DN2222
Applied Numerical Methods
- part 2:
Numerical Linear Algebra

Lecture 4
Least Squares Problem
&
Singular Value Decomposition

2011-11-08

Linear systems: $Ax = b$

- If A is $n \times n$ and non-singular then $x = A^{-1}b$ uniquely.
- If A is $m \times n$ and $m < n$ the problem is *underdetermined* (and thus usually have infinitely many solutions).
- If A is $m \times n$ and $m > n$ the problem is *overdetermined* and then normally have no solution. This is the topic for today.
- We will try to find the "best approximate solution" to the overdetermined system.

Example 3.1 Ruhe p21

Given m pairs of data points $(t_1, y_1), \dots (t_m, y_m)$ from a sample of radioactive decay. The intensity is modeled by

$$y = \sum_{j=1}^{n} \alpha_j \mathrm{e}^{-\lambda_j t}, \qquad \alpha_j, \lambda_j \geq 0$$

The residual is then

$$r = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix} - \begin{pmatrix} \mathrm{e}^{-\lambda_1 t_1} & \dots & \mathrm{e}^{-\lambda_n t_1} \\ \mathrm{e}^{-\lambda_1 t_2} & \dots & \mathrm{e}^{-\lambda_n t_2} \\ \vdots & \vdots & \vdots \\ \mathrm{e}^{-\lambda_1 t_m} & \dots & \mathrm{e}^{-\lambda_n t_m} \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_n \end{pmatrix} = y - A(\lambda, t)x(\alpha)$$

- $r = y - A(\lambda, t)x(\alpha)$
- The *residual*, r, is the difference between the *observation* vector y and the product of the *design* or *system* matrix A and the *parameter* vector x.
- The task at hand is to compute parameters α and λ such that the residual is minimized in an appropriate norm.
- If the λ are known and only the α need to be determined the system is linear, otherwise non-linear.

Choice of norm

- If the errors are independent we choose the Euclidian norm $||r||_2 = (r^T r)^{1/2} = (\sum_{i=1}^{m} r_i^2)^{1/2}$. It is the most common and we talk about the *least squares method.*

- If we want to minimze the maximal residual we use the infimum norm $||r||_\infty = \max_i |r_i|$. This is typical for polynomial interpolation.
- Another choice is the 1-norm $||r||_1 = \sum_i |r_i|$. This is typically used when avoidance of outliers is important.

Surveyers work

- The least squares method was invented by Gauss trying to improve accuracy in the German surveyers and astronomers measurement.
- In 1974-78 the US National Geodetic Survey updated its database in the same manner - solving the biggest least squares problem ever: about 6 million equations and 400000 unknowns.

Solutions

- **Normal equations**. Fast but not very accurate. Adequate when the condition number is small.
- **QR decomposition**. Twice the amount of work but more accuarate. The standard method.
- **SVD**. Even more work but works even if A is not full rank.

Normal equations

- To derive the normal equations we need to minimize $||r||_2^2 = r^T r = (b - Ax)^T (b - Ax)$
- Leads to $A^T Ax = A^T b$ or $x = (A^T A)^{-1} A^T b$
- Proof: Let $x' = x + e$ then

$$||Ax' - b||_2^2 = (Ax' - b)^T (Ax' - b) = (Ae + Ax - b)^T (Ae + Ax - b)$$

$$= (Ae)^T (Ae) + (Ax - b)^T (Ax - b) + 2(Ae)^T (Ax - b)$$

$$= ||Ae||_2^2 + ||Ax - b||_2^2 + 2e^T (A^T Ax - A^T b) = ||Ae||_2^2 + ||Ax - b||_2^2$$

- This is equivalent to the Pythagorean theorem. The solution is optimized whrn the residual is orthogonal to the space spanned by the columns of A.
- Since $A^T A$ is symmetric and positive definite we can use Cholesky factorization. The cost for Cholesky is $\frac{1}{3}n^3$ and the cost for obtaining $A^T A$ from A is $n^2 m$.
- Since $m > n$ forming $A^T A$ dominates the cost!

QR Decomposition

- **Thm 3.1** (Dp107) Let A be $m \times n$ with $m > n$ and rank(A)=n. Then there exists a unique $m \times n$ orthogonal matrix Q ($Q^T Q = I_n$) and a unique $n \times n$ upper triangular matrix R with positive diagonal elements $r_{ii} > 0$ such that $A = QR$.

- First proof uses Gram-Schmidt orthogonalization process. If apply GS to the columns of $A = [a_1, a_2, \dots a_n$ one gets a sequence of orthonormal vectors q_i obtained from a linear combination of a_1 to a_i.

