

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; there are no links on the course web page pointing to these files. 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.

Be sure you are using the most recent version of CVX, which is Version 2.0 (beta), build 937. You can check this using the command `cvx_version`.

Be sure to check your email often during the exam, just in case we need to send out an important announcement.

1. *Minimum time speed profile along a road.* A vehicle of mass $m > 0$ moves along a road in \mathbf{R}^3 , which is piecewise linear with given knot points $p_1, \dots, p_{N+1} \in \mathbf{R}^3$, starting at p_1 and ending at p_{N+1} . We let $h_i = (p_i)_3$, the z -coordinate of the knot point; these are the heights of the knot points (above sea-level, say). For your convenience, these knot points are equidistant, *i.e.*, $\|p_{i+1} - p_i\|_2 = d$ for all i . (The points give an arc-length parametrization of the road.) We let $s_i > 0$ denote the (constant) vehicle speed as it moves along road segment i , from p_i to p_{i+1} , for $i = 1, \dots, N$, and $s_{N+1} \geq 0$ denote the vehicle speed after it passes through knot point p_{N+1} . Our goal is to minimize the total time to traverse the road, which we denote T .

We let $f_i \geq 0$ denote the total fuel burnt while traversing the i th segment. This fuel burn is turned into an increase in vehicle energy given by ηf_i , where $\eta > 0$ is a constant that includes the engine efficiency and the energy content of the fuel. While traversing the i th road segment the vehicle is subject to a drag force, given by $C_D s_i^2$, where $C_D > 0$ is the coefficient of drag, which results in an energy loss $dC_D s_i^2$.

We derive equations that relate these quantities via energy balance:

$$\frac{1}{2}ms_{i+1}^2 + mgh_{i+1} = \frac{1}{2}ms_i^2 + mgh_i + \eta f_i - dC_D s_i^2, \quad i = 1, \dots, N,$$

where $g = 9.8$ is the gravitational acceleration. The lefthand side is the total vehicle energy (kinetic plus potential) after it passes through knot point p_{i+1} ; the righthand side is the total vehicle energy after it passes through knot point p_i , plus the energy gain from the fuel burn, minus the energy lost to drag. To set up the first vehicle speed s_1 requires an additional initial fuel burn f_0 , with $\eta f_0 = \frac{1}{2}ms_1^2$.

Fuel is also used to power the on-board system of the vehicle. The total fuel used for this purpose is f_{ob} , where $\eta f_{\text{ob}} = TP$, where $P > 0$ is the (constant) power consumption of the on-board system. We have a fuel capacity constraint: $\sum_{i=0}^N f_i + f_{\text{ob}} \leq F$, where $F > 0$ is the total initial fuel.

The problem data are $m, d, h_1, \dots, h_{N+1}, \eta, C_D, P$, and F . (You don't need the knot points p_i .)

- (a) Explain how to find the fuel burn levels f_0, \dots, f_N that minimize the time T , subject to the constraints.
- (b) Carry out the method described in part (a) for the problem instance with data given in `min_time_speed_data.m`. Give the optimal time T^* , and compare it to the time T^{unif} achieved if the fuel for propulsion were burned uniformly, *i.e.*, $f_0 = \dots = f_N$. For each of these cases, plot speed versus distance along the road, using the plotting code in the data file as a template.

2. *Polynomial approximation of inverse using eigenvalue information.* We seek a polynomial of degree k , $p(a) = c_0 + c_1a + c_2a^2 + \dots + c_k a^k$, for which

$$p(A) = c_0I + c_1A + c_2A^2 \dots + c_kA^k$$

is an approximate inverse of the nonsingular matrix A , for all $A \in \mathcal{A} \subset \mathbf{R}^{n \times n}$. When $\hat{x} = p(A)b$ is used as an approximate solution of the linear equation $Ax = b$, the associated residual norm is $\|A(p(A)b) - b\|_2$. We will judge our polynomial (*i.e.*, the coefficients c_0, \dots, c_k) by the worst case residual over $A \in \mathcal{A}$ and b in the unit ball:

$$R^{\text{wc}} = \sup_{A \in \mathcal{A}, \|b\|_2 \leq 1} \|A(p(A)b) - b\|_2.$$

The set of matrices we take is $\mathcal{A} = \{A \in \mathbf{S}^n \mid \sigma(A) \subseteq \Omega\}$, where $\sigma(A)$ is the set of eigenvalues of A (*i.e.*, its spectrum), and $\Omega \subset \mathbf{R}$ is a union of a set of intervals (that do not contain 0).

