CMDA 3606 · MATHEMATICAL MODELING II

Problem Set 7

Posted 20 March 2019. Due at 5pm on Thursday, 28 March 2019.

Basic guidelines: Students may discuss the problems on this assignment, but each student must submit his or her individual writeup and code. (In particular, you *must write up your own individual MATLAB code.*) Students may consult class notes and other online resources for general information; cite all your sources and list those with whom you have discussed the problems.

1. [35 points: 10 points each for (a), (b); 7 points for (c); 8 points for (d)]

   (a) Consider the matrix and vector

   $$\mathbf{A} = \begin{bmatrix} 1 & 0 & -1 \\ -1 & 0 & 1 \end{bmatrix}, \qquad \mathbf{b} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}.$$

   (i) Compute (by hand) the full SVD of $\mathbf{A}$.

   (ii) Let r denote the rank of $\mathbf{A}$. Compute the pseudoinverse

   $$\mathbf{A}^+ = \sum_{j=1}^{r} \frac{1}{\sigma_j} \mathbf{v}_j \mathbf{u}_j^T.$$

   (iii) Any solution $\mathbf{x}$ of the least squares problem

   $$\min_{\mathbf{x} \in \mathbb{R}^3} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|$$

   must have the form $\mathbf{x} = \mathbf{A}^+\mathbf{b} + \mathbf{c}$, where $\mathbf{c} \in \mathcal{N}(\mathbf{A})$.
   Write down a formula for all such solutions $\mathbf{x}$ for this $\mathbf{A}$ and $\mathbf{b}$.

   (iv) Confirm that, among all the solutions identified in (iii), $\|\mathbf{x}\|$ is minimized by $\mathbf{x} = \mathbf{A}^+\mathbf{b}$.

   (b) Now for any $\varepsilon > 0$, consider the modified problem

   $$\mathbf{A}_\varepsilon = \begin{bmatrix} 1 & \varepsilon & -1 \\ -1 & \varepsilon & 1 \end{bmatrix}, \qquad \mathbf{b} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}.$$

   (i) Compute (by hand) the full SVD of $\mathbf{A}_\varepsilon$.

   (ii) Write down a formula for the pseudoinverse $\mathbf{A}_\varepsilon^+$.

   (iii) Write down all $\mathbf{x}_\varepsilon$ that solve

   $$\min_{\mathbf{x}_\varepsilon \in \mathbb{R}^3} \|\mathbf{b} - \mathbf{A}_\varepsilon \mathbf{x}_\varepsilon\|$$

   and specify the $\mathbf{x}_\varepsilon$ that minimizes $\|\mathbf{x}_\varepsilon\|$.

   (c) Now compare your results from parts (a) and (b): Let $\mathbf{x}$ and $\mathbf{x}_\varepsilon$ denote the minimum norm solutions from (a.iv) and (b.iii).
   Does $\mathbf{A}_\varepsilon^+ \to \mathbf{A}^+$ as $\varepsilon \to 0$? Does $\mathbf{x}_\varepsilon \to \mathbf{x}$ as $\varepsilon \to 0$?
   Based on your answer to these questions: Is the minimum-norm solution to a least squares problem a continuous function of the entries of $\mathbf{A}$? What are the implications of this observation for more practical problems where $\mathbf{A}$ could be polluted by rounding errors, data noise, etc.?

   (d) For arbitrary $\mathbf{A} \in \mathbb{R}^{m \times n}$, show that $\mathbf{A}^+ = \lim_{\varepsilon \to 0} (\mathbf{A}^T\mathbf{A} + \varepsilon \mathbf{I})^{-1}\mathbf{A}^T$, where

   $$\mathbf{A}^+ = \sum_{j=1}^{r} \frac{1}{\sigma_j} \mathbf{v}_j \mathbf{u}_j^T.$$

2. [35 points: 5 points per part]

Consider the matrix and vector

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 0 & 2 \\ 2 & 0 \end{bmatrix}, \qquad \mathbf{b} = \begin{bmatrix} 1 \\ -2 \\ 0 \end{bmatrix}.$$

(This problem was designed to be solved by hand; you can use MATLAB to assist if you need to.)

