EE263 Prof. S. Boyd
Oct. 27 – 28 or Oct. 28 – 29, 2006.

Midterm exam

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

Please read the following instructions carefully.

- You may use any books, notes, or computer programs (*e.g.*, Matlab), but you may not discuss the exam with anyone until Oct. 30, after everyone has taken the exam. The only exception is that you can ask the TAs or Stephen Boyd for clarification, by emailing to the staff email address. We've tried pretty hard to make the exam unambiguous and clear, so we're unlikely to say much.
- Please address email inquiries to `ee263-aut0607-staff@lists`. This forwards the mail to the professor and the TAs. In particular, please do not use Stephen Boyd's or the TAs' individual email addresses.
- Since you have 24 hours, we expect your solutions to be legible, neat, and clear. Do not hand in your rough notes, and please try to simplify your solutions as much as you can. We will deduct points from solutions that are technically correct, but much more complicated than they need to be.
- Please check your email a few times during the exam, just in case we need to send out a clarification or other announcement. It's unlikely we'll need to do this, but you never know.
- Attach the official exam cover page (available when you pick up or drop off the exam) to your exam, and assemble your solutions to the problems in order, *i.e.*, problem 1, problem 2, . . . , problem 6. Start each solution on a new page.
- Please make a copy of your exam before handing it in. We have never lost one, but it might occur.
- When a problem involves some computation (say, using Matlab), we do not want just the final answers. We want a clear discussion and justification of exactly what you did, the Matlab source code that produces the result, and the final numerical result. Be sure to show us your verification that your computed solution satisfies whatever properties it is supposed to, at least up to numerical precision. For example, if you compute a vector x that is supposed to satisfy $Ax = b$ (say), show us the Matlab code that checks this, and the result. (This might be done by the Matlab code `norm(A*x-b)`; be sure to show us the result, which should be very small.) *We will not check your numerical solutions for you, in cases where there is more than one solution.*

- In the portion of your solutions where you explain the mathematical approach, you *cannot* refer to Matlab operators, such as the backslash operator. (You can, of course, refer to inverses of matrices, or any other standard mathematical construct.)
- Some of the problems are described in a practical setting, such as digital circuit design, image processing, and optimal control. *You do not need to understand anything about the application area to solve these problems.* We've taken special care to make sure all the information and math needed to solve the problem is given in the problem description.
- We *do not* expect you to be able to solve all parts of all problems, so don't worry if you cannot finish them all.
- Four of the problems require you to download and run a Matlab file to generate the data needed. These files can be found at the URL

  `http://www.stanford.edu/class/ee263/matlab/FILENAME`

  where you should substitute the particular filename (given in the problem) for `FILENAME`. *There are no links on the course web page pointing to these files, so you'll have to type in the whole URL yourself.*
- Please respect the honor code. Although we encourage you to work on homework assignments in small groups, *you cannot discuss the midterm with anyone*, with the exception of Stephen Boyd and the TAs, until everyone has taken it.
- Finally, a few hints:
  - Problems may be easier (or harder) than they might at first appear.
  - None of the problems require long calculations or any serious programming.

1. *Point of closest convergence of a set of lines.* We have m lines in $\mathbf{R}^n$, described as

$$\mathcal{L}_i = \{p_i + tv_i \mid t \in \mathbf{R}\}, \quad i = 1, \ldots, m,$$

where $p_i \in \mathbf{R}^n$, and $v_i \in \mathbf{R}^n$, with $\|v_i\| = 1$, for $i = 1, \ldots, m$. We define the distance of a point $z \in \mathbf{R}^n$ to a line $\mathcal{L}$ as

$$\mathbf{dist}(z, \mathcal{L}) = \min\{\|z - u\| \mid u \in \mathcal{L}\}.$$

(In other words, $\mathbf{dist}(z, \mathcal{L})$ gives the closest distance between the point z and the line $\mathcal{L}$.)

