

Welcome to Ma221! Lec 13, Fall 24

Stability of applying orthogonal transformations

Summary: Any Algorithm that just
multiplies by orthog. matrices is
back ward stable

Q' "nearly orthogonal" $\|Q' - Q\| = O(\epsilon)$
 $Q^T Q = I$

$$\begin{aligned} (*) \quad fl(Q' \cdot A) &= Q \cdot A + G \quad \|G\| = O(\epsilon) \|A\| \\ &= Q(A + Q^T G) = Q(A + G') \\ &\quad \|G'\|_2 = \|G\|_2 \end{aligned}$$

Multiplication by Q' is back ward stable
= exact transform of $A + G'$

$$\begin{aligned} \text{Eg: } Q' &= I - 2uu^T \quad \|u\|_2 = 1 \\ \text{in practice } &|\|u\|_2 - 1| = O(\epsilon) \end{aligned}$$

What if I multiply by many Q'_i ?

$$\begin{aligned} &fl(Q'_3(Q'_2(Q'_1 A))) \\ &= fl(Q'_3(Q'_2(Q_1(A + G_1)))) \\ &= fl(Q'_3(Q_2(Q_1(A + G_1) + G_2))) \end{aligned}$$

$$= \| Q_3 (Q_2 (Q_1 A + G_1) + G_2) + G_3 \|$$

$$= Q_3 Q_2 Q_1 A + \underbrace{Q_3 G_2 + Q_3 \cdot Q_2 \cdot G_1 + G_3}_G$$

$$\|G\| \leq \|G_2\| + \|G_1\| + \|G_3\| = O(\epsilon) \cdot \|A\|$$

$$= Q_3 Q_2 Q_1 (A + \underbrace{(Q_3 Q_2 Q_1)^T}_G G)$$

$$\|G'\| = \|G\| = O(\epsilon) \|A\|$$

\Rightarrow Perturbation for LS applies to solving using $A=QR$ via Householder

One more fast but sometimes unstable alg for QR
Cholesky QR:

$$\text{Factor } A^T A = L L^T = R^T R \quad (R=L^T)$$

$$\text{Form } Q = A R^{-1}$$

Can fail if A ill-conditioned

but used if A known to be well-conditioned

Dealing with (nearly) low rank matrices

Motivations: Real data often

low rank (nearly redundant)

① take precautions for inaccurate LS

② Use it to solve LS faster
or compress data
Use LS as example of compression,
other uses too

Solving a LS problem when A rank deficient

Thm: $A^{m \times n}$, $m \geq n$ rank $r < n$

$$A = U \Sigma V^T = \begin{matrix} m \\ \left[\begin{array}{ccc} U_1 & U_2 & U_3 \end{array} \right] \end{matrix} \begin{matrix} r & n-r \\ \left[\begin{array}{cc} \Sigma_1 & 0 \\ 0 & \Sigma_2 \\ 0 & 0 \end{array} \right] \end{matrix} \begin{matrix} n \\ \left[\begin{array}{cc} V_1 & V_2 \end{array} \right]^T \end{matrix}$$

Σ_1 : full rank

$\Sigma_2 = 0$ (later: tiny)

The set of vectors minimizing

$\|Ax - b\|_2$ is

$$\left\{ x = V_1 \Sigma_1^{-1} U_1^T b + \underbrace{V_2 y_2}_{\text{any vector in null_space}(A)}, \text{ any } y_2 \in \mathbb{R}^{n-r} \right\}$$

Unique x minimizing both $\|Ax - b\|_2$
and $\|x\|_2$ is gotten from $y_2 = 0$

$$\Rightarrow x = \underbrace{V_1 \Sigma_1^{-1} U_1^T}_A b$$

Def: $A^+ = V_1 \Sigma_1^{-1} U_1^T$ is Moore-Penrose
pseudoinverse of A (includes
full rank case $r = n$)

