STABILITY AND NUMERICAL ASPECTS OF LEAST SQUARES

Dr. Matthieu R Bloch

Monday, November 22, 2021

LOGISTICS

General announcements

- Assignment 6 to be posted… (grading traffic jam)
- 4 lectures left! (No lecture on Wednesday November 24, 2021)

Midterm 2

- 99% graded, grades released after Thanksiving weekend

LOGISTICS

General announcements

- Assignment 6 to be posted... (grading traffic jam)
- 4 lectures left! (No lecture on Wednesday November 24, 2021)

Midterm 2

- 99% graded, grades released after Thanksiving weekend
- Midterm solution during office hours on Tuesday November 23, 2021

Assignment 5

- Grading finalized

Assignment 6 and 7

- Posted this week, due date on Monday December 6 and Monday December 13

WHAT'S ON THE AGENDA FOR TODAY?

Last time:

- Stability of least squares

Today:

- More on stability and numerical considerations

Reading: lecture notes 14/15/16

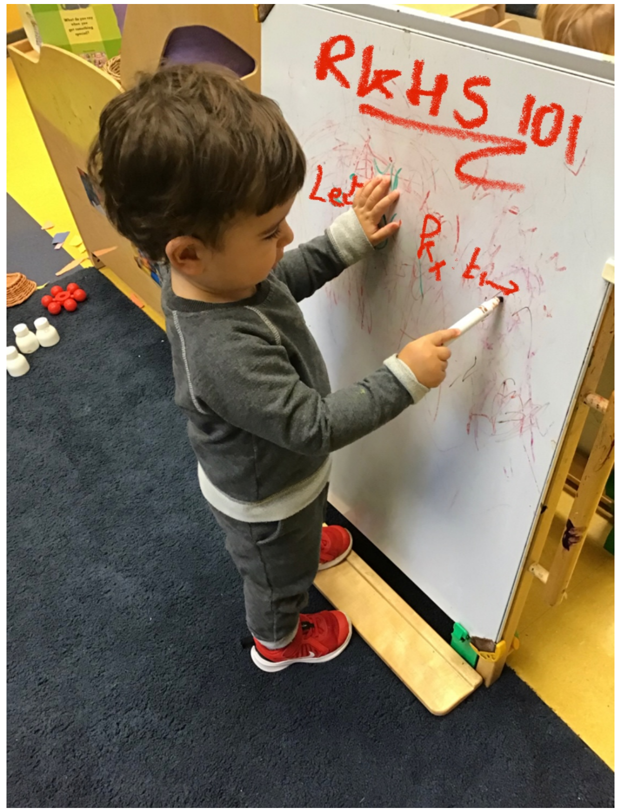

Toddlers can do it!

STABILITY OF LEAST SQUARES

What if we observe $\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{e}$ and we apply the pseudo inverse? $\hat{\mathbf{x}} = \mathbf{A}^+\mathbf{y}$

We can separate the error analysis into two components

$$\hat{\mathbf{x}} - \mathbf{x}_0 = \underbrace{\mathbf{A}^+\mathbf{A}\mathbf{x}_0 - \mathbf{x}_0}_{\text{null space error}} + \underbrace{\mathbf{A}^+\mathbf{e}}_{\text{noise error}}$$

We will express the error in terms of the SVD $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\intercal$ With

- $\{\mathbf{v}_i\}_{i=1}^r$ orthobasis of $\mathrm{row}(\mathbf{A})$, augmented by $\{\mathbf{v}_i\}_{i=1}^{r+1} \in \ker \mathbf{A}$ to form an orthobasis of $\mathbb{R}^n$
- $\{\mathbf{u}_i\}_{i=1}^r$ orthobasis of $\mathrm{col}(\mathbf{A})$, augmented by $\{\mathbf{u}\}_{i=1}^{r+1} \in \ker \mathbf{A}^\intercal$ to form an orthobasis of $\mathbb{R}^m$

The null space error is given by

$$\left\|\mathbf{A}^+\mathbf{A}\mathbf{x}_0 - \mathbf{x}_0\right\|_2^2 = \sum_{i=r+1}^{n} \left|\langle \mathbf{v}_i, x_0\rangle\right|^2$$

The noise error is given by

$$\left\|\mathbf{A}^+\mathbf{e}\right\|_2^2 = \sum_{i=1}^{r} \frac{1}{\sigma_i^2}\left|\langle \mathbf{e}, \mathbf{u}_i\rangle\right|^2$$

STABLE RECONSTRUCTION BY TRUNCATION

How do we mitigate the effect of small singular values in reconstruction?

