EE364a Homework 5 additional problems

1. *Optimal operation of a hybrid vehicle.* Solve the instance of the hybrid vehicle operation problem described in exercise 4.65 in *Convex Optimization*, with problem data given in the file `hybrid_veh_data.m`, and fuel use function $F(p) = p + \gamma p^2$ (for $p \geq 0$).

   Hint. You will actually formulate and solve a *relaxation* of the original problem. You may find that some of the equality constraints you relaxed to inequality constraints do not hold for the solution found. This is not an error: it just means that there is no incentive (in terms of the objective) for the inequality to be tight. You can fix this in (at least) two ways. One is to go back and adjust certain variables, without affecting the objective and maintaining feasibility, so that the relaxed constraints hold with equality. Another simple method is to add to the objective a term of the form

   $$\epsilon \sum_{t=1}^{T} \max\{0, -P_{\rm mg}(t)\},$$

   where ϵ is small and positive. This makes it more attractive to use the brakes to extract power from the wheels, even when the battery is (or will be) full (which removes any fuel incentive).

   Find the optimal fuel consumption, and compare to the fuel consumption with a non-hybrid version of the same vehicle (*i.e.*, one without a battery). Plot the braking power, engine power, motor/generator power, and battery energy versus time.

   How would you use optimal dual variables for this problem to find $\partial F_{\rm total}/\partial E_{\rm batt}^{\rm max}$, *i.e.*, the partial derivative of optimal fuel consumption with respect to battery capacity? (You can just assume that this partial derivative exists.) You do not have to give a long derivation or proof; you can just state how you would find this derivative from optimal dual variables for the problem. Verify your method numerically, by changing the battery capacity a small amount and re-running the optimization, and comparing this to the prediction made using dual variables.

2. *Heuristic suboptimal solution for Boolean LP.* This exercise builds on exercises 4.15 and 5.13, which involve the Boolean LP

   $$\begin{array}{ll} \mbox{minimize} & c^T x \\ \mbox{subject to} & Ax \preceq b \\ & x_i \in \{0,1\}, \quad i = 1,\ldots,n, \end{array}$$

   with optimal value $p^\star$. Let $x^{\rm rlx}$ be a solution of the LP relaxation

   $$\begin{array}{ll} \mbox{minimize} & c^T x \\ \mbox{subject to} & Ax \preceq b \\ & 0 \preceq x \preceq \mathbf{1}, \end{array}$$

so $L = c^T x^{\rm rlx}$ is a lower bound on $p^\star$. The relaxed solution $x^{\rm rlx}$ can also be used to guess a Boolean point $\hat{x}$, by rounding its entries, based on a threshold $t \in [0, 1]$:

$$\hat{x}_i = \left\{ \begin{array}{ll} 1 & x_i^{\rm rlx} \geq t \\ 0 & \mbox{otherwise,} \end{array} \right.$$

for $i = 1, \ldots, n$. Evidently $\hat{x}$ is Boolean (*i.e.*, has entries in $\{0, 1\}$). If it is feasible for the Boolean LP, *i.e.*, if $A\hat{x} \preceq b$, then it can be considered a guess at a good, if not optimal, point for the Boolean LP. Its objective value, $U = c^T \hat{x}$, is an upper bound on $p^\star$. If U and L are close, then $\hat{x}$ is nearly optimal; specifically, $\hat{x}$ cannot be more than $(U - L)$-suboptimal for the Boolean LP.

This rounding need not work; indeed, it can happen that for all threshold values, $\hat{x}$ is infeasible. But for some problem instances, it can work well.

Of course, there are many variations on this simple scheme for (possibly) constructing a feasible, good point from $x^{\rm rlx}$.

Finally, we get to the problem. Generate problem data using

```
rand('state',0);
n=100;
m=300;
A=rand(m,n);
b=A*ones(n,1)/2;
c=-rand(n,1);
```

You can think of x_i as a job we either accept or decline, and $-c_i$ as the (positive) revenue we generate if we accept job i. We can think of $Ax \preceq b$ as a set of limits on m resources. A_{ij}, which is positive, is the amount of resource i consumed if we accept job j; b_i, which is positive, is the amount of resource i available.

Find a solution of the relaxed LP and examine its entries. Note the associated lower bound L. Carry out threshold rounding for (say) 100 values of t, uniformly spaced over $[0, 1]$. For each value of t, note the objective value $c^T \hat{x}$ and the maximum constraint violation $\max_i (A\hat{x} - b)_i$. Plot the objective value and the maximum violation versus t. Be sure to indicate on the plot the values of t for which $\hat{x}$ is feasible, and those for which it is not.

Find a value of t for which $\hat{x}$ is feasible, and gives minimum objective value, and note the associated upper bound U. Give the gap $U - L$ between the upper bound on $p^\star$ and the lower bound on $p^\star$. If you define vectors `obj` and `maxviol`, you can find the upper bound as `U=min(obj(find(maxviol<=0)))`.

3. *Maximizing house profit in a gamble and imputed probabilities.* A set of n participants bet on which one of m outcomes, labeled $1, \ldots, m$, will occur. Participant i offers to purchase up to $q_i > 0$ gambling contracts, at price $p_i > 0$, that the true outcome will

be in the set $S_i \subset \{1, \ldots, m\}$. The house then sells her x_i contracts, with $0 \le x_i \le q_i$. If the true outcome j is in S_i, then participant i receives \$1 per contract, *i.e.*, x_i. Otherwise, she loses, and receives nothing. The house collects a total of $x_1p_1 + \cdots + x_np_n$, and pays out an amount that depends on the outcome j,

$$\sum_{j \in S_i} x_i.$$

The difference is the house profit.