in practice, Σ_2 contains all singular values $<$ some user defined threshold

Proof: $\|Ax - b\|_2 = \|U \Sigma V^T x - b\|_2$

$= \|\Sigma V^T x - U^T b\|_2$ since $UU^T = I$

$= \|\Sigma y - U^T b\|_2$ where $y = V^T x$

$\|x\|_2 = \|y\|_2$ so ok to minimize either one

$= \left\| \begin{bmatrix} \Sigma_1 y_1 - U_1^T b \\ -U_2^T b \\ -U_3^T b \end{bmatrix} \right\|_2$ where $y = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix}$

minimized by $y_1 = \Sigma_1^{-1} U_1^T b$

$\|y_1\|_2^2 + \|y_2\|_2^2 = \|y\|_2^2$ minimized by $y_2 = 0$

$x = Vy = [V_1, V_2] \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = V_1 y_1 + V_2 y_2$

$= V_1 \Sigma_1^{-1} U_1^T b$ QED

Solving LS when A nearly rank deficient using truncated SVD

A sing $\Rightarrow \kappa(A) = \frac{\sigma_{\max}}{\sigma_{\min}} = \infty$

$\arg \min_x \| \begin{bmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \|_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$

$$0 < \epsilon < 1 \quad \operatorname{argmin}_x \left\| \begin{bmatrix} 1 & 0 \\ 0 & \epsilon \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \right\|_2 = \begin{bmatrix} 1 \\ 1/\epsilon \end{bmatrix}$$

what does a solution mean if it can change discontinuously?

Often A not known exactly just up to some tolerance $\|A - A'\|_2 < \epsilon$,
What to do?

Def: Truncated SVD
if $A = U \Sigma V^T$

$$A(\epsilon) = U \Sigma(\epsilon) V^T$$

$$\Sigma(\epsilon) = \operatorname{diag}(\sigma_1(\epsilon), \sigma_2(\epsilon), \dots, \sigma_n(\epsilon))$$

$$\sigma_i(\epsilon) = \begin{cases} \sigma_i & \text{if } \sigma_i \geq \epsilon \\ 0 & \text{if } \sigma_i < \epsilon \end{cases}$$

$A(\epsilon)$ = lowest rank matrix within distance ϵ of A

Solve LS using $A(\epsilon)$: reduces condition number from $\kappa(A) = \frac{\sigma_{\max}}{\sigma_{\min}}$

$$\text{to } \kappa(A_{\text{TOL}}) = \frac{\sigma_{\max}}{\epsilon}$$

Replacing A by "easier problem"
less sensitive, but larger residual;
process called regularization, several mechanisms

Lemma: $x_1 = \operatorname{argmin}_x \|A(\tau_0)x - b_1\|_2$

$x_2 = \operatorname{argmin}_x \|A(\tau_0)x - b_2\|_2$

Then $\|x_1 - x_2\|_2 \leq \frac{\|b_1 - b_2\|_2}{\tau_0}$

proof: $\|x_1 - x_2\|_2 = \|(A(\tau_0))^+ (b_1 - b_2)\|_2$

$$= \|\sqrt{\Sigma(\tau_0)^+} U^T (b_1 - b_2)\|_2$$

$$= \|\Sigma(\tau_0)^+ U^T (b_1 - b_2)\|_2$$

$$\leq \|\Sigma(\tau_0)^+\|_2 \cdot \|b_1 - b_2\|_2$$

$$\leq \frac{1}{\sigma_k} \|b_1 - b_2\|_2 \quad \sigma_k \text{ smallest } \geq \tau_0$$

$$\leq \frac{1}{\tau_0} \|b_1 - b_2\|_2$$

How does $A(\tau_0)$ depend on τ_0 ?