$$\hat{\mathbf{x}} = \mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^\intercal \mathbf{y} = \sum_{i=1}^{r} \frac{1}{\sigma_i} \langle \mathbf{y}, \mathbf{u}_i \rangle \mathbf{v}_i$$

Truncate the SVD to $r' < r$

$$\mathbf{A}_t \triangleq \sum_{i=1}^{r'} \sigma_i \mathbf{u}_i \mathbf{v}_i^\intercal \qquad \mathbf{A}_t^+ = \sum_{i=1}^{r'} \frac{1}{\sigma_i} \mathbf{u}_i \mathbf{v}_i^\intercal$$

Reconstruct $\hat{\mathbf{x}}_t = \sum_{i=1}^{r'} \frac{1}{\sigma_i} \langle \mathbf{y}, \mathbf{u}_i \rangle \mathbf{v}_i = \mathbf{A}_t$

Error analysis:

$$\left\| \hat{\mathbf{x}}_t - \mathbf{x}_0 \right\|_2^2 = \sum_{i=r+1}^{n} \left| \langle \mathbf{x}_0, \mathbf{v}_i \rangle \right|^2 + \underbrace{\sum_{i=r'+1}^{r} \left| \langle \mathbf{x}_0, \mathbf{v}_i \rangle \right|^2 + \sum_{i=1}^{r'} \frac{1}{\sigma_i^2} \left| \langle \mathbf{e}, \mathbf{u}_i \rangle \right|^2}_{\text{Tradeoff}}$$

STABLE RECONSTRUCTION BY TRUNCATION

How do we mitigate the effect of small singular values in reconstruction?

$$\hat{\mathbf{x}} = \mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^\intercal \mathbf{y} = \sum_{i=1}^{r} \frac{1}{\sigma_i} \langle \mathbf{y}, \mathbf{u}_i \rangle \mathbf{v}_i$$

Truncate the SVD to $r' < r$

$$\mathbf{A}_t \triangleq \sum_{i=1}^{r'} \sigma_i \mathbf{u}_i \mathbf{v}_i^\intercal \qquad \mathbf{A}_t^+ = \sum_{i=1}^{r'} \frac{1}{\sigma_i} \mathbf{u}_i \mathbf{v}_i^\intercal$$

Reconstruct $\hat{\mathbf{x}}_t = \sum_{i=1}^{r'} \frac{1}{\sigma_i} \langle \mathbf{y}, \mathbf{u}_i \rangle \mathbf{v}_i = \mathbf{A}_t$

Error analysis:

$$\left\|\hat{\mathbf{x}}_t - \mathbf{x}_0\right\|_2^2 = \sum_{i=r+1}^{n} \left|\langle \mathbf{x}_0, \mathbf{v}_i \rangle\right|^2 + \sum_{i=r'+1}^{r} \left|\langle \mathbf{x}_0, \mathbf{v}_i \rangle\right|^2 + \sum_{i=1}^{r\,'} \frac{1}{\sigma_i^2} \left|\langle \mathbf{e}, \mathbf{u}_i \rangle\right|^2$$

STABLE RECONSTRUCTION BY REGULARIZATION

Regularization means changing the problem to solve

$$\min_{\mathbf{x}\in\mathbb{R}^n} \|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_2^2 \qquad \lambda > 0$$

The solution is

$$\hat{\mathbf{x}} = \underbrace{(\mathbf{A}^\intercal\mathbf{A}+\lambda\mathbf{I})^{-1}\mathbf{A}^\intercal\mathbf{y}}_{1} = \mathbf{V}(\mathbf{\Sigma}^2+\lambda\mathbf{I})^{-1}\mathbf{\Sigma}\mathbf{U}^\intercal\mathbf{y} = \sum^{r} \frac{\sigma_i}{\sigma_i^2+\lambda}\langle u_i, y\rangle v_i$$

(recall normal equations $(A^\intercal A + \lambda I)x = A^\intercal y$)

Scaling by $\frac{\sigma_i}{\sigma_i^2+\lambda}$ for every $i \in [\![1;r]\!]$

If $\sigma_i^2 \ll \lambda$ then $\frac{\sigma_i}{\sigma_i^2+\lambda} \approx \frac{\sigma_i}{\lambda} \ll 1$

If $\sigma_i^2 \gg \lambda$ then $\frac{\sigma_i}{\sigma_i^2+\lambda} \approx \frac{1}{\sigma_i} \ll 1$

