DS-GA 1014 - Homework 10

Eric Niblock

November 15th, 2020

1. **(2 points). Let $A \in \mathbb{R}^{n\times m}$ and $\vec{y} \in \mathbb{R}^n$. We consider the least square problem:**

$$\textbf{minimize} \quad ||A\vec{x} - \vec{y}||^2 \quad \textbf{with respect to } \vec{x} \in \mathbb{R}^m$$

We know from the lecture that $\vec{x}^{LS} \stackrel{def}{=} A^\dagger \vec{y}$ is a solution of the above.

(a) **Show that $\vec{x}^{LS} \perp Ker(A)$.**

We know that we can express A and $A^\dagger$ as the following,

$$\begin{aligned} A &= U\Sigma V^T \\ A^\dagger &= V\Sigma^{\text{`}} U^T \end{aligned} \tag{1}$$

Now, we know from previous work that the first r columns of $U \in \mathbb{R}^{n\times n}$, $\vec{u}_1, ..., \vec{u}_r$, form a basis for the image of A, and the last $m-r$ rows of $V \in \mathbb{R}^{m\times m}$, $\vec{v}_{r+1}, ..., \vec{v}_m$, form a basis for the kernel of A. Furthermore, for $i \in \{1, ..., r\}$, we must have that $v_i \perp Ker(A)$, since the rows of an orthogonal matrix V are all orthogonal.

We know that $\vec{x}^{LS}$ can be expressed as a linear combination of $\vec{v}_1, ..., \vec{v}_r$, since these vectors form a basis for the image of $A^\dagger$, as shown from (1), and we know that $\vec{x}^{LS}$ is in the image of $A^\dagger$ from the statement. Since $\vec{x}^{LS}$ can be expressed as a linear combination of $\vec{v}_1, ..., \vec{v}_r \perp Ker(A)$, it is clear that $\vec{x}^{LS} \perp Ker(A)$.

(b) **Deduce that $\vec{x}^{LS}$ is the solution of the least square equation that has the smallest (Euclidean) norm.**

We know that the full set of solutions to the minimization problem is given by,

$$\{A^\dagger \overrightarrow{\mathbf{y}} + \overrightarrow{\mathbf{v}} \mid \overrightarrow{\mathbf{v}} \in Ker(A)\} \tag{2}$$

Since we have that $\overrightarrow{\mathbf{x}}^{LS} \perp Ker(A)$, we also have that, $A^\dagger \overrightarrow{\mathbf{y}} \perp \overrightarrow{\mathbf{v}}$ for all $\overrightarrow{\mathbf{v}} \in Ker(A)$. Now, take an alternate solution from the set, $\overrightarrow{\mathbf{x}}^* = A^\dagger \overrightarrow{\mathbf{y}} + \overrightarrow{\mathbf{v}}$. Our goal is to show that,

$$||\overrightarrow{\mathbf{x}}^{LS}|| < ||\overrightarrow{\mathbf{x}}^*|| \tag{3}$$

$$||A^\dagger \overrightarrow{\mathbf{y}}|| < ||A^\dagger \overrightarrow{\mathbf{y}} + \overrightarrow{\mathbf{v}}|| \tag{4}$$

But since $A^\dagger \overrightarrow{\mathbf{y}} \perp \overrightarrow{\mathbf{v}}$ for all $\overrightarrow{\mathbf{v}} \in Ker(A)$, we can write,

$$||A^\dagger \overrightarrow{\mathbf{y}}|| < ||A^\dagger \overrightarrow{\mathbf{y}} + \overrightarrow{\mathbf{v}}|| = ||A^\dagger \overrightarrow{\mathbf{y}}|| + ||\overrightarrow{\mathbf{v}}|| \tag{5}$$

Which is clealy true. Therefore, $\overrightarrow{\mathbf{x}}^{LS}$ is the solution of the least square equation that has the smallest (Euclidean) norm.

