  $i$ th data point. (So the labels can take the values  $1, \dots, K$ .) Our classifier will use  $K$  affine functions,  $f_k(x) = a_k^T x + b_k$ ,  $k = 1, \dots, K$ , which we also collect into affine function from  $\mathbf{R}^n$  into  $\mathbf{R}^K$  as  $f(x) = Ax + b$ . (The rows of  $A$  are  $a_k^T$ .) Given feature vector  $x$ , we guess the label  $\hat{y} = \operatorname{argmax}_k f_k(x)$ . We assume that exact ties never occur, or if they do, an arbitrary choice can be made. Note that if a multiple of  $\mathbf{1}$  is added to  $b$ , the classifier does not change. Thus, without loss of generality, we can assume that  $\mathbf{1}^T b = 0$ .

To correctly classify the data examples, we need  $f_{y_i}(x_i) > \max_{k \neq y_i} f_k(x_i)$  for all  $i$ . This is a set of homogeneous strict inequalities in  $a_k$  and  $b_k$ , which are feasible if and only if the set of nonstrict inequalities  $f_{y_i}(x_i) \geq 1 + \max_{k \neq y_i} f_k(x_i)$  are feasible. This motivates the loss function

$$L(A, b) = \sum_{i=1}^m \left( 1 + \max_{k \neq y_i} f_k(x_i) - f_{y_i}(x_i) \right)_+,$$

where  $(u)_+ = \max\{u, 0\}$ . The multi-label SVM chooses  $A$  and  $b$  to minimize

$$L(A, b) + \mu \|A\|_F^2,$$

subject to  $\mathbf{1}^T b = 0$ , where  $\mu > 0$  is a regularization parameter. (Several variations on this are possible, such as regularizing  $b$  as well, or replacing the Frobenius norm squared with the sum of norms of the columns of  $A$ .)

- (a) Show how to find  $A$  and  $b$  using convex optimization. Be sure to justify any changes of variables or reformulation (if needed), and convexity of the objective and constraints in your formulation.
- (b) Carry out multi-label SVM on the data given in `multi_label_svm_data.m`. Use the data given in `X` and `y` to fit the SVM model, for a range of values of  $\mu$ . This data set includes an additional set of data, `Xtest` and `ytest`, that you can use to test the SVM models. Plot the test set classification error rate (*i.e.*, the fraction of data examples in the test set for which  $\hat{y} \neq y$ ) versus  $\mu$ .

You don't need to try more than 10 or 20 values of  $\mu$ , and we suggest choosing them uniformly on a log scale, from (say)  $10^{-2}$  to  $10^2$ .

5. *De-leveraging.* We consider a multi-period portfolio optimization problem, with  $n$  assets and  $T$  time periods, where  $x_t \in \mathbf{R}^n$  gives the holdings (say, in dollars) at time  $t$ , with negative entries denoting, as usual, short positions. For each time period the return vector has mean  $\mu \in \mathbf{R}^n$  and covariance  $\Sigma \in \mathbf{S}_{++}^n$ . (These are known.)

The initial portfolio  $x_0$  maximizes the risk-adjusted expected return  $\mu^T x - \gamma x^T \Sigma x$ , where  $\gamma > 0$ , subject to the leverage limit constraint  $\|x\|_1 \leq L^{\text{init}}$ , where  $L^{\text{init}} > 0$  is the given initial leverage limit. (There are several different ways to measure leverage; here we use the sum of the total short and long positions.) The final portfolio  $x_T$  maximizes the risk-adjusted return, subject to  $\|x\|_1 \leq L^{\text{new}}$ , where  $L^{\text{new}} > 0$  is the given final leverage limit (with  $L^{\text{new}} < L^{\text{init}}$ ). This uniquely determines  $x_0$  and  $x_T$ , since the objective is strictly concave.