Proof of $(A^TA+\lambda I)^{-1}A^Ty = V(\Sigma^2+\lambda I)^{-1}\Sigma U^Ty$

Recall that $\tilde{V} = [V \mid V_0] \in \mathbb{R}^{n\times n}$

Recall $A = U\Sigma V^T \in \mathbb{R}^{m\times n}$, $A^TA \in \mathbb{R}^{n\times n}$

Introduce
$$\tilde{\Sigma} = \begin{pmatrix} \sigma_1 & & & & & \\ & \ddots & & & 0 & \\ & & \sigma_r & & & \\ & & & 0 & & \\ & 0 & & & \ddots & \\ & & & & & 0 \end{pmatrix} \in \mathbb{R}^{n\times n}$$

(in general $V^TV = I$ but $VV^T = I$)

Note that $A^TA + \lambda I_n = \tilde{V}\tilde{\Sigma}^2\tilde{V}^T + \lambda\tilde{V}\tilde{V}^T = \tilde{V}(\tilde{\Sigma}^2 + \lambda I)\tilde{V}^T$

Hence $(A^TA+\lambda I)^{-1}A^Ty = \tilde{V}\underbrace{(\tilde{\Sigma}^2+\lambda I)^{-1}}_{2}\underbrace{\tilde{V}^TV}_{1}\Sigma U^Ty$

Note $\tilde{V}^TV = [V \mid V_0]^TV = \begin{bmatrix} I_{r\times r} \\ \hline 0_{n-r\times r} \end{bmatrix}$ so that $\tilde{V}(\tilde{\Sigma}^2+\lambda I_n)^{-1}\tilde{V}^TV = V(\Sigma^2+\lambda I_r)^{-1}$

Error analysis: $\hat{x}_{tk} = \sum_{i=1}^{r} \frac{\sigma_i}{\sigma_i^2+\lambda} \langle u_i, y\rangle v_i$ w/ $y = Ax_0 + e$

$$= \sum_{i=1}^{r} \frac{\sigma_i}{\sigma_i^2+\lambda} \langle u_i, Ax_0\rangle v_i + \sum_{i=1}^{r} \frac{\sigma_i}{\sigma_i^2+\lambda} \langle u_i, e\rangle v_i \quad \text{w/} \quad \boxed{A = \sum_{i=1}^{r} \sigma_i u_i v_i^T}$$

$$= \underbrace{\sum_{i=1}^{r} \frac{\sigma_i^2}{\sigma_i^2+\lambda} \langle v_i, x_0\rangle v_i}_{1} + \underbrace{\sum_{i=1}^{r} \frac{\sigma_i}{\sigma_i^2+\lambda} \langle u_i, e\rangle v_i}_{2}$$

Recall $x_0 = \sum_{i=1}^{n} \langle v_i, x_0\rangle v_i$; we will have the same kernel error as before

1 is a source of "regularization" error w/ magnitude squared $\sum_{i=1}^{r} \left(1 - \frac{\sigma_i^2}{\sigma_i^2+\lambda}\right)^2 |\langle v_i, x_0\rangle|^2 = \sum_{i=1}^{r} |\langle v_i, x_0\rangle|^2 \frac{\lambda^2}{(\sigma_i^2+\lambda)^2}$

$\hookrightarrow \left\| \sum_{i=1}^{r} \langle v_i, x_0\rangle v_i - \sum_{i=1}^{r} \frac{\sigma_i^2}{\sigma_i^2+\lambda} \langle v_i, x_0\rangle v_i \right\|_2^2$

If $\sigma_i^2 \ll \lambda$ the error is $|\langle v_i, x_0\rangle|^2$

If $\sigma_i^2 \gg \lambda$ the error is $\frac{\lambda^2}{\sigma_i^2} \ll 1$

2 is a source of noise error w/ magnitude squared $\sum_{i=1}^{r} \left(\frac{\sigma_i}{\sigma_i^2+\lambda}\right)^2 |\langle u_i, e\rangle|^2 \leq \frac{1}{4\lambda} \sum_{i=1}^{r} |\langle u_i, e\rangle|^2 \leq \frac{1}{4\lambda} \|e\|_2^2$

Proof of $\left(\frac{\sigma_i}{\sigma_i^2+\lambda}\right)^2 \le \frac{1}{4\lambda}$

Set $f(x) = \left(\frac{x}{x^2+\lambda}\right)^2 \quad x \in (0;+\infty) \qquad f(0) = \frac{1}{\lambda^2}$

