

ECE 3803, Fall 2021

Homework #3

Due Tuesday, September 21, at 11:59pm

1. Prepare a one paragraph summary of what we talked about in class in the last week. 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. Recall the Tikhonov regularized least squares problem

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \delta\|\mathbf{x}\|_2^2.$$

In this problem we will justify the formulas for the solution to this problem given on page 50 in the notes.

- (a) Show that the Tikhonov regularized least squares problem can be rewritten in the format of a standard least squares problem

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \|\tilde{\mathbf{y}} - \tilde{\mathbf{A}}\mathbf{x}\|_2^2$$

for an appropriate choice of $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{y}}$.

- (b) Derive the solution $\hat{\mathbf{x}} = (\mathbf{A}^T \mathbf{A} + \delta \mathbf{I})^{-1} \mathbf{A}^T \mathbf{y}$ from the formula for the solution to the standard least squares problem together with the $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{y}}$ computed in part (a). [Hint: a useful fact here is that if

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}$$

then

$$\tilde{\mathbf{A}}^T \tilde{\mathbf{A}} = \mathbf{A}^T \mathbf{A} + \mathbf{B}^T \mathbf{B}.$$

You do not need to prove this, but you should convince yourself that it is true.]

- (c) Derive the formula $\hat{\mathbf{x}} = \mathbf{V}(\mathbf{\Sigma}^2 + \delta \mathbf{I})^{-1} \mathbf{\Sigma} \mathbf{U}^T \mathbf{y}$ by plugging $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ into the formula in part (b). In this problem, assume that $R = N$.¹ [Hint: recall that if $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are invertible matrices, then $(\mathbf{A}\mathbf{B}\mathbf{C})^{-1} = \mathbf{C}^{-1} \mathbf{B}^{-1} \mathbf{A}^{-1}$.]
- (d) In the notes I also claimed that another formula for the solution to the Tikhonov regularized least squares problem was $\hat{\mathbf{x}} = \mathbf{A}^T (\mathbf{A} \mathbf{A}^T + \delta \mathbf{I})^{-1} \mathbf{y}$. Show that by plugging $\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$, you arrive at the same formula as in part (c). In this problem, assume that $R = M$.

¹This result is actually true for any M , N , and R , but I'm only asking you to prove it in the full-rank case because it's a slightly simpler argument. The fact that this makes the argument is a hint. What extra fact about the SVD holds when $R = N$?

3. In the notes we showed that when using the truncated SVD, we could decompose the error that we incur as a result of noise into “noise error” and “approximation error” terms of the form

$$\hat{\mathbf{x}}_{\text{trunc}} - \hat{\mathbf{x}}_{\text{clean}} = \sum_{r=1}^{R'} \frac{1}{\sigma_r} \langle \mathbf{e}, \mathbf{u}_r \rangle \mathbf{v}_r + \sum_{r=R'+1}^R -\langle \mathbf{x}, \mathbf{v}_r \rangle \mathbf{v}_r.$$

Show that a similar decomposition is possible for Tikhonov regularization. Specifically, derive an expression of the form

$$\hat{\mathbf{x}}_{\text{tik}} - \hat{\mathbf{x}}_{\text{clean}} = \sum_{r=1}^R \frac{\sigma_r}{\sigma_r^2 + \delta} \langle \mathbf{e}, \mathbf{u}_r \rangle \mathbf{v}_r + \sum_{r=1}^R \frac{\sigma_r}{\sigma_r^2 + \delta} \langle \mathbf{x}, \mathbf{v}_r \rangle \mathbf{v}_r.$$

4. In class we saw that when the smallest singular value σ_R of a matrix \mathbf{A} gets very small, the least squares solution can be extremely sensitive to even small perturbations. The two strategies that we discussed for handling this were using a truncated SVD and Tikhonov regularization.

In this problem I would like you to explore using these ideas to “stabilize” the reconstruction problem from problem 7(c) in Homework #2.