The question is how to move from  $x_0$  to  $x_T$ , *i.e.*, how to choose  $x_1, \dots, x_{T-1}$ . We will do this so as to maximize the objective

$$J = \sum_{t=1}^T (\mu^T x_t - \gamma x_t^T \Sigma x_t - \phi(x_t - x_{t-1})),$$

which is the total risk-adjusted expected return, minus the total transaction cost. The transaction cost function  $\phi$  has the form

$$\phi(u) = \sum_{i=1}^n (\kappa_i |u_i| + \lambda_i u_i^2),$$

where  $\kappa \succeq 0$  and  $\lambda \succeq 0$  are known parameters. We will require that  $\|x_t\|_1 \leq L^{\text{init}}$ , for  $t = 1, \dots, T-1$ . In other words, the leverage limit is the initial leverage limit up until the deadline  $T$ , when it drops to the new lower value.

- (a) Explain how to find the portfolio sequence  $x_1^*, \dots, x_{T-1}^*$  that maximizes  $J$  subject to the leverage limit constraints.
- (b) Find the optimal portfolio sequence  $x_t^*$  for the problem instance with data given in `deleveraging_data.m`. Compare this sequence with two others:  $x_t^{\text{lp}} = x_0$  for  $t = 1, \dots, T-1$  (*i.e.*, one that does all trading at the last possible period), and the linearly interpolated portfolio sequence

$$x_t^{\text{lin}} = (1 - t/T)x_0 + (t/T)x_T, \quad t = 1, \dots, T-1.$$

For each of these three portfolio sequences, give the objective value obtained, and plot the risk and transaction cost adjusted return,

$$\mu^T x_t - \gamma x_t^T \Sigma x_t - \phi(x_t - x_{t-1}),$$

and the leverage  $\|x_t\|_1$ , versus  $t$ , for  $t = 0, \dots, T$ . Also, for each of the three portfolio sequences, generate a single plot that shows how the holdings  $(x_t)_i$  of the  $n$  assets change over time, for  $i = 1, \dots, n$ .

Give a *very short* (one or two sentence) intuitive explanation of the results.

6. *Worst-case variance.* Suppose  $Z$  is a random variable on  $\mathbf{R}^n$  with covariance matrix  $\Sigma \in \mathbf{S}_+^n$ . Let  $c \in \mathbf{R}^n$ . The variance of  $Y = c^T Z$  is  $\mathbf{var}(Y) = c^T \Sigma c$ . We define the *worst-case variance* of  $Y$ , denoted  $\mathbf{wcv}\mathbf{ar}(Y)$ , as the maximum possible value of  $c^T \tilde{\Sigma} c$ , over all  $\tilde{\Sigma} \in \mathbf{S}_+^n$  that satisfy  $\Sigma_{ii} = \tilde{\Sigma}_{ii}$ ,  $i = 1, \dots, n$ . In other words, the worst-case variance of  $Y$  is the maximum possible variance, if we are allowed to arbitrarily change the correlations between  $Z_i$  and  $Z_j$ . Of course we have  $\mathbf{wcv}\mathbf{ar}(Y) \geq \mathbf{var}(Y)$ .

- (a) Find a simple expression for  $\mathbf{wcv}\mathbf{ar}(Y)$  in terms of  $c$  and the diagonal entries of  $\Sigma$ . You must justify your expression.
- (b) *Portfolio optimization.* Explain how to find the portfolio  $x \in \mathbf{R}^n$  that maximizes the expected return  $\mu^T x$  subject to a limit on risk,  $\mathbf{var}(r^T x) = x^T \Sigma x \leq R$ , and a limit on worst-case risk  $\mathbf{wcv}\mathbf{ar}(r^T x) \leq R^{\text{wc}}$ , where  $R > 0$  and  $R^{\text{wc}} > R$  are given. Here  $\mu = \mathbf{E} r$  and  $\Sigma = \mathbf{E}(r - \mu)(r - \mu)^T$  are the (given) mean and covariance of the (random) return vector  $r \in \mathbf{R}^n$ .
- (c) Carry out the method of part (b) for the problem instance with data given in `wc_risk_portfolio_opt_data.m`. Also find the optimal portfolio when the worst-case risk limit is ignored. Find the expected return and worst-case risk for these two portfolios.

