

ECE 6270, Spring 2021

Homework #1

Due Thursday, Jan 28, at 11:59pm Suggested Reading: B&V, Chapter 1 and Sections 4.3, 4.4, 4.6.

1. Prepare a one paragraph summary of what we talked about in class so far. I do not want just a bulleted list of topics, I want you to use complete sentences and establish context (Why is what we have learned relevant? How does it connect with other classes?). The more insight you give, the better.
2. Familiarize yourself with Python.
 - (a) Install (if necessary) the Anaconda distribution of Python at <https://www.anaconda.com/products/individual>.
 - (b) Review the Python tutorial at <https://cs4540-f18.github.io/notes/python-basics>.
 - (c) Check out and consider bookmarking this site: <https://numpy.org/doc/stable/user/numpy-for-matlab-users.html>
 - (d) If you feel it would be helpful, especially if you are new to Python, consider taking the time to go through one (or both) of these more basic Python tutorials:
 - <https://developers.google.com/edu/python/>
 - <https://www.learnpython.org/>
3. In the problems below and on subsequent homeworks you can/will make frequent use of the CVXPY package. You will probably need to install this package as it does not come standard with most python distributions. This should be easy – simply open up a Python prompt and type `pip install cvxpy`. If you are on Windows you may also need to install the Visual Studio build tools if you don't already have them. If you are having any trouble, see <https://www.cvxpy.org/install/> for details. Install CVXPY and poke around on the CVXPY website to get a feeling for what it can do.
4. In our first class we mentioned a nonconvex optimization problem that nevertheless has a simple solution:

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{maximize}} \quad \mathbf{x}^T \mathbf{A} \mathbf{x} \quad \text{subject to} \quad \|\mathbf{x}\|_2 = 1.$$

Here we consider a special case of this problem (that is not hard to extend to the general case). Suppose that \mathbf{D} is a *diagonal* $N \times N$ matrix whose diagonal elements are positive. Show that a maximizer $\hat{\mathbf{x}}$ to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{maximize}} \quad \mathbf{x}^T \mathbf{D} \mathbf{x} \quad \text{subject to} \quad \|\mathbf{x}\|_2 = 1$$

has a 1 in the entry corresponding to the largest diagonal element of \mathbf{D} , and is 0 elsewhere. You should approach this by rigorously arguing that any feasible $\mathbf{x}' \neq \hat{\mathbf{x}}$ cannot achieve a higher value for the objective function.

5. A significant motivation in the early development of optimization algorithms involved applications in logistics and supply chain management. This is a somewhat contrived example. Suppose you are the manager of a factory that produces two products: toilet paper and paper towels. The profit for paper towels is \$10 per ton of pulp, but the profit for toilet paper is \$100 per ton of pulp. Your factory has a supply of 8 000 tons of pulp per month. In addition, you face the following conditions:
- You have a government contract that requires you to produce 200 tons of paper towels each month.¹
 - You must ship all of the products that you produce using a fleet of trucks that has a capacity of 100 000 ton-miles (i.e., you could transport 100 000 tons 1 mile, or 1 ton 100 000 miles, or somewhere in between).
 - Paper towels are delivered to a warehouse 10 miles from your factory, while toilet paper is delivered to a warehouse 35 miles from your factory.

You would like to decide how much of each product to produce to maximize your profit. Write down a mathematical formulation of this as an optimization problem. What kind of optimization problem is this (linear, quadratic, semidefinite, ...)?

6. Consider the optimization program

$$\min_{\mathbf{x} \in \mathbb{R}^N} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_1,$$

where \mathbf{A} is a given $M \times N$ matrix and $\mathbf{y} \in \mathbb{R}^M$. Show how this can be converted into an equivalent linear program.

7. One important optimization problem that arises in the context of machine learning has the form

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \sum_{m=1}^M \max(\mathbf{a}_m^T \mathbf{x}, z_m) + \tau \|\mathbf{x}\|_2^2.$$

Show how to express this program as a quadratic program with linear constraints. (Hint: $\max(\mathbf{a}^T \mathbf{x}, z) \leq u$ is the same as saying $\mathbf{a}^T \mathbf{x} \leq u$ and $z \leq u$.)

8. You hopefully have some prior experience with least squares problems in which we aim to find an \mathbf{x} that minimizes $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$. While minimizing the ℓ_2 norm of the error is *often* a good idea, a big part of why it is so popular is just that it is so easy to solve the underlying optimization problem. However, with modern tools we can easily explore situations where other ℓ_p norms might be more appropriate. In particular, we will consider minimizing $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_1$ and $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_\infty$.