(a) Find the unique vector $\mathbf{x} = (\mathbf{A}^T\mathbf{A})^{-1}\mathbf{A}^T\mathbf{b}$ that solves

$$\min_{\mathbf{x}\in\mathbb{R}^n} \|\mathbf{b} - \mathbf{A}\mathbf{x}\|.$$

Recall that $\mathbf{b}$ can be written in the form $\mathbf{b} = \mathbf{b}_R + \mathbf{b}_N$, where $\mathbf{b}_R \in \mathcal{R}(\mathbf{A})$ and $\mathbf{b}_N \in \mathcal{N}(\mathbf{A}^T)$, and that the least squares solution $\mathbf{x}$ satisfies $\mathbf{A}\mathbf{x} = \mathbf{b}_R$.

(b) For the particular $\mathbf{b}$ given above, compute $\mathbf{b}_R$ and $\mathbf{b}_N$.

(c) Compute $\|\mathbf{b}_N\|$ (recall that $\min_{\mathbf{x}\in\mathbb{R}^n} \|\mathbf{b} - \mathbf{A}\mathbf{x}\| = \|\mathbf{b}_N\|$).

The standard least squares algorithm you solved in part (a) can be viewed as a finding the vector $\mathbf{e}$ of smallest norm so that $\mathbf{b} + \mathbf{e} \in \mathcal{R}(\mathbf{A})$. (In this case, $\mathbf{e}$ is simply $\mathbf{b}_N$.) In data fitting applications, this set-up implicitly assumes that $\mathbf{b}$ contains all errors in the data, and $\mathbf{A}$ is known perfectly. In cases where $\mathbf{A}$ might be polluted with error too, we could instead find the smallest $\mathbf{e}$ *and* $\mathbf{E}$ such that $\mathbf{b} + \mathbf{e} \in \mathcal{R}(\mathbf{A} + \mathbf{E})$. This is called the *Total Least Squares* (TLS) problem. We measure the size of the errors $\mathbf{E}$ and $\mathbf{e}$ simultaneously as $\|[\mathbf{E}\ \mathbf{e}]\|$, where $\|[\mathbf{E}\ \mathbf{e}]\|$ denotes the largest singular value of the $m \times (n+1)$ matrix $[\mathbf{E}\ \mathbf{e}]$.

To solve the TLS problem, we seek the minimal $\|[\mathbf{E}\ \mathbf{e}]\|$ such that

$$(\mathbf{A} + \mathbf{E})\mathbf{x} = (\mathbf{b} + \mathbf{e})$$

has a solution. This equation is equivalent to

$$[\,\mathbf{A} + \mathbf{E} \quad \mathbf{b} + \mathbf{e}\,] \begin{bmatrix} \mathbf{x} \\ -1 \end{bmatrix} = \mathbf{0}.$$

So, we seek the smallest $\|[\mathbf{E}\ \mathbf{e}]\|$ that gives $[\mathbf{A} + \mathbf{E}\ \ \mathbf{b} + \mathbf{e}]$ a nontrivial null space.

If the SVD of $[\mathbf{A}\ \ \mathbf{b}]$ is given by

$$[\mathbf{A}\ \ \mathbf{b}] = \sum_{j=1}^{n+1} \sigma_j \mathbf{u}_j \mathbf{v}_j^T,$$

then the smallest change $[\mathbf{E}\ \mathbf{e}]$ that makes $[\mathbf{A} + \mathbf{E}\ \ \mathbf{b} + \mathbf{e}]$ have such a nullspace is given by

$$[\mathbf{E}\ \mathbf{e}] = -\sigma_{n+1}\mathbf{u}_{n+1}\mathbf{v}_{n+1}^T,$$

giving the error $\|[\mathbf{E}\ \mathbf{e}]\| = \sigma_{n+1}$.

(d) Find the singular value decomposition of $[\mathbf{A}\ \mathbf{b}]$ for the $\mathbf{A}$ and $\mathbf{b}$ given above.

(e) Compute $[\mathbf{E}\ \mathbf{e}] = -\sigma_3\mathbf{u}_3\mathbf{v}_3^T$ for this example.

(f) Find the solution $\mathbf{x}$ to $(\mathbf{A} + \mathbf{E})\mathbf{x} = (\mathbf{b} + \mathbf{e})$ for this example.