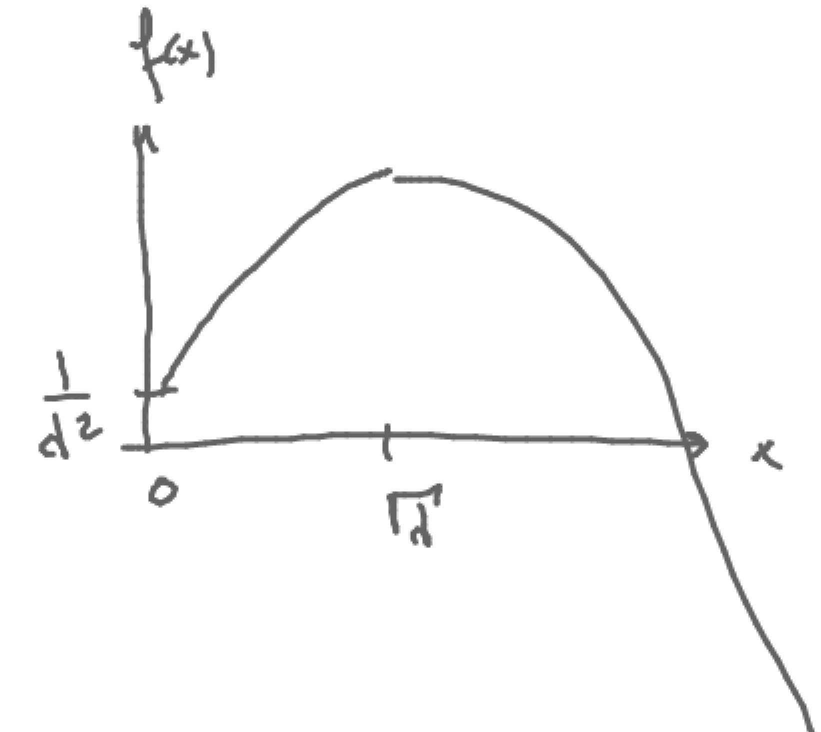

$$f'(x) = 2\,\frac{x}{x^2+\lambda} \times \frac{x^2+\lambda-2x^2}{(x^2+\lambda)^2} = \frac{2x}{(x^2+\lambda)^3} \times (\lambda - x^2)$$

$$f'(x) \geq 0 \iff \lambda - x^2 \geq 0$$
$$\iff x \leq \sqrt{\lambda}$$

Note that $f(\sqrt{\lambda}) = \left(\frac{\sqrt{\lambda}}{2\lambda}\right)^2 = \left(\frac{1}{2\sqrt{\lambda}}\right)^2 = \frac{1}{4\lambda}$

NUMERICAL METHODS

We have seen several solutions to systems of linear equations $\mathbf{A}\mathbf{x} = \mathbf{y}$ so far

- $\mathbf{A}$ full column rank: $\hat{\mathbf{x}} = (\mathbf{A}^\intercal\mathbf{A})^{-1}\mathbf{A}^\intercal\mathbf{y}$
- $\mathbf{A}$ full row rank: $\hat{\mathbf{x}} = \mathbf{A}^\intercal(\mathbf{A}\mathbf{A}^\intercal)^{-1}\mathbf{y}$
- Ridge regression: $\hat{\mathbf{x}} = (\mathbf{A}^\intercal\mathbf{A} + \delta\mathbf{I})^{-1}\mathbf{A}^\intercal\mathbf{y}$
- Kernel regression: $\hat{\mathbf{x}} = (\mathbf{K} + \delta\mathbf{I})^{-1}\mathbf{y}$
- Ridge regression in Hilbert space: $\hat{\mathbf{x}} = (\mathbf{A}^\intercal\mathbf{A} + \delta\mathbf{G})^{-1}\mathbf{A}^\intercal\mathbf{y}$

Extension: constrained least-squares

$$\min_{\mathbf{x}\in\mathbb{R}^n} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \text{ s.t. } \mathbf{x} = \mathbf{B}\boldsymbol{\alpha} \text{ for some } \boldsymbol{\alpha}$$

- The solution is $\hat{\mathbf{x}} = \mathbf{B}(\mathbf{B}^\intercal\mathbf{A}^\intercal\mathbf{A}\mathbf{B})^{-1}\mathbf{B}^\intercal\mathbf{A}^\intercal\mathbf{y}$

NUMERICAL METHODS

We have seen several solutions to systems of linear equations $\mathbf{A}\mathbf{x} = \mathbf{y}$ so far