We seek a point $z^\star \in \mathbf{R}^n$ that minimizes the sum of the squares of the distances to the lines,

$$\sum_{i=1}^{m} \mathbf{dist}(z, \mathcal{L}_i)^2.$$

The point $z^\star$ that minimizes this quantity is called the *point of closest convergence.*

(a) Explain how to find the point of closest convergence, given the lines (*i.e.*, given $p_1, \ldots, p_m$ and $v_1, \ldots, v_m$). If your method works provided some condition holds (such as some matrix being full rank), say so. If you can relate this condition to a simple one involving the lines, please do so.

(b) Find the point $z^\star$ of closest convergence for the lines with data given in the Matlab file `line_conv_data.m`. This file contains $n \times m$ matrices `P` and `V` whose columns are the vectors $p_1, \ldots, p_m$, and $v_1, \ldots, v_m$, respectively. The file also contains commands to plot the lines and the point of closest convergence (once you have found it). Please include this plot with your solution.

2. *Estimating direction and amplitude of a light beam.* A light beam with (nonnegative) amplitude a comes from a direction $d \in \mathbf{R}^3$, where $\|d\| = 1$. (This means the beam travels in the direction $-d$.) The beam falls on $m \geq 3$ photodetectors, each of which generates a scalar signal that depends on the beam amplitude and direction, and the direction in which the photodetector is pointed. Specifically, photodetector i generates an output signal p_i, with

$$p_i = a\alpha \cos\theta_i + v_i,$$

where θ_i is the angle between the beam direction d and the outward normal vector q_i of the surface of the ith photodetector, and α is the photodetector sensitivity. You can interpret $q_i \in \mathbf{R}^3$, which we assume has norm one, as the direction the ith photodetector is pointed. We assume that $|\theta_i| < 90^\circ$, *i.e.*, the beam illuminates the top of the photodetectors. The numbers v_i are small measurement errors.

You are given the photodetector direction vectors $q_1, \ldots, q_m \in \mathbf{R}^3$, the photodetector sensitivity α, and the noisy photodetector outputs, $p_1, \ldots, p_m \in \mathbf{R}$. Your job is to estimate the beam direction $d \in \mathbf{R}^3$ (which is a unit vector), and a, the beam amplitude.

To describe unit vectors $q_1, \ldots, q_m$ and d in $\mathbf{R}^3$ we will use azimuth and elevation, defined as follows:

$$q = \begin{bmatrix} \cos\phi\cos\theta \\ \cos\phi\sin\theta \\ \sin\phi \end{bmatrix}.$$

Here ϕ is the elevation (which will be between 0° and 90°, since all unit vectors in this problem have positive 3rd component, *i.e.*, point upward). The azimuth angle θ, which varies from 0° to 360°, gives the direction in the plane spanned by the first and second coordinates. If $q = e_3$ (*i.e.*, the direction is directly up), the azimuth is undefined.

(a) Explain how to do this, using a method or methods from this class. The simpler the method the better. If some matrix (or matrices) needs to be full rank for your method to work, say so.

(b) Carry out your method on the data given in `beam_estim_data.m`. This mfile defines `p`, the vector of photodetector outputs, a vector `det_az`, which gives the azimuth angles of the photodetector directions, and a vector `det_el`, which gives the elevation angles of the photodetector directions. Note that both of these are given in *degrees*, not radians. Give your final estimate of the beam amplitude a and beam direction d (in azimuth and elevation, in degrees).

3. *Minimum energy input with way-point constraints.* We consider a vehicle that moves in $\mathbf{R}^2$ due to an applied force input. We will use a discrete-time model, with time index $k = 1, 2, \ldots$; time index k corresponds to time $t = kh$, where $h > 0$ is the sample interval. The position at time index k is denoted by $p(k) \in \mathbf{R}^2$, and the velocity by $v(k) \in \mathbf{R}^2$, for $k = 1, \ldots, K+1$. These are related by the equations

$$p(k+1) = p(k) + hv(k), \quad v(k+1) = (1-\alpha)v(k) + (h/m)f(k), \quad k = 1, \ldots, K,$$

where $f(k) \in \mathbf{R}^2$ is the force applied to the vehicle at time index k, $m > 0$ is the vehicle mass, and $\alpha \in (0, 1)$ models drag on the vehicle: In the absence of any other force, the vehicle velocity decreases by the factor $1 - \alpha$ in each time index. (These formulas are approximations of more accurate formulas that we will see soon, but for the purposes of this problem, we consider them exact.) The vehicle starts at the origin, at rest, *i.e.*, we have $p(1) = 0$, $v(1) = 0$. (We take $k = 1$ as the initial time, to simplify indexing.)