- (a) Explain how to find coefficients c_0^*, \dots, c_k^* that minimize R^{wc} . Your solution can involve expressions that involve the supremum of a polynomial (with scalar argument) over an interval.
- (b) Carry out your method for $k = 4$ and $\Omega = [-0.6, -0.3] \cup [0.7, 1.8]$. You can replace the supremum of a polynomial over Ω by a maximum over uniformly spaced (within each interval) points in Ω , with spacing 0.01. Give the optimal value $R^{\text{wc}*}$ and the optimal coefficients $c^* = (c_0^*, \dots, c_k^*)$.

Remarks. (Not needed to solve the problem.)

- The approximate inverse $p(A)b$ would be computed by recursively, requiring the multiplication of A with a vector k times.
- This approximate inverse could be used as a preconditioner for an iterative method.
- The Cayley-Hamilton theorem tells us that the inverse of any (invertible) matrix is a polynomial of degree $n - 1$ of the matrix. Our hope here, however, is to get a single polynomial, of relatively low degree, that serves as an approximate inverse for many different matrices.

3. *Fitting a generalized additive regression model.* A generalized additive model has the form

$$f(x) = \alpha + \sum_{j=1}^n f_j(x_j),$$

for $x \in \mathbf{R}^n$, where $\alpha \in \mathbf{R}$ is the offset, and $f_j : \mathbf{R} \rightarrow \mathbf{R}$, with $f_j(0) = 0$. The functions f_j are called the *regressor functions*. When each f_j is linear, *i.e.*, has the form $w_j x_j$, the generalized additive model is the same as the standard (linear) regression model. Roughly speaking, a generalized additive model takes into account nonlinearities in each regressor x_j , but not nonlinear interactions among the regressors. To visualize a generalized additive model, it is common to plot each regressor function (when n is not too large).

We will restrict the functions f_j to be piecewise-affine, with given knot points $p_1 < \dots < p_K$. This means that f_j is affine on the intervals $(-\infty, p_1]$, $[p_1, p_2]$, \dots , $[p_{K-1}, p_K]$, $[p_K, \infty)$, and continuous at p_1, \dots, p_K . Let C denote the total (absolute value of) change in slope across all regressor functions and all knot points. The value C is a measure of nonlinearity of the regressor functions; when $C = 0$, the generalized additive model reduces to a linear regression model.

Now suppose we observe samples or data $(x^{(1)}, y^{(1)}), \dots, (x^{(N)}, y^{(N)}) \in \mathbf{R}^n \times \mathbf{R}$, and wish to fit a generalized additive model to the data. We choose the offset and the regressor functions to minimize

$$\frac{1}{N} \sum_{i=1}^N (y^{(i)} - f(x^{(i)}))^2 + \lambda C,$$

where $\lambda > 0$ is a regularization parameter. (The first term is the mean-square error.)

- (a) Explain how to solve this problem using convex optimization.
- (b) Carry out the method of part (a) using the data in the file `gen_add_reg_data.m`. This file contains the data, given as an $N \times n$ matrix \mathbf{X} (whose rows are $(x^{(i)})^T$), a column vector \mathbf{y} (which give $y^{(i)}$), a vector \mathbf{p} that gives the knot points, and the scalar `lambda`.

Give the mean-square error achieved by your generalized additive regression model. Compare the estimated and true regressor functions in a 3×3 array of plots (using the plotting code in the data file as a template), over the range $-10 \leq x_i \leq 10$. The true regressor functions (to be used only for plotting, of course) are given in the cell array `f`.

Hints.

- You can represent each regressor function f_j as a linear combination of the basis functions $b_0(u) = u$ and $b_i(u) = (u - p_k)_+ - (-p_k)_+$ for $k = 1, 2, \dots, K$, where $(a)_+ = \max\{a, 0\}$.
- You might find the matrix $\mathbf{XX} = [b_0(\mathbf{X}) \quad b_1(\mathbf{X}) \quad \dots \quad b_K(\mathbf{X})]$ useful.

4. *Maximum likelihood estimation for an affinely transformed distribution.* Let z be a random variable on \mathbf{R}^n with density $p_z(u) = \exp -\phi(\|u\|_2)$, where $\phi : \mathbf{R} \rightarrow \mathbf{R}$ is convex and increasing. Examples of such distributions include the standard normal $\mathcal{N}(0, \sigma^2 I)$, with $\phi(u) = (u)_+^2 + \alpha$, and the multivariable Laplacian distribution, with $\phi(u) = (u)_+ + \beta$, where α and β are normalizing constants, and $(a)_+ = \max\{a, 0\}$. Now let x be the random variable $x = Az + b$, where $A \in \mathbf{R}^{n \times n}$ is nonsingular. The distribution of x is parametrized by A and b .

Suppose x_1, \dots, x_N are independent samples from the distribution of x . Explain how to find a maximum likelihood estimate of A and b using convex optimization. If you make any further assumptions about A and b (beyond invertibility of A), you must justify it.