- (a) First consider the truncated SVD approach. Specifically, set $M = 21$ and form the system of equations $\mathbf{y} = \mathbf{X}\mathbf{a}$ required to interpolate a polynomial of degree 20. However, rather than computing $\hat{\mathbf{a}} = \mathbf{X}^\dagger \mathbf{y} = \mathbf{X}^{-1} \mathbf{y}$, compute the SVD of \mathbf{X} and then compute

$$\hat{\mathbf{a}}_{\text{trunc}} = \mathbf{V}_{R'} \boldsymbol{\Sigma}_{R'}^{-1} \mathbf{U}_{R'}^T \mathbf{y}$$

for choices of R' ranging from 1 all the way up to 21. Comment on which value of R' gives the best estimate of the original $f(x)$? How does this compare to the result you obtained in problem 7(c) in the last homework where you instead chose a smaller value of M ? Look at the singular values of \mathbf{X} (or the singular values for the \mathbf{X} matrix formed when M is set to be smaller than 21) and comment if you observe anything interesting.

- (b) Next repeat this process but using Tikhonov regularization. Consider a range of values for the parameter δ and comment on which value works best. Compare these results to those from part (a).

5. In this problem you will write Python functions that implement two variants of gradient descent for the least squares problem. Each function should take as input a generic matrix \mathbf{A} and vector \mathbf{y} , an initial guess \mathbf{x}_0 , and stopping criteria (I would include both a tolerance ϵ as well as a maximum number of iterations, just in case). While you wouldn’t normally save the results of each iteration, to help you understand what is going on, your code should return an array containing the sequence of iterates $\mathbf{x}_0, \mathbf{x}_1, \dots$

As I said above, I would like you to implement two variants. One should be what was described in the notes. (You can implement either version 1 or 2 – the problems we will be working on here are not going to be big enough for the computational difference to be noticeable. Remember that the “inputs” to those algorithms should be $\mathbf{H} = \mathbf{A}^T \mathbf{A}$ and $\mathbf{b} = \mathbf{A}^T \mathbf{y}$.)

The other variant should be gradient descent with α_k set to a fixed constant. In the context of least squares, we can work out analytically what a good value for each step size is, but we will see later that in many optimization problems, finding such an analytical solution may not be possible. Here we will see what kind of impact this can have. Note that your version with a fixed step size will need an additional input (the step size α).

- (a) First let's test both versions on the simplest possible least squares problem I can think of

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}.$$

Hopefully just by looking at this problem you can see that the solution to

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$$

will just be $\hat{\mathbf{x}} = \mathbf{y} = [1, 1]^T$. Verify that this is indeed what both versions of your code produces when using an initial estimate of $\mathbf{x}_0 = \mathbf{0}$. Report how many iterations are required for both versions, reporting the results for several different values of the step size in the fixed step size version. What step size gives you the best performance here?

- (b) Let's make the problem a little more interesting and consider

$$\mathbf{A} = \begin{bmatrix} 1 & 0.99 \\ 0.99 & 1 \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} 1.99 \\ 1.99 \end{bmatrix}.$$

It may not be *quite* as obvious this time, but it should be clear that in this case we still have a solution of $[1, 1]^T$. Compute the condition number $\kappa(\mathbf{A})$. Based on our discussion in class, what does this tell us about how many iterations we might expect to need compared to the previous problem? Try out both versions of the algorithm and report how many iterations are required. If the answer is not what you expected, can you hypothesize why?

- (c) Let's consider one more tweak of this problem:

$$\mathbf{A} = \begin{bmatrix} 1 & 0.99 \\ 1 & 1.01 \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} 1.99 \\ 2.01 \end{bmatrix}.$$

The solution remains $[1, 1]^T$. Again, compute the condition number $\kappa(\mathbf{A})$, and try out both versions of the algorithm, reporting how many iterations are required. What is different from the previous problem? (If you are struggling to explain what is happening, it may help to plot the iterates together with a contour plot of the least squares objective function, and/or to experiment with different starting points for \mathbf{x}_0 .)

- (d) In part (c), you may need to set the tolerance parameter ϵ to be very small in order to avoid having the algorithm terminate at a point \mathbf{x}_k where \mathbf{x}_k is still quite far from the true solution. Explain why this is necessary?

- (e) Let's consider one last version of this problem:

$$\mathbf{A} = \begin{bmatrix} 2 & -1 \\ 0 & -1 \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} -1 \\ -1 \end{bmatrix}.$$

Again, the solution remains $[1, 1]^T$. As before, compute the condition number $\kappa(\mathbf{A})$, and try out both versions of the algorithm, reporting how many iterations are required. Comment on how many iterations are required compared to part (c), and if possible, provide an intuitive explanation for the difference.