The problem is to find forces $f(1), \ldots, f(K) \in \mathbf{R}^2$ that minimize the cost function

$$J = \sum_{k=1}^{K} \|f(k)\|^2,$$

subject to *way-point constraints*

$$p(k_i) = w_i, \quad i = 1, \ldots, M,$$

where k_i are integers between 1 and K. (These state that at the time $t_i = hk_i$, the vehicle must pass through the location $w_i \in \mathbf{R}^2$.) Note that there is no requirement on the vehicle velocity at the way-points.

(a) Explain how to solve this problem, given all the problem data (*i.e.*, h, α, m, K, the way-points $w_1, \ldots, w_M$, and the way-point indices $k_1, \ldots, k_M$).

(b) Carry out your method on the specific problem instance with data $h = 0.1$, $m = 1$, $\alpha = 0.1$, $K = 100$, and the $M = 4$ way-points

$$w_1 = \begin{bmatrix} 2 \\ 2 \end{bmatrix}, \quad w_2 = \begin{bmatrix} -2 \\ 3 \end{bmatrix}, \quad w_3 = \begin{bmatrix} 4 \\ -3 \end{bmatrix}, \quad w_4 = \begin{bmatrix} -4 \\ -2 \end{bmatrix},$$

with way-point indices $k_1 = 10$, $k_2 = 30$, $k_3 = 40$, and $k_4 = 80$.

Give the optimal value of J.

Plot $f_1(k)$ and $f_2(k)$ versus k, using

```
subplot(211); plot(f(1,:));
subplot(212); plot(f(2,:));
```

We assume here that `f` is a $2 \times K$ matrix, with columns $f(1), \ldots, f(K)$.

Plot the vehicle trajectory, using `plot(p(1,:),p(2,:))`. Here `p` is a $2 \times (K+1)$ matrix with columns $p(1), \ldots, p(K+1)$.

4. *Digital circuit gate sizing.* A digital circuit consists of a set of n (logic) gates, interconnected by wires. Each gate has one or more inputs (typically between one and four), and one output, which is connected via the wires to other gate inputs and possibly to some external circuitry. When the output of gate i is connected to an input of gate j, we say that gate i *drives* gate j, or that gate j is in the *fan-out* of gate i. We describe the topology of the circuit by the *fan-out list* for each gate, which tells us which other gates the output of a gate connects to. We denote the fan-out list of gate i as $\mathrm{FO}(i) \subseteq \{1, \dots, n\}$. We can have $\mathrm{FO}(i) = \emptyset$, which means that the output of gate i does not connect to the inputs of any of the gates $1, \dots, n$ (presumably the output of gate i connects to some external circuitry). It's common to order the gates in such a way that each gate only drives gates with higher indices, *i.e.*, we have $\mathrm{FO}(i) \subseteq \{i+1, \dots, n\}$. We'll assume that's the case here. (This means that the gate interconnections form a directed acyclic graph.)

   To illustrate the notation, a simple digital circuit with $n = 4$ gates, each with 2 inputs, is shown below. For this circuit we have

$$\mathrm{FO}(1) = \{3, 4\}, \quad \mathrm{FO}(2) = \{3\}, \quad \mathrm{FO}(3) = \emptyset, \quad \mathrm{FO}(4) = \emptyset.$$

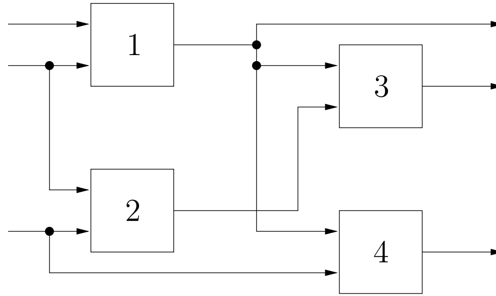

   The 3 input signals arriving from the left are called *primary inputs*, and the 3 output signals emerging from the right are called *primary outputs* of the circuit. (You don't need to know this, however, to solve this problem.)