*Remark.* If a portfolio is highly leveraged, and the correlations in the returns change drastically, you (the portfolio manager) can be in big trouble, since you are now exposed to much more risk than you thought you were. And yes, this (almost exactly) has happened.

7. *Smallest confidence ellipsoid.* Suppose the random variable  $X$  on  $\mathbf{R}^n$  has log-concave density  $p$ . Formulate the following problem as a convex optimization problem: Find an ellipsoid  $\mathcal{E}$  that satisfies  $\mathbf{Prob}(X \in \mathcal{E}) \geq 0.95$  and is smallest, in the sense of minimizing the sum of the squares of its semi-axis lengths. You do not need to worry about how to solve the resulting convex optimization problem; it is enough to formulate the smallest confidence ellipsoid problem as the problem of minimizing a convex function over a convex set involving the parameters that define  $\mathcal{E}$ .



8. *Lyapunov analysis of a dynamical system.* We consider a discrete-time time-varying linear dynamical system with state  $x_t \in \mathbf{R}^n$ . The state propagates according to the linear recursion  $x_{t+1} = A_t x_t$ , for  $t = 0, 1, \dots$ , where the matrices  $A_t$  are unknown but satisfy  $A_t \in \mathcal{A} = \{A^{(1)}, \dots, A^{(K)}\}$ , where  $A^{(1)}, \dots, A^{(K)}$  are known. (In computer science, this would be called a non-deterministic linear automaton.) We call the sequence  $x_0, x_1, \dots$  a *trajectory* of the system. There are infinitely many trajectories, one for each sequence  $A_0, A_1, \dots$

The *Lyapunov exponent*  $\kappa$  of the system is defined as

$$\kappa = \sup_{A_0, A_1, \dots} \limsup_{t \rightarrow \infty} \|x_t\|_2^{1/t}.$$

(If you don't know what sup and lim sup mean, you can replace them with max and lim, respectively.) Roughly speaking, this means that all trajectories grow no faster than  $\kappa^t$ . When  $\kappa < 1$ , the system is called *exponentially stable*.

It is a hard problem to determine the Lyapunov exponent of the system, or whether the system is exponentially stable, given the data  $A^{(1)}, \dots, A^{(K)}$ . In this problem we explore a powerful method for computing an upper bound on the Lyapunov exponent.

- (a) Let  $P \in \mathbf{S}_{++}^n$  and define  $V(x) = x^T P x$ . Suppose  $V$  satisfies

$$V(A^{(i)}x) \leq \gamma^2 V(x) \text{ for all } x \in \mathbf{R}^n, i = 1, \dots, K.$$

Show that  $\kappa \leq \gamma$ . Thus  $\gamma$  is an upper bound on the Lyapunov exponent  $\kappa$ . (The function  $V$  is called a quadratic *Lyapunov function* for the system.)

- (b) Explain how to use convex or quasiconvex optimization to find a matrix  $P \in \mathbf{S}_{++}^n$  with the smallest value of  $\gamma$ , *i.e.*, with the best upper bound on  $\kappa$ . You must justify your formulation.
- (c) Carry out the method of part (b) for the specific problem with data given in `lyap_exp_bound_data.m`. Report the best upper bound on  $\kappa$ , to a tolerance of 0.01. The data  $A^{(i)}$  are given as a cell array; `A{i}` gives  $A^{(i)}$ .
- (d) *Approximate worst-case trajectory simulation.* The quadratic Lyapunov function found in part (c) can be used to generate sequences of  $A_t$  that tend to result in large values of  $\|x_t\|_2^{1/t}$ . Start from a random vector  $x_0$ . At each  $t$ , generate  $x_{t+1}$  by choosing  $A_t = A^{(i)}$  that maximizes  $V(A^{(i)}x_t)$ , where  $P$  is computed from part (c). Do this for 50 time steps, and generate 5 such trajectories. Plot  $\|x_t\|_2^{1/t}$  and  $\gamma$  against  $t$  to verify that the bound you obtained in the previous part is valid. Report the lower bound on the Lyapunov exponent that the trajectories suggest.