(a) *Optimal house strategy.* How should the house decide on x so that its worst-case profit (over the possible outcomes) is maximized? (The house determines x after examining all the participant offers.)

(b) *Imputed probabilities.* Suppose $x^\star$ maximizes the worst-case house profit. Show that there exists a probability distribution π on the possible outcomes (*i.e.*, $\pi \in \mathbf{R}^m_+$, $\mathbf{1}^T\pi = 1$) for which $x^\star$ also maximizes the expected house profit. Explain how to find π.

Hint. Formulate the problem in part (a) as an LP; you can construct π from optimal dual variables for this LP.

Remark. Given π, the 'fair' price for offer i is $p_i^{\mathrm{fair}} = \sum_{j \in S_i} \pi_j$. All offers with $p_i > p_i^{\mathrm{fair}}$ will be completely filled (*i.e.*, $x_i = q_i$); all offers with $p_i < p_i^{\mathrm{fair}}$ will be rejected (*i.e.*, $x_i = 0$).

Remark. This exercise shows how the probabilities of outcomes (*e.g.*, elections) can be guessed from the offers of a set of gamblers.

(c) *Numerical example.* Carry out your method on the simple example below with $n = 5$ participants, $m = 5$ possible outcomes, and participant offers

Participant i	p_i	q_i	S_i
1	0.50	10	{1,2}
2	0.60	5	{4}
3	0.60	5	{1,4,5}
4	0.60	20	{2,5}
5	0.20	10	{3}

Compare the optimal worst-case house profit with the worst-case house profit, if all offers were accepted (*i.e.*, $x_i = q_i$). Find the imputed probabilities.

4. *Minimax rational fit to the exponential.* (See exercise 6.9.) We consider the specific problem instance with data

$$t_i = -3 + 6(i-1)/(k-1), \quad y_i = e^{t_i}, \quad i = 1, \ldots, k,$$

where $k = 201$. (In other words, the data are obtained by uniformly sampling the exponential function over the interval $[-3, 3]$.) Find a function of the form

$$f(t) = \frac{a_0 + a_1t + a_2t^2}{1 + b_1t + b_2t^2}$$

that minimizes $\max_{i=1,\ldots,k} |f(t_i) - y_i|$. (We require that $1 + b_1 t_i + b_2 t_i^2 > 0$ for $i = 1, \ldots, k$.)

Find optimal values of a_0, a_1, a_2, b_1, b_2, and give the optimal objective value, computed to an accuracy of 0.001. Plot the data and the optimal rational function fit on the same plot. On a different plot, give the fitting error, *i.e.*, $f(t_i) - y_i$.

Hint. You can use `strcmp(cvx_status,'Solved')`, after `cvx_end`, to check if a feasibility problem is feasible.

5. *Maximum likelihood estimation of x and noise mean and covariance.* Consider the maximum likelihood estimation problem with the linear measurement model

$$y_i = a_i^T x + v_i, \quad i = 1, \ldots, m$$

(as discussed on page 7-3 of the lecture notes and page 352 of the textbook). The vector $x \in \mathbf{R}^n$ is a vector of unknown parameters, y_i are the measurement values, and v_i are independent and identically distributed measurement errors.

In this problem we make the assumption that the *normalized* probability density function of the errors is given (normalized to have zero mean and unit variance), but not their mean and variance. In other words, the density of the measurement errors v_i is

$$p(z) = \frac{1}{\sigma} f(\frac{z - \mu}{\sigma}),$$

where f is a given, normalized density. The parameters μ and σ are the mean and standard deviation of the distribution p, and are not known.

The maximum likelihood estimates of x, μ, σ are the maximizers of the log-likelihood function

$$\sum_{i=1}^m \log p(y_i - a_i^T x) = -m \log \sigma + \sum_{i=1}^m \log f(\frac{y_i - a_i^T x - \mu}{\sigma}),$$

where y is the observed value. Show that if f is log-concave, then the maximum likelihood estimates of x, μ, σ can be determined by solving a convex optimization problem.

6. *Maximum likelihood estimation for exponential family.* A probability distribution on $\mathcal{D} \subseteq \mathbf{R}^n$, parametrized by $\theta \in \mathbf{R}^m$, is called an *exponential family* if it has the form

$$p_\theta(x) = a(\theta) \exp(\theta^T c(x))$$

for $x \in \mathcal{D}$, where $c : \mathbf{R}^n \to \mathbf{R}^m$, and

$$a(\theta) = \left(\int_{\mathcal{D}} \exp(\theta^T c(x))\, dx \right)^{-1}.$$

(We consider only values of θ for which the integral above is finite.) Many families of distributions have this form, for appropriate choice of the parameter θ.

(a) When $c(x) = x$ and $\mathcal{D} = \mathbf{R}^n_+$, what is the associated family of distributions? What is the set of valid values of θ?

(b) Explain how to represent the normal family $\mathcal{N}(\mu, \Sigma)$ as an exponential family. *Hint.* Use parameter $(z, Y) = (\Sigma^{-1}\mu, \Sigma^{-1})$. With this parameter, $\theta^T c(x)$ has the form $z^T c_1(x) + \mathbf{Tr}\, Y C_2(x)$, where $C_2(x) \in \mathbf{S}^n$.

(c) Show that for any $x \in \mathcal{D}$, the log-likelihood function $\log p_\theta(x)$ is concave in θ. This means that maximum-likelihood estimation for an exponential family leads to a convex optimization problem. You don't have to give a formal proof of concavity of $\log p_\theta(x)$; if you like, you can approximate the integral appearing in the expression as a (finite) Riemann sum, show concavity of this approxmation, and then just state 'take the limit'.