   Each gate has a (real) *scale factor* or *size* x_i. These scale factors are the design variables in the gate sizing problem. They must satisfy $1 \leq x_i \leq x^{\max}$, where $x^{\max}$ is a given maximum allowed gate scale factor (typically on the order of 100). The total area of the circuit has the form

$$A = \sum_{i=1}^{n} a_i x_i,$$

   where a_i are positive constants.

   Each gate has an *input capacitance* C_i^{in}, which depends on the scale factor x_i as

$$C_i^{\mathrm{in}} = \alpha_i x_i,$$

   where α_i are positive constants.

Each gate has a *delay* d_i, which is given by

$$d_i = \beta_i + \gamma_i C_i^{\rm load}/x_i,$$

where β_i and γ_i are positive constants, and $C_i^{\rm load}$ is the *load capacitance* of gate i. Note that the gate delay d_i is always larger than β_i, which can be intepreted as the minimum possible delay of gate i, achieved only in the limit as the gate scale factor becomes large.

The load capacitance of gate i is given by

$$C_i^{\rm load} = C_i^{\rm ext} + \sum_{j \in {\rm FO}(i)} C_j^{\rm in},$$

where $C_i^{\rm ext}$ is a positive constant that accounts for the capacitance of the interconnect wires and external circuitry.

We will follow a simple design method, which assigns an equal delay T to all gates in the circuit, *i.e.*, we have $d_i = T$, where $T > 0$ is given. For a given value of T, there may or may not exist a feasible design (*i.e.*, a choice of the x_i, with $1 \leq x_i \leq x^{\rm max}$) that yields $d_i = T$ for $i = 1, \ldots, n$. We can assume, of course, that $T > \max_i \beta_i$, *i.e.*, T is larger than the largest minimum delay of the gates.

Finally, we get to the problem.

(a) Explain how to find a design $x^\star \in \mathbf{R}^n$ that minimizes T, subject to a given area constraint $A \leq A^{\rm max}$. You can assume the fanout lists, and all constants in the problem description are known; your job is to find the scale factors x_i. Be sure to explain how you determine if the design problem is feasible, *i.e.*, whether or not there is an x that gives $d_i = T$, with $1 \leq x_i \leq x^{\rm max}$, and $A \leq A^{\rm max}$.

Your method can involve any of the methods or concepts we have seen so far in the course. It can also involve a simple search procedure, *e.g.*, trying (many) different values of T over a range.

Note: this problem concerns the general case, and not the simple example shown above.

(b) Carry out your method on the particular circuit with data given in the file `gate_sizing_data.m`. The fan-out lists are given as an $n \times n$ matrix `F`, with i, j entry one if $j \in {\rm FO}(i)$, and zero otherwise. In other words, the ith row of `F` gives the fanout of gate i. The jth entry in the ith row is 1 if gate j is in the fan-out of gate i, and 0 otherwise.

Comments and hints.

- You do not need to know anything about digital circuits; *everything* you need to know is stated above.
- Yes, this problem *does* belong on the EE263 midterm.

5. *Oh no. It's the dreaded theory problem.* In the list below there are 11 statements about two square matrices A and B in $\mathbf{R}^{n\times n}$.

   (a) $\mathcal{R}(B) \subseteq \mathcal{R}(A)$.
   (b) there exists a matrix $Y \in \mathbf{R}^{n\times n}$ such that $B = YA$.
   (c) $AB = 0$.
   (d) $BA = 0$.
   (e) $\mathbf{rank}([\ A \quad B\ ]) = \mathbf{rank}(A)$.
   (f) $\mathcal{R}(A) \perp \mathcal{N}(B^T)$.
   (g) $\mathbf{rank}(\begin{bmatrix} A \\ B \end{bmatrix}) = \mathbf{rank}(A)$.
   (h) $\mathcal{R}(A) \subseteq \mathcal{N}(B)$.
   (i) there exists a matrix $Z \in \mathbf{R}^{n\times n}$ such that $B = AZ$.
   (j) $\mathbf{rank}([\ A \quad B\ ]) = \mathbf{rank}(B)$.
   (k) $\mathcal{N}(A) \subseteq \mathcal{N}(B)$.