2. (2 points). Let $A \in \mathbb{R}^{n \times d}$ and $\overrightarrow{\mathbf{y}} \in \mathbb{R}^{n}$. The Ridge regression adds a ℓ_2 penalty to the least square problem:

$$\textbf{minimize} \quad ||A\overrightarrow{\mathbf{x}} - \overrightarrow{\mathbf{y}}||^2 + \lambda||\overrightarrow{\mathbf{x}}||^2 \quad \textbf{with respect to } \overrightarrow{\mathbf{x}} \in \mathbb{R}^{m}$$

for some penalization parameter $\lambda > 0$. Show that the above admits a unique solution given by

$$\overrightarrow{\mathbf{x}}^{Ridge} = (A^T A + \lambda Id_n)^{-1} A^T \overrightarrow{\mathbf{y}}$$

We define $f(\overrightarrow{\mathbf{x}}) = ||A\overrightarrow{\mathbf{x}} - \overrightarrow{\mathbf{y}}||^2 + \lambda||\overrightarrow{\mathbf{x}}||^2$. We know that $f(\overrightarrow{\mathbf{x}})$ is convex because the sum of convex functions in convex. Therefore, it must admit a minimum. Then, we have that,

$$\nabla f(\overrightarrow{\mathbf{x}}) = A^T(A\overrightarrow{\mathbf{x}} - \overrightarrow{\mathbf{y}}) + \lambda\overrightarrow{\mathbf{x}} \tag{6}$$

We then set the gradient equal to zero and solve for $\overrightarrow{\mathbf{x}}$. This yields,

$$A^T A\overrightarrow{\mathbf{x}} - A^T\overrightarrow{\mathbf{y}} + \lambda\overrightarrow{\mathbf{x}} = 0 \tag{7}$$

$$A^T A\overrightarrow{\mathbf{x}} + \lambda\overrightarrow{\mathbf{x}} = A^T\overrightarrow{\mathbf{y}} \tag{8}$$

$$(A^T A + \lambda Id_n)\overrightarrow{\mathbf{x}} = A^T\overrightarrow{\mathbf{y}} \tag{9}$$

Then, we know that $(A^T A + \lambda Id_n)$ is invertible for some choice of $\lambda > 0$, because as a previous result, we had that, for any symmetric matrix M, there exists $\lambda > 0$ such that the matrix $M + \lambda Id_n$ is positive definite. We know that $A^T A$ is symmetric, and therefore, there is some $\lambda > 0$ which forces $(A^T A + \lambda Id_n)$ to be positive definite. Furthermore, any positive definite matrix is invertible. So, for some choice of λ we have,

$$\overrightarrow{\mathbf{x}}^{Ridge} = (A^T A + \lambda Id_n)^{-1} A^T\overrightarrow{\mathbf{y}} \tag{10}$$

3. (3 points). Recall that $||M||_{Sp}$ denotes the spectral norm of a matrix M.

(a) Let $A \in \mathbb{R}^{n\times m}$. Show that for all $\overrightarrow{\mathbf{x}} \in \mathbb{R}^m$,

$$||A\overrightarrow{\mathbf{x}}|| \leq ||A||_{Sp}||\overrightarrow{\mathbf{x}}||$$

Since we know that every $A = U\Sigma V^T$, with orthogonal matrices U and V^T, and diagonal matrix Σ, we have that,

$$A^T A = V\Sigma^T U^T U\Sigma V^T = V\Sigma^{(2)}V^T \tag{11}$$

Where $\Sigma^{(2)}$ holds the square of the diagonal values of Σ, and hence the eigenvalues of $A^T A$. Then,

$$||A\overrightarrow{\mathbf{x}}|| = \overrightarrow{\mathbf{x}}^T A^T A\overrightarrow{\mathbf{x}} \tag{12}$$

Since $\overrightarrow{\mathbf{x}} \in \mathbb{R}^m$ and $V \in \mathbb{R}^{m\times m}$ and orthogonal, we have that all $\overrightarrow{\mathbf{x}}$ can be expressed as a combination of the columns of V. So,

$$\begin{aligned}
||A\overrightarrow{\mathbf{x}}||^2 = \overrightarrow{\mathbf{x}}^T A^T A\overrightarrow{\mathbf{x}} &= \left(\sum_{i=1}^{min(m,n)} c_i \overrightarrow{\mathbf{u}}_i^T\right) A^T A \left(\sum_{i=1}^{min(m,n)} c_i \overrightarrow{\mathbf{u}}_i\right) \\
&= \left(\sum_{i=1}^{min(m,n)} c_i \overrightarrow{\mathbf{u}}_i^T\right) \left(\sum_{i=1}^{min(m,n)} c_i (A^T A\overrightarrow{\mathbf{u}}_i)\right) \\
&= \left(\sum_{i=1}^{min(m,n)} c_i \overrightarrow{\mathbf{u}}_i^T\right) \left(\sum_{i=1}^{min(m,n)} c_i \lambda_i \overrightarrow{\mathbf{u}}_i\right) \\
&= \sum_{i=1}^{min(m,n)} c_i^2 \lambda_i \overrightarrow{\mathbf{u}}_i^T \overrightarrow{\mathbf{u}}_i = \sum_{i=1}^{min(m,n)} c_i^2 \lambda_i
\end{aligned} \tag{13}$$