Hint. The density of $x = Az + b$ is given by

$$p_x(v) = \frac{1}{|\det A|} p_z(A^{-1}(v - b)).$$

5. *Functions of a random variable with log-concave density.* Suppose the random variable X on \mathbf{R}^n has log-concave density, and let $Y = g(X)$, where $g : \mathbf{R}^n \rightarrow \mathbf{R}$. For each of the following statements, either give a counterexample, or show that the statement is true.
- (a) If g is affine and not constant, then Y has log-concave density.
 - (b) If g is convex, then $\mathbf{Prob}(Y \leq a)$ is a log-concave function of a .
 - (c) If g is concave, then $\mathbf{E}((Y - a)_+)$ is a convex and log-concave function of a . (This quantity is called the tail expectation of Y ; you can assume it exists. We define $(s)_+$ as $(s)_+ = \max\{s, 0\}$.)

6. *Affine policy.* We consider a family of LPs, parametrized by the random variable u , which is uniformly distributed on $\mathcal{U} = [-1, 1]^p$,

$$\begin{aligned} & \text{minimize} && c^T x \\ & \text{subject to} && Ax \preceq b(u), \end{aligned}$$

where $x \in \mathbf{R}^n$, $A \in \mathbf{R}^{m \times n}$, and $b(u) = b_0 + Bu \in \mathbf{R}^m$ is an affine function of u . You can think of u_i as representing a deviation of the i th parameter from its nominal value. The parameters might represent (deviations in) levels of resources available, or other varying limits.

The problem is to be solved many times; in each time, the value of u (*i.e.*, a sample) is given, and then the decision variable x is chosen. The mapping from u into the decision variable $x(u)$ is called the *policy*, since it gives the decision variable value for each value of u . When enough time and computing hardware is available, we can simply solve the LP for each new value of u ; this is an optimal policy, which we denote $x^*(u)$.

In some applications, however, the decision $x(u)$ must be made very quickly, so solving the LP is not an option. Instead we seek a suboptimal policy, which is affine: $x^{\text{aff}}(u) = x_0 + Ku$, where x_0 is called the *nominal decision* and $K \in \mathbf{R}^{n \times p}$ is called the *feedback gain matrix*. (Roughly speaking, x_0 is our guess of x before the value of u has been revealed; Ku is our modification of this guess, once we know u .) We determine the policy (*i.e.*, suitable values for x_0 and K) ahead of time; we can then evaluate the policy (that is, find $x^{\text{aff}}(u)$ given u) very quickly, by matrix multiplication and addition.

We will choose x_0 and K in order to minimize the expected value of the objective, while insisting that for any value of u , feasibility is maintained:

$$\begin{aligned} & \text{minimize} && \mathbf{E} c^T x^{\text{aff}}(u) \\ & \text{subject to} && Ax^{\text{aff}}(u) \preceq b(u) \quad \forall u \in \mathcal{U}. \end{aligned}$$

The variables here are x_0 and K . The expectation in the objective is over u , and the constraint requires that $Ax^{\text{aff}}(u) \preceq b(u)$ hold almost surely.

- (a) Explain how to find optimal values of x_0 and K by solving a standard explicit convex optimization problem (*i.e.*, one that does not involve an expectation or an infinite number of constraints, as the one above does.) The numbers of variables or constraints in your formulation should not grow exponentially with the problem dimensions n , p , or m .
- (b) Carry out your method on the data given in `affine_pol_data.m`. To evaluate your affine policy, generate 100 independent samples of u , and for each value, compute the objective value of the affine policy, $c^T x^{\text{aff}}(u)$, and of the optimal policy, $c^T x^*(u)$. Scatter plot the objective value of the affine policy (y -axis) versus the objective value of the optimal policy (x -axis), and include the line $y = x$ on the plot. Report the average values of $c^T x^{\text{aff}}(u)$ and $c^T x^*(u)$ over your samples. (These are estimates of $\mathbf{E} c^T x^{\text{aff}}(u)$ and $\mathbf{E} c^T x^*(u)$. The first number, by the way, can be found exactly.)

7. *Cost-comfort trade-off in air conditioning.* A heat pump (air conditioner) is used to cool a residence to temperature T_t in hour t , on a day with outside temperature T_t^{out} , for $t = 1, \dots, 24$. These temperatures are given in degrees Kelvin, and we will assume that $T_t^{\text{out}} \geq T_t$.

A total amount of heat $Q_t = \alpha(T_t^{\text{out}} - T_t)$ must be removed from the residence in hour t , where α is a positive constant (related to the quality of thermal insulation).