   Your job is to collect them into (the largest possible) groups of equivalent statements. Two statements are equivalent if each one implies the other. For example, the statement 'A is onto' is equivalent to '$\mathcal{N}(A) = \{0\}$' (when A is square, which we assume here), because every square matrix that is onto has zero nullspace, and vice versa. Two statements are not equivalent if there exist (real) square matrices A and B for which one holds, but the other does not. A group of statements is equivalent if any pair of statements in the group is equivalent.

   We want *just* your answer, which will consist of lists of mutually equivalent statements. We will not read any justification. If you add any text to your answer, as in 'c and e are equivalent, provided A is nonsingular', we will mark your response as wrong.

   Put your answer in the following specific form. List each group of equivalent statements on a line, in (alphabetic) order. Each new line should start with the first letter not listed above. For example, you might give your answer as

   a, c, d, h
   b, i
   e
   f, g, j, k.

   This means you believe that statements a, c, d, and h are equivalent; statements b and i are equivalent; and statements f, g, j, and k are equivalent. You also believe that the first group of statements is not equivalent to the second, or the third, and so on.

   We will take points off for false groupings (*i.e.*, listing statements in the same line when they are not equivalent) as well as for missed groupings (*i.e.*, when you list equivalent statements in different lines).

6. *Smooth interpolation on a 2D grid.* This problem concerns arrays of real numbers on an $m \times n$ grid. Such as array can represent an image, or a sampled description of a function defined on a rectangle. We can describe such an array by a matrix $U \in \mathbf{R}^{m\times n}$, where U_{ij} gives the real number at location i, j, for $i = 1, \ldots, m$ and $j = 1, \ldots, n$. We will think of the index i as associated with the y axis, and the index j as associated with the x axis.

   It will also be convenient to describe such an array by a vector $u = \mathbf{vec}(U) \in \mathbf{R}^{mn}$. Here $\mathbf{vec}$ is the function that stacks the columns of a matrix on top of each other:

$$\mathbf{vec}(U) = \begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix},$$

   where $U = [u_1 \cdots u_n]$. To go back to the array representation, from the vector, we have $U = \mathbf{vec}^{-1}(u)$. (This looks complicated, but isn't; $\mathbf{vec}^{-1}$ just arranges the elements in a vector into an array.)

   We will need two linear functions that operate on $m \times n$ arrays. These are simple approximations of partial differentiation with respect to the x and y axes, respectively. The first function takes as argument an $m \times n$ array U and returns an $m \times (n-1)$ array V of forward (rightward) differences:

$$V_{ij} = U_{i,j+1} - U_{ij}, \quad i = 1, \ldots, m, \quad j = 1, \ldots, n-1.$$

   We can represent this linear mapping as multiplication by a matrix $D_x \in \mathbf{R}^{m(n-1)\times mn}$, which satisfies

$$\mathbf{vec}(V) = D_x\mathbf{vec}(U).$$

   (This looks scarier than it is—each row of the matrix D_x has exactly one $+1$ and one -1 entry in it.)

   The other linear function, which is a simple approximation of partial differentiation with respect to the y axis, maps an $m \times n$ array U into an $(m-1) \times n$ array W, is defined as

$$W_{ij} = U_{i+1,j} - U_{ij}, \quad i = 1, \ldots, m-1, \quad j = 1, \ldots, n.$$

   We define the matrix $D_y \in \mathbf{R}^{(m-1)n\times mn}$, which satisfies $\mathbf{vec}(W) = D_y\mathbf{vec}(U)$.