- $\mathbf{A}$ full column rank: $\hat{\mathbf{x}} = (\mathbf{A}^\intercal \mathbf{A})^{-1} \mathbf{A}^\intercal \mathbf{y}$
- $\mathbf{A}$ full row rank: $\hat{\mathbf{x}} = \mathbf{A}^\intercal (\mathbf{A}\mathbf{A}^\intercal)^{-1} \mathbf{y}$
- Ridge regression: $\hat{\mathbf{x}} = (\mathbf{A}^\intercal \mathbf{A} + \delta \mathbf{I})^{-1} \mathbf{A}^\intercal \mathbf{y}$
- Kernel regression: $\hat{\mathbf{x}} = (\mathbf{K} + \delta \mathbf{I})^{-1} \mathbf{y}$
- Ridge regression in Hilbert space: $\hat{\mathbf{x}} = (\mathbf{A}^\intercal \mathbf{A} + \delta \mathbf{G})^{-1} \mathbf{A}^\intercal \mathbf{y}$

Extension: constrained least-squares

$$\min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \text{ s.t. } \mathbf{x} = \mathbf{B}\boldsymbol{\alpha} \text{ for some } \boldsymbol{\alpha}$$

- The solution is $\hat{\mathbf{x}} = \mathbf{B}(\mathbf{B}^\intercal \mathbf{A}^\intercal \mathbf{A}\mathbf{B})^{-1} \mathbf{B}^\intercal \mathbf{A}^\intercal \mathbf{y}$

All these problems involve a symmetric positive definite system of equations.

- Many methods to achieve this based on matrix factorization

EASY SYSTEMS

Diagonal system

- $\mathbf{A} \in \mathbb{R}^{n\times n}$ invertible and diagonal
- $O(n)$ complexity

$$Ax = y = \begin{pmatrix} a_1 & & 0 \\ & \ddots & \\ 0 & & a_n \end{pmatrix} x = \begin{pmatrix} a_1 x_1 \\ \vdots \\ a_n x_n \end{pmatrix} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}$$

The solution is $x_i = \frac{y_i}{a_i}$

EASY SYSTEMS

Diagonal system

- $\mathbf{A} \in \mathbb{R}^{n\times n}$ invertible and diagonal
- $O(n)$ complexity

Orthogonal system

- $\mathbf{A} \in \mathbb{R}^{n\times n}$ invertible and orthogonal
- $O(n^2)$ complexity

$y = Ax$ so that $A^T y = x$ b/c $A^T A = I$ $O(n^2)$ is the cost of matrix multiplication

EASY SYSTEMS

Diagonal system

- $\mathbf{A} \in \mathbb{R}^{n \times n}$ invertible and diagonal
- $O(n)$ complexity

Orthogonal system

- $\mathbf{A} \in \mathbb{R}^{n \times n}$ invertible and orthogonal
- $O(n^2)$ complexity

Lower triangular system

- $\mathbf{A} \in \mathbb{R}^{n \times n}$ invertible and lower diagonal
- $O(n^2)$ complexity

$$Ax = y = \begin{pmatrix} a_{11} & & 0 \\ & \ddots & \\ * & & a_{nn} \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} \quad \text{ie.} \quad \begin{array}{l} a_{11} x_1 = y_1 \;\checkmark \\ a_{21} x_1 + a_{22} x_2 = y_2 \;\checkmark \\ \vdots \\ a_{n1} x_1 + \dots + a_{nn} x_n = y_n \;\checkmark \end{array} \Bigg\downarrow \text{ substitution}$$

EASY SYSTEMS

Diagonal system

- $\mathbf{A} \in \mathbb{R}^{n\times n}$ invertible and diagonal
- $O(n)$ complexity

Orthogonal system

- $\mathbf{A} \in \mathbb{R}^{n\times n}$ invertible and orthogonal
- $O(n^2)$ complexity

Lower triangular system

- $\mathbf{A} \in \mathbb{R}^{n\times n}$ invertible and lower diagonal
- $O(n^2)$ complexity

General strategy: factorize $\mathbf{A}$ to recover some of the structures above

FACTORIZATIONS

LU factorization

$A \in \mathbb{R}^{n\times n}$ can be written $A = LU$ w/ $L \in \mathbb{R}^{n\times n}$ lower diagonal
$U \in \mathbb{R}^{n\times n}$ upper diagonal

If A invertible then U and L are as well so that $Ax = b \Leftrightarrow LUx = b \Rightarrow x = U^{-1}L^{-1}b$ $(*)$

We can solve by ① solve $Lw = b$ for w
② solve $Ux = w$ for x

triangular systems $O(n^2)$ (instead of $O(n^3)$ if inverting A d

FACTORIZATIONS

LU factorization

Cholesky factorization

QR decomposition

SVD and eigenvalue decompositions

COMPUTING EIGENVALUE DECOMPOSITIONS FOR SYMETRIC MATRICES

Many techniques: we shall only discuss one based on *power iterations*