The electrical energy required to pump out this heat is given by $E_t = Q_t/\gamma_t$, where

$$\gamma_t = \eta \frac{T_t}{T_t^{\text{out}} - T_t}$$

is the *coefficient of performance* of the heat pump and $\eta \in (0, 1]$ is the efficiency constant. The efficiency is typically around 0.6 for a modern unit; the theoretical limit is $\eta = 1$. (When $T_t = T_t^{\text{out}}$, we take $\gamma_t = \infty$ and $E_t = 0$.)

Electrical energy prices vary with the hour, and are given by $P_t > 0$ for $t = 1, \dots, 24$. The total energy cost is $C = \sum_t P_t E_t$. We will assume that the prices are known.

Discomfort is measured using a piecewise-linear function of temperature,

$$D_t = (T_t - T^{\text{ideal}})_+,$$

where T^{ideal} is an ideal temperature, below which there is no discomfort. The total daily discomfort is $D = \sum_{t=1}^{24} D_t$. You can assume that $T^{\text{ideal}} < T_t^{\text{out}}$.

To get a point on the optimal cost-comfort trade-off curve, we will minimize $C + \lambda D$, where $\lambda > 0$. The variables to be chosen are T_1, \dots, T_{24} ; all other quantities described above are given.

Show that this problem has an analytical solution of the form $T_t = \psi(P_t, T_t^{\text{out}})$, where $\psi : \mathbf{R}^2 \rightarrow \mathbf{R}$. The function ψ can depend on the constants $\alpha, \eta, T^{\text{ideal}}, \lambda$. Give ψ explicitly. You are free (indeed, encouraged) to check your formula using CVX, with made up values for the constants.

Disclaimer. The focus of this course is *not* on deriving 19th century pencil and paper solutions to problems. But every now and then, a practical problem will actually have an analytical solution. This is one of them.

8. *Least-cost road grading.* A road is to be built along a given path. We must choose the height of the roadbed (say, above sea level) along the path, minimizing the total cost of grading, subject to some constraints. The cost of grading (*i.e.*, moving earth to change the height of the roadbed from the existing elevation) depends on the difference in height between the roadbed and the existing elevation. When the roadbed is below the existing elevation it is called a *cut*; when it is above it is called a *fill*. Each of these incurs engineering costs; for example, fill is created in a series of *lifts*, each of which involves dumping just a few inches of soil and then compacting it. Deeper cuts and higher fills require more work to be done on the road shoulders, and possibly, the addition of reinforced concrete structures to stabilize the earthwork. This explains why the marginal cost of cuts and fills increases with their depth/height.

We will work with a discrete model, specifying the road height as h_i , $i = 1, \dots, n$, at points equally spaced a distance d from each other along the given path. These are the variables to be chosen. (The heights h_1, \dots, h_n are called a *grading plan*.) We are given e_i , $i = 1, \dots, n$, the existing elevation, at the points. The grading cost is

$$C = \sum_{i=1}^n (\phi^{\text{fill}}((h_i - e_i)_+) + \phi^{\text{cut}}((e_i - h_i)_+)),$$

where ϕ^{fill} and ϕ^{cut} are the fill and cut cost functions, respectively, and $(a)_+ = \max\{a, 0\}$. The fill and cut functions are increasing and convex. The goal is to minimize the grading cost C .

The road height is constrained by given limits on the first, second, and third derivatives:

$$\begin{aligned} |h_{i+1} - h_i|/d &\leq D^{(1)}, & i = 1, \dots, n-1 \\ |h_{i+1} - 2h_i + h_{i-1}|/d^2 &\leq D^{(2)}, & i = 2, \dots, n-1 \\ |h_{i+1} - 3h_i + 3h_{i-1} - h_{i-2}|/d^3 &\leq D^{(3)}, & i = 3, \dots, n-1, \end{aligned}$$

where $D^{(1)}$ is the maximum allowable road slope, $D^{(2)}$ is the maximum allowable curvature, and $D^{(3)}$ is the maximum allowable third derivative.

- (a) Explain how to find the optimal grading plan.
- (b) Find the optimal grading plan for the problem with data given in `road_grading_data.m`, and fill and cut cost functions

$$\phi^{\text{fill}}(u) = 2(u)_+^2 + 30(u)_+, \quad \phi^{\text{cut}} = 12(u)_+^2 + (u)_+.$$

Plot $h_i - e_i$ for the optimal grading plan and report the associated cost.

- (c) Suppose the optimal grading problem with $n = 1000$ can be solved on a particular machine (say, with one, or just a few, cores) in around one second. Assuming the author of the software took EE364a, about how long will it take to solve the optimal grading problem with $n = 10000$? Give a very brief justification of your answer, no more than a few sentences.