   We define the *roughness* of an array U as

$$R = \|D_x\mathbf{vec}(U)\|^2 + \|D_y\mathbf{vec}(U)\|^2.$$

   The roughness measure R is the sum of the squares of the differences of each element in the array and its neighbors. Small R corresponds to smooth, or smoothly varying, U. The roughness measure R is zero precisely for constant arrays, *i.e.*, when U_{ij} are all equal.

Now we get to the problem, which is to interpolate some unknown values in an array in the smoothest possible way, given the known values in the array. To define this precisely, we partition the set of indices $\{1, \dots, mn\}$ into two sets: I_{known} and I_{unknown}. We let $k \geq 1$ denote the number of known values (*i.e.*, the number of elements in I_{known}), and $mn - k$ the number of unknown values (the number of elements in I_{unknown}). We are given the values u_i for $i \in I_{\text{known}}$; the goal is to guess (or estimate or assign) values for u_i for $i \in I_{\text{unknown}}$. We'll choose the values for u_i, with $i \in I_{\text{unknown}}$, so that the resulting U is as smooth as possible, *i.e.*, so it minimizes R. Thus, the goal is to fill in or interpolate missing data in a 2D array (an image, say), so the reconstructed array is as smooth as possible.

We give the k known values in a vector $w_{\text{known}} \in \mathbf{R}^k$, and the $mn - k$ unknown values in a vector $w_{\text{unknown}} \in \mathbf{R}^{mn-k}$. The complete array is obtained by putting the entries of w_{known} and w_{unknown} into the correct positions of the array. We describe these operations using two matrices $Z_{\text{known}} \in \mathbf{R}^{mn \times k}$ and $Z_{\text{unknown}} \in \mathbf{R}^{mn \times (mn-k)}$, that satisfy

$$\mathbf{vec}(U) = Z_{\text{known}} w_{\text{known}} + Z_{\text{unknown}} w_{\text{unknown}}.$$

(This looks complicated, but isn't: Each row of these matrices is a unit vector, so multiplication with either matrix just stuffs the entries of the w vectors into particular locations in $\mathbf{vec}(U)$. In fact, the matrix $[Z_{\text{known}} \; Z_{\text{unknown}}]$ is an $mn \times mn$ permutation matrix.)

In summary, you are given the problem data w_{known} (which gives the known array values), Z_{known} (which gives the locations of the known values), and Z_{unknown} (which gives the locations of the unknown array values, in some specific order). Your job is to find w_{unknown} that minimizes R.

(a) Explain how to solve this problem. You are welcome to use any of the operations, matrices, and vectors defined above in your solution (*e.g.*, $\mathbf{vec}$, $\mathbf{vec}^{-1}$, D_x, D_y, Z_{known}, Z_{unknown}, w_{known}, $\dots$). If your solution is valid provided some matrix is (or some matrices are) full rank, say so.

(b) Carry out your method using the data created by `smooth_interpolation.m`. The file gives m, n, w_{known}, Z_{known} and Z_{unknown}. This file also creates the matrices D_x and D_y, which you are welcome to use. (This was *very* nice of us, by the way.) You are welcome to look at the code that generates these matrices, but you do not need to understand it. For this problem instance, around 50% of the array elements are known, and around 50% are unknown.

    The mfile also includes the original array `Uorig` from which we removed elements to create the problem. This is just so you can see how well your smooth reconstruction method does in reconstructing the original array. Of course, you cannot use `Uorig` to create your interpolated array `U`.

    To visualize the arrays use the Matlab command `imagesc()`, with matrix argument. If you prefer a grayscale image, or don't have a color printer, you can

issue the command `colormap gray`. The mfile that gives the problem data will plot the original image `Uorig`, as well as an image containing the known values, with zeros substituted for the unknown locations. This will allow you to see the pattern of known and unknown array values.

Compare `Uorig` (the original array) and `U` (the interpolated array found by your method), using `imagesc()`. Hand in complete source code, as well as the plots. Be sure to give the value of roughness R of U.

Hints:

- In Matlab, $\mathbf{vec}(U)$ can be computed as `U(:)`;
- $\mathbf{vec}^{-1}(u)$ can be computed as `reshape(u,m,n)`.