And therefore, we are entitled to write,

$$\lambda_{min} \sum_{i=1}^{min(m,n)} c_i^2 \leq \sum_{i=1}^{min(m,n)} c_i^2 \lambda_i = ||A\overrightarrow{\mathbf{x}}||^2 \leq \lambda_{max} \sum_{i=1}^{min(m,n)} c_i^2 \tag{14}$$

Where the first part of the equation merely served as a point of comparison. Then, taking the square root,

$$||A\overrightarrow{\mathbf{x}}|| \leq \sqrt{\lambda_{max} \sum_{i=1}^{min(m,n)} c_i^2} = ||A||_{Sp}||\overrightarrow{\mathbf{x}}|| \tag{15}$$

(b) Show that for all $A \in \mathrm{I\!R}^{\mathrm{n \times m}}$ and $B \in \mathrm{I\!R}^{\mathrm{m \times k}}$:

$$||AB||_{Sp} \leq ||A||_{Sp}||B||_{Sp}$$

From the previous result, we know if we take some $\overrightarrow{\mathbf{u}} \in \mathrm{I\!R}^{\mathrm{k}}$ such that $||\overrightarrow{\mathbf{u}}|| = 1$, then we have,

$$||AB\overrightarrow{\mathbf{u}}|| \leq ||A||_{Sp}||B\overrightarrow{\mathbf{u}}|| \leq ||A||_{Sp}||B||_{Sp}||\overrightarrow{\mathbf{u}}|| = ||A||_{Sp}||B||_{Sp} \tag{16}$$

So, we have, for any $\overrightarrow{\mathbf{u}} \in \mathrm{I\!R}^{\mathrm{k}}$ such that $||\overrightarrow{\mathbf{u}}|| = 1$,

$$||AB\overrightarrow{\mathbf{u}}|| \leq ||A||_{Sp}||B||_{Sp} \tag{17}$$

Now if we take the value of $\overrightarrow{\mathbf{u}} \in \mathrm{I\!R}^{\mathrm{k}}$ such that $||\overrightarrow{\mathbf{u}}|| = 1$ which maximizes $||AB\overrightarrow{\mathbf{u}}||$, this is precisely the definition of the spectral norm of $||AB||$. Thus,

$$||AB||_{Sp} \leq ||A||_{Sp}||B||_{Sp} \tag{18}$$

(c) Is it true that for all $n, m, k \geq 1$, all $A \in \mathrm{I\!R}^{\mathrm{n \times m}}$ and $B \in \mathrm{I\!R}^{\mathrm{m \times k}}$:

$$||AB||_F \leq ||A||_F||B||_F$$

Give a proof or a counter-example.

First, note that,

$$||AB||_F^2 = \sum_{i=1}^{n}\sum_{j=1}^{k}(\overrightarrow{\mathbf{a}}_i^T\overrightarrow{\mathbf{b}}_j)^2 \tag{19}$$

Where $\overrightarrow{\mathbf{a}}_i^T$ and $\overrightarrow{\mathbf{b}}_j$ are the i-th row and j-th column of A and B, respectively. By Cauchy-Shwartz, we see that,

$$\begin{aligned}\sum_{i=1}^{n}\sum_{j=1}^{k}(\overrightarrow{\mathbf{a}}_i^T\overrightarrow{\mathbf{b}}_j)^2 &\leq \sum_{i=1}^{n}\sum_{j=1}^{k}||\overrightarrow{\mathbf{a}}_i||^2||\overrightarrow{\mathbf{b}}_j||^2 \\ &= \sum_{i=1}^{n}||\overrightarrow{\mathbf{a}}_i||^2\sum_{j=1}^{k}||\overrightarrow{\mathbf{b}}_j||^2 \\ &= ||A||_F^2||B||_F^2\end{aligned} \tag{20}$$

So, summarizing,

$$||AB||_F^2 = \sum_{i=1}^{n}\sum_{j=1}^{k}(\overrightarrow{\mathbf{a}}_i^T\overrightarrow{\mathbf{b}}_j)^2 \leq ||A||_F^2||B||_F^2 \tag{21}$$