piecewise constant, changes when $\tau_0 = \sigma_k$

Other advantage of $A(\tau_0)$: compressible

only need $\underbrace{m \cdot k}_{U_1} + \underbrace{k}_{\Sigma(\tau_0)} + \underbrace{n \cdot k}_{V_1}$ words

instead of $m \cdot n$

can be \ll SVD

Solving (nearly) low rank LS

using Tikhonov regularization

or ridge regression

Replace $\operatorname{argmin}_x \|Ax - b\|_2^2$

with $\operatorname{argmin}_x \|Ax - b\|_2^2 + \lambda \|x\|_2^2$, $\lambda > 0$

λ "penalizes" very large x , user parameter

$$\begin{aligned} & \operatorname{argmin}_x \|Ax - b\|_2^2 + \lambda \|x\|_2^2 \\ &= \operatorname{argmin}_x \left\| \underbrace{\begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix}}_{\text{full rank if } \lambda > 0} x - \begin{bmatrix} b \\ 0 \end{bmatrix} \right\|_2^2 \end{aligned}$$

$$NE \Rightarrow x = \left(\begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix}^T \begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix} \right)^{-1} \begin{bmatrix} A \\ \sqrt{\lambda} I \end{bmatrix}^T \begin{bmatrix} b \\ 0 \end{bmatrix}$$

$$(*) = \underbrace{(A^T A + \lambda I)^{-1}}_{\text{pos def } \forall \lambda > 0} A^T b$$

How does λ change SVD

plug $A = U \Sigma V^T$ into (*)

$$\begin{aligned} x &= V \left(\Sigma (\Sigma^2 + \lambda I)^{-1} \right) U^T b \\ &= V \cdot \operatorname{diag} \left(\frac{\sigma_i}{\sigma_i^2 + \lambda} \right) U^T b \end{aligned}$$

if $\lambda = 0 \Rightarrow$ usual answer via SVD

$$\sigma_i \gg \lambda^{1/2} \Rightarrow \frac{\sigma_i}{\sigma_i^2 + \lambda} \sim \frac{1}{\sigma_i}, \text{ like SVD}$$

$$\sigma_i \leq \lambda^{1/2} \Rightarrow \frac{\sigma_i}{\sigma_i^2 + \lambda} \leq \frac{1}{2\lambda^{1/2}}$$

i.e. $\lambda^{1/2}$ and tol in $A(\text{tol})$
 play similar roles, but
 solution of ridge regression is
 continuous in λ

Solving low rank LS using QR

QR with column pivoting

Suppose we did $A = QR$ exactly on

A : $\text{rank}(A) = r < n$, what would R look like?

If leading r columns of A were

linearly independent (true for "almost all"

low rank A)

$$R = \begin{matrix} & \begin{matrix} r & n-r \end{matrix} \\ \begin{matrix} r \\ n-r \end{matrix} & \begin{bmatrix} R_{11} & R_{12} \\ 0 & 0 \end{bmatrix} \end{matrix} \quad \begin{matrix} R_{11} \text{ full rank} \\ R_{22} = 0 \end{matrix}$$

If A nearly low rank, hope that

$$\|R_{22}\| < \text{tol}, \text{ set } R_{22} = 0$$

Assuming this works, solve LS:

$$A = \begin{matrix} \begin{matrix} n & m-n \end{matrix} \\ \begin{matrix} m \\ m \end{matrix} \end{matrix} \hat{Q} \hat{R} = \begin{matrix} \begin{matrix} n & m-n \end{matrix} \\ \begin{matrix} m \\ m \end{matrix} \end{matrix} \begin{bmatrix} \hat{Q} & \hat{Q}' \\ 0 \end{bmatrix} \begin{bmatrix} R \\ 0 \end{bmatrix}$$

$$\begin{matrix} \begin{matrix} n & m-n \end{matrix} \\ \begin{matrix} m \\ m \end{matrix} \end{matrix} [\hat{Q}, \hat{Q}'] = \begin{bmatrix} \hat{Q}_1 & \hat{Q}_2 & \hat{Q}' \end{bmatrix}_m$$

$$\begin{aligned}
& \operatorname{argmin}_x \|Ax - b\|_2 \\
&= \operatorname{argmin}_x \left\| \begin{bmatrix} Q_1 & Q_2 & Q' \end{bmatrix} \begin{bmatrix} R \\ 0 \end{bmatrix} x - b \right\|_2 \\
&= \operatorname{argmin}_x \left\| \begin{array}{c} r \\ n-r \\ m-n \end{array} \begin{bmatrix} R_{11} & R_{12} \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{array}{c} r \\ n-r \end{array} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} - \begin{bmatrix} Q_1^T b \\ Q_2^T b \\ Q'^T b \end{bmatrix} \right\|_2 \\
&= \operatorname{argmin}_x \left\| \begin{bmatrix} R_{11} x_1 + R_{12} x_2 - Q_1^T b \\ -Q_2^T b \\ -Q'^T b \end{bmatrix} \right\|_2
\end{aligned}$$

solution: $x_1 = R_{11}^{-1} Q_1^T b - R_{11}^{-1} R_{12} x_2$
for any x_2

How to pick x_2 to minimize $\|x\|_2$?