(g) How does the error $\|[\mathbf{E}\ \mathbf{e}]\| = \sigma_{n+1}$ for this example compare to the error $\|\mathbf{b}_N\|$ for the standard least squares problem you computed in part (c)?

3. [30 points: 8 points each for (a), (b), (c); 6 points for (d)]

In many engineering applications, one can explore a physical system by *measuring* values $y_1, y_2, \ldots \in \mathbb{R}$, where y_k is believed to obey some unknown model of the form

$$y_k = \sum_{j=1}^{r} \gamma_j \lambda_j^k.$$

(For example, the y_k could record the vibration of a point on a bridge, measured by an accelerometer at regularly spaced time intervals $t_k = k\Delta t$.)

The first step to modeling this system is to figure out the correct value of r (the *order* of the model).

Given measurements of the $2n-1$ values $y_1, \ldots, y_{2n-1}$, one can compute the order r as follows. Form the *Hankel matrix*

$$\mathbf{H}_n = \begin{bmatrix} y_1 & y_2 & y_3 & \cdots & y_n \\ y_2 & y_3 & \iddots & \iddots & \vdots \\ y_3 & \iddots & \iddots & \iddots & y_{2n-3} \\ \vdots & \iddots & \iddots & y_{2n-3} & y_{2n-2} \\ y_n & \cdots & y_{2n-3} & y_{2n-2} & y_{2n-1} \end{bmatrix}.$$

Then, so long as $n \geq r$, classical theory dictates that

$$r = \text{rank}(\mathbf{H}) = \#\text{ of nonzero singular values of }\mathbf{H}.$$

(I would gladly explain the neat theory justifying this statement to any curious students.)

Imagine an engineer approaches you, explains this theory, and asks you to use it to compute the order of his/her system.

(a) Test this theory out for a small example. Let

$$y_k = 1^k + \left(\frac{1}{2}\right)^k,$$

so that $r = 2$. In MATLAB, form and compute the singular values of the matrices

$$\mathbf{H}_3 = \begin{bmatrix} y_1 & y_2 & y_3 \\ y_2 & y_3 & y_4 \\ y_3 & y_4 & y_5 \end{bmatrix}, \quad \mathbf{H}_4 = \begin{bmatrix} y_1 & y_2 & y_3 & y_4 \\ y_2 & y_3 & y_4 & y_5 \\ y_3 & y_4 & y_5 & y_6 \\ y_4 & y_5 & y_6 & y_7 \end{bmatrix}, \quad \mathbf{H}_5 = \begin{bmatrix} y_1 & y_2 & y_3 & y_4 & y_5 \\ y_2 & y_3 & y_4 & y_5 & y_6 \\ y_3 & y_4 & y_5 & y_6 & y_7 \\ y_4 & y_5 & y_6 & y_7 & y_8 \\ y_5 & y_6 & y_7 & y_8 & y_9 \end{bmatrix}.$$

Confirm that, in each case, you get only two nonzero singular values.

(b) Now consider an example that is slightly more realistic. Let

$$y_k = \sum_{j=1}^{20} \left(\frac{1}{j}\right)^k,$$

so that $r = 20$. For this data, compute $\mathbf{H}_{50}$ and create a plot of the singular values:

```
semilogy(svd(H50),'.','markersize',20)
```

Does this calculation confirm that $\text{rank}(\mathbf{H}_{50}) = 20$?

Estimate a value for r based on your results, and explain why you decided on that value.

(c) Suppose we add noise to the previous example, so that now we compute

$$y_k = \sum_{j=1}^{20} \left(\frac{1}{j}\right)^k + \eta_k,$$

where $\eta_k = \texttt{1e} - 8 * \texttt{randn}$ in MATLAB. (Each measurement y_k is now polluted by random Gaussian noise, standard deviation 10^{-8}: use a different call to `randn` for each k.)

Like in part (b), form $\mathbf{H}_{50}$ (now with the noisy measurements) and plot the singular values. Now what value would you estimate for r?

(d) Investigate the `rank` command in MATLAB. The command

```
type rank
```

will display the code that MATLAB uses to compute rank. Explain the five lines of this code.