4. **(3 points). Consider the 5×4 matrix A and $\vec{y} \in \mathbb{R}^5$ given by:**

$$A = \begin{bmatrix} 1.1 & -2.3 & 1.7 & 4.5 \\ 1.7 & 1.6 & 3.8 & 0.3 \\ 1.0 & 0.1 & 1.3 & 0.2 \\ -0.5 & -0.4 & 0 & -1.3 \\ -0.5 & 2.9 & -0.3 & 2.0 \end{bmatrix} \quad \textbf{\textit{and}} \quad \vec{y} = \begin{bmatrix} -13.8 \\ -2.7 \\ 9.6 \\ -2.4 \\ 3.9 \end{bmatrix} \tag{22}$$

In each of the following questions, it is intended that you solve the problem using the programming language of your choice and only report the numerical answer to 3 decimal places, without including your code files in your submission.

(a) **Compute the minimizer $\vec{x}^* \in \mathbb{R}^4$ of**

$$||A\vec{x} - \vec{y}||$$

We find during our use of Python that,

```
import numpy as np

A = np.array([[1.1,-2.3,1.7,4.5],[1.7,1.6,3.8,0.3],[1,0.1,1.3,0.2],[-0.5,-0.4,0,-1.3],[-0.5,2.9,-0.3,2]])
y = np.array([-13.8,-2.7,9.6,-2.4,3.9])
```

```
A.T@A
```

```
array([[ 5.6 , -0.96,  9.78,  5.31],
       [-0.96, 16.43,  1.43, -3.53],
       [ 9.78,  1.43, 19.11,  8.45],
       [ 5.31, -3.53,  8.45, 26.07]])
```

```
inv = np.linalg.inv(A.T@A)
xmin = inv@A.T@y
print('X* is given by: ', xmin)
```

```
X* is given by:  [14.07409328  3.60631035 -7.87274476 -1.74716021]
```

(b) **Find a vector $\vec{v} \in \mathbb{R}^5$ with $v_1 > 0$ and $||\vec{v}|| = 1$ such that the minimizer of**

$$||A\vec{x} - (\vec{y} + \vec{v})||$$

is also $\vec{x}^*$.

Here, we find the solution for vector $\vec{v}$, and check out result,

```
import scipy.linalg
v = scipy.linalg.null_space(A.T)
print('v is given by: ',v.T[0])
```

```
v is given by:  [ 0.03521722 -0.26599821  0.78205836  0.51802759  0.21917308]
```

And then we check our solution and find we yield the same minimum as the previous part,

```
newy = y+v.T[0]
xminnew = inv@A.T@newy
print('The new minimum is given by: ', xmin)
```

```
The new minimum is given by:  [14.07409328  3.60631035 -7.87274476 -1.74716021]
```

(c) Find a vector $\vec{w} \in \mathbb{R}^5$ with $w_1 > 0$ and $||\vec{w}|| = 1$ such that the minimizer $\vec{x}'$ of

$$||A\vec{x} - (\vec{y} + \vec{w})||$$

maximizes the error $||\vec{x}^* - \vec{x}'||$ and also give the resulting error. That is, we are trying to corrupt the vector $\vec{y}$ with a fixed amount of noise $\vec{w}$ that maximally modifies the least squares solution.

We find the solution given below, and then we check this solution programmatically,

```
inv = np.linalg.inv(A.T@A)
Adag = inv@A.T
M = Adag.T@Adag
l,v = np.linalg.eigh(M)
print('w is given by: ',-1*v[:,-1])
```

```
w is given by:  [ 0.17718915  0.17894147 -0.52073406  0.80126178  0.15296916]
```

Here, we check the solution using 10,000 random vectors of unit norm as a comparison, ensuring that each random vector produces a norm less than our theorized maximum norm.

```python
from random import gauss

def make_rand_vector(dims):
    vec = [gauss(0, 1) for i in range(dims)]
    mag = sum(x**2 for x in vec) ** .5
    if vec[0] < 0:
        vec[0] = -1*vec[0]
    return np.array([x/mag for x in vec])
```

```python
maxnorm = np.linalg.norm(inv@A.T@y - inv@A.T@(y+v[:,-1]))
for i in range(10000):
    t = make_rand_vector(5)
    if np.linalg.norm(inv@A.T@y - inv@A.T@(y+t)) > maxnorm:
        print(np.linalg.norm(inv@A.T@y - inv@A.T@(y+t)))
        print(t)
        break
    if i==9999:
        print('Success!')
```

```
Success!
```