- (a) Download the file `hw01_prob8.py`. This file sets up a simple regression problem in which \mathbf{y} consists of noisy observations of a smooth function $f(t)$ and considers three noise models: Gaussian noise, sparse Gaussian noise, and uniform noise. Compute and plot the least squares solution for each case.

¹Assume that 1 ton of paper towels or 1 ton of toilet paper each require 1 ton of raw pulp.

(b) Now estimate the solution for the Chebyshev approximation problem:

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_\infty.$$

You do not have to convert this to a linear program first; CVXPY has built-in tools for handling the ℓ_∞ norm. Compute and plot the resulting estimate of \mathbf{x} for each of the three noise cases.

- (c) Next repeat part (b) for $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_1$. Again, compute and plot the results for each of the three noise cases.
- (d) Compare the results that you obtained in the previous parts. Which method seems best suited for each case?

9. In this problem we will explore the idea of *group testing* as a strategy for testing a large population for a rare disease by pooling samples together. Suppose that we have a population of N people and we collect saliva samples from each of them. We are looking for genetic signatures of a particular virus in these samples. Let x_n denote the concentration of this material in the sample for the n^{th} person. We will assume that for healthy people $x_n = 0$, but for infected people, $x_n > 0$. Our testing procedure will be to form a series of M mixtures of samples from different subsets of people, and then only run tests on these mixtures. The goal here is to set $M < N$, and the question is then whether we can identify the infected people from the results of these tests.

We will consider the following approach:² we will form mixtures by constructing *random* combinations of samples, and we will attempt to recover the original \mathbf{x} using a simple convex optimization problem.

To mathematically represent the sampling/testing process, assume that we will ultimately run M tests, each of which will tell us the concentration of viral material in the combined sample being tested. For each person, their sample will be divided into $K < M$ equal portions, which will be assigned at random to the M tests. We will do this independently for each of the N people. We can ultimately represent the concentration of viral material in each of the mixed samples that we will ultimately test as a vector $\mathbf{y} \in \mathbb{R}^M$. We can write \mathbf{y} as

$$\mathbf{y} = \mathbf{A}\mathbf{x},$$

where \mathbf{A} is a matrix that represents the assignment of people to mixed samples/tests. Specifically, \mathbf{A} is a $M \times N$ matrix where each column is constructed independently by picking K entries at random, setting them to 1, and setting the remaining entries to 0.

Suppose there is no noise in our tests, so that we can estimate \mathbf{y} perfectly. Our inference problem is now to estimate \mathbf{x} given knowledge of \mathbf{y} and \mathbf{A} . In general, since $M < N$, recovering \mathbf{x} is impossible. However, when \mathbf{x} is sparse, meaning that it has only a few nonzeros (in this case meaning that most of the population is negative), then recovering \mathbf{x} is possible, although this fact was only broadly appreciated within the last 15 years or so.

²Just to be clear, this is not exactly what Georgia Tech is doing, but it is similar in spirit. The approach that Georgia Tech is somewhat simpler and builds on ideas that date back to World War II, when they were used to test soldiers for syphilis.

We will try to estimate \mathbf{x} by solving the following optimization problem:

$$\underset{\mathbf{x}}{\text{minimize}} \quad \|\mathbf{x}\|_1 \quad \text{subject to} \quad \mathbf{Ax} = \mathbf{y}, \quad \mathbf{x} \geq \mathbf{0}.$$

Below we will explore when and how well this works.

- (a) Suppose that you are testing a population of size $N = 1000$, but you can only process $M = 100$ tests. Assume that only 1% of the population is positive (meaning that there will be 10 infected individuals). Each person's sample will be split and added to $K = 10$ different batches. The file `hw01_prob09.py` contains code that sets up this problem. Use CVXPY to solve the optimization problem above and verify that this approach correctly identifies the 10 infected individuals.
- (b) Experiment with K . In practice, you might not want to divide a person's sample into too many tests. How low can you set K before things begin to fail?
- (c) Suppose that the prevalence of the disease begins to grow beyond 1%. How widespread can the disease become before the approach begins to fail (holding M and N fixed). As the disease becomes more widespread, you may need to adjust K . What value of K seems to work best when the spread of the disease is just below the threshold where the approach begins to fail?