Problem Set 3

All numbered problems come from Stephen Boyd and Lieven Vandenberghe's Additional exercises for Convex Optimization. Use the pdf and data on the course website.

Problems

Choose at least four of the five problems below (not counting the survey). If you do all five correctly, you will get 0.5pts of extra credit on your homework score, which cannot exceed 12pts total.

1. *Quantile regression.* In this problem, we are going to consider norm minimization using the tilted absolute penalty function with parameter $\tau \in [0, 1]$

$$p_\tau^{\text{tlt}}(u) = \tau u_- + (1-\tau)u_+ = \begin{cases} -\tau u & u < 0 \\ (1-\tau)u & u \geq 0. \end{cases}$$

With $\tau = 1/2$, we recover the standard absolute value penalty function. With $\tau > 1/2$, we attach a higher penalty to under-estimate than over-estimate, so predictions are 'high', and vice versa for $\tau < 1/2$.

a) Plot $p_\tau^{\text{tlt}}(u)$ for $\tau = 0.05$, 0.5, and 0.95.

In this problem, we will assume that the data is generated by a linear model with additive, zero mean noise, *i.e.*, , $y = Ax + v$ and we solve the problem

$$\text{minimize} \quad \sum_{i=1}^m p_\tau^{\text{tlt}}(a_i^T x - y_i),$$

where $x \in \mathbf{R}^n$ is our variable. Define the residual as $r = Ax^\star - y$.

b) What fraction of residuals r_i do you expect to be positive? How about negative?

c) An autoregressive model is a model of the form

$$y_t = \sum_{i=1}^n x_i y_{t-i} + v_t,$$

where n is the 'memory' of the predictor. Fit three autoregressive models to the data `y_train` in `ar_data.jl` using the tilted absolute penalty with $\tau = 0.1$, 0.5, and 0.9 Plot a histogram of the residuals for each model (on the training data). On a single plot, show each model's one-step-ahead predictions on the test data, `y_test` (your plot should have four lines: one for each predictor and one for the data itself).

Notes: You can model the autoregressive model as $\hat{y} = Ax$, where $A \in \mathbf{R}^{N-n \times n}$, where N is the number of training data points of y. Look at the file `ar_data.jl` to see how the data is generated. Since there is an n-sample lag in the model, your output plot will be of length `n_test - n`.

2. *5.21 Robust LP with polyhedral cost uncertainty.* (p. 58)

3. *21.19 Typesetting LaTeX.* (p. 266)

4. *7.38 Bounding the median.* (p. 114)

   Hint: **med** is not convex... apologies I know I said that I didn't plan to assign quasiconvex optimization problems in lecture. In terms of the optimization variable, the probability distribution p, it may be helpful to work with the indices instead $\mathbf{med}^{\text{ind}}(p) = \min\{k \mid \sum_{i=1}^{k} p_i = \mathbf{1}_k^T p \geq 1/2\}$. (Note that there is a one to one mapping between indices k and values x_k.) Also, since this problem is small, you don't need to worry about performing bisection search for full credit; solving up to N problems is fine.

5. *17.11 Planning production with uncertain demand.* (p. 194)

6. Please fill out the (anonymous) course survey at `https://forms.gle/7o6DKZEFwHKFfx9D7`