Ex: $A = \begin{bmatrix} e & 1 \\ 0 & 0 \end{bmatrix}$ $0 < e \ll 1$ $R_{11} = e, R_{12} = 1$

$$x = \begin{bmatrix} (b_1 - x_2)/e \\ x_2 \end{bmatrix} \quad e \text{ tiny} \Rightarrow \text{very sensitive to changes in } b_1, x_2$$

permute columns of A :

$$AP = \begin{bmatrix} 1 & e \\ 0 & 0 \end{bmatrix} \Rightarrow x = \begin{bmatrix} b_1 - e x_2 \\ x_2 \end{bmatrix}$$

insensitive to changes in x_2, b_1

What would a perfect R factor look like?

Compare $\begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix}$ to $\begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix}$

Def: Rank Revealing QR factorization

(RRQR for short) is $A \cdot P = QR$

P permutation chosen so that

① R_{22} is "small", ideally $\|R_{22}\|_2 = O(\sigma_{k+1})$

R_{22} "contains smallest $n-k$ sing. values"

② R_{11} is "large", ideally $\sigma_{\min}(R_{11})$
not much smaller than σ_k

If in addition

③ $\|R_{11}^{-1} R_{12}\|_2$ not "too large"
then $AP = QR$ called "strong RRQR"

Thm (informal) if (1), (2), (3) hold

$$\sigma_i(A) \geq \sigma_i(R_{11}) \geq \frac{\sigma_i(A)}{\sqrt{1 + \|R_{11}^{-1} R_{12}\|_2^2}}$$

for $1 \leq i \leq k$

$$\sigma_i(A) \leq \sigma_{\max}(R_{22}) \sqrt{1 + \|R_{11}^{-1} R_{12}\|_2^2}$$

$i = k+1, \dots, n$

Leading k columns of $A \cdot P$ contain most information about $\text{range}(A)$

$$AP = \begin{bmatrix} A_1 & A_2 \end{bmatrix} = QR = \begin{bmatrix} Q_1 & Q_2 \end{bmatrix} \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix}$$

$$\approx \begin{bmatrix} Q_1 & Q_2 \end{bmatrix} \cdot \begin{bmatrix} R_{11} & R_{12} \\ 0 & 0 \end{bmatrix}$$

$$= Q_1 \begin{bmatrix} R_{11} & R_{12} \end{bmatrix}$$

$$= Q_1 R_{11} \begin{bmatrix} I & R_{11}^{-1} R_{12} \end{bmatrix}$$

$$= A_1 \begin{bmatrix} I & R_{11}^{-1} R_{12} \end{bmatrix}$$

How to compute P ?

(matlab demo, see notes)

Algorithm: QR with column pivoting
QRCP for short

Analogous to partial pivoting

often works, like GEPP, can be off by a factor of 2^n

first step: pick largest column of A ,
move to front

take one step of QR

among remaining columns, pick one with largest component orthogonal to first column

repeat

for $i = 1$ to $\min(m-1, n)$ or R_{22} small enough



"truncated QR"

choose largest column (in norm) in trailing matrix

$$j = \underset{j \geq i}{\operatorname{argmax}} \|A(i:m, j)\|_2$$

if $i \neq j$, swap cols i and j

multiply $A(i:m, i:n)$ by Householder matrix to zero out $A(i+1:m, i)$



Need to compute column norms carefully to be fast + accurate

in LAPACK: $g \leftarrow qp^3$, $q \leftarrow qp^3 r k$

Matlab: $[Q,R,P] = qr(A)$