- Unfortunately GS is numerically unstable in floating point arithmetic when the columns of A are nearly dependent.

- Modified Gram-Schmidt (MGS) is more stable but could still end up with a Q which is far from orthogonal.

$$\begin{aligned} x &= (A^T A)^{-1} A^T b \\ &= (R^T Q^T Q R)^{-1} R^T Q^T b \\ &= (R^T R)^{-1} R^T Q^T b \\ &= R^{-1} Q^T b \end{aligned}$$

- The cost for QR-decomposition is about $2n^2 m - \frac{2}{3} n^3$, about twice the cost of normal equations if $m >> n$ and about the same if $m = n$.

Singular Value Decomposition

- SVD is used for many things, not only least squares.

- Thm 3.2 (Dp109) Let A be an arbitrary $m \times n$ matrix with $m \geq n$. Then we can write $A = U \Sigma V^T$, where U is an $m \times n$ such that $U^T U = I$, V is an $n \times n$ such that $V^T V = I$, and $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$, where $\sigma_1 \geq \sigma_2 \cdots \geq \sigma_n \geq 0$. The columns $u_1, \dots, u_n$ of U are called left singular vectors. The columns $v_1, \dots, v_n$ of V are called left singular vectors. The σ_i are called singular values.

- Proof of Thm 3.2: We assume that SVD exists for an $(m-1) \times (n-1)$ matrix and then prove it for an $n \times m$. SVD has a large number of properties:

- If A is symmetric, then $\sigma_i = |\lambda_i|$ and $v_i = sign(\lambda_i) u_i$.

- The eigenvalues of $A^T A$ are σ_i^2. The right singular vectors are the corresponding orthogonal eigenvectors.

- The eigenvalues of AA^T are σ_i^2 and $m - n$ zeroes. The left singular vectors are the corresponding orthogonal eigenvectors.

- If A has full rank, the least squares solution of $Ax = b$ is $x = V \Sigma^{-1} U^T b$.

- $||A||_2 = |\sigma_1|$.

- If A is square and non-singular then $||A^{-1}||_2^{-1} = \sigma_n$ and $||A||_2 ||A^{-1}||_2 = \frac{\sigma_1}{\sigma_n}$.

- Suppose that A is $m \times n$ and has rank n with $m > n$, then $A^+ = (A^T A)^{-1} A^T = R^{-1} Q^T = V \Sigma^{-1} U^T$ is called the (Morse-Penrose) pseudo-inverse of A.

- If $m < n$ then $A^+ = A^T(AA^T)^{-1}$.

Rank Deficient least Squares Problems

- What happens if A is rank deficient (or nearly)?
- This occurs often, like signals in noisy data (Lab3), digital image restoration or compression, etc.
- Rank deficient problems are very ill-conditioned.
- Making an ill-conditioned problem well-conditioned by imposing extra conditions on the solution is called *regularization.*
- If A is rank deficient the least squares solution is not unique.
- Prop 3.1 (Dp125) Let A be an $m \times n$ matrix with rank(A)=$r < n$. Then there is an $n - r$ dimensional set of vectors that all minimizes $||(||Ax - b)$.
- Proof: Let z be such that $Az = 0$ then if x minimizes $||Ax - b||$ then so does $x + z$.
- If, due to round-off, some σ_i has a small value rather than zero, Then the unique solution is likely to be very large.
- Thus: If A is nearly rank deficient (σ_{min} is small) the solution x is ill-conditioned and possibly very large.

Prop 3.3 (Dp126) When A is exactly singular, the x tht minimizes $||Ax - b||_2$ can be characterized as follows: Let $A = U\Sigma V^T$ have rank $r < n$. Then write

$$A = [U1, U2] \begin{bmatrix} \Sigma_1 & 0 \\ 0 & 0 \end{bmatrix} [V1, V2]^T = U_1\Sigma_1 V_1^T$$

where Σ_1 is $r \times r$ and non-singular and U_1 and V_1 have r columns. Let $\sigma = \sigma_{min}(\Sigma_1)$. Then:

- all solutions x can be written as $x = V_1\Sigma_1^{-1}U_1^T b + V_2 z$, z being any vector.

Proof: Choose $\tilde{U}$ such that $W = [U_1, U_2, \tilde{U}]$ is an orthogonal matrix.

$$||Ax - b||_2^2 = ||W^T(Ax - b)||_2^2 = || \begin{bmatrix} U_1^T \\ U_2^T \\ \tilde{U}^T \end{bmatrix} (U_1\Sigma_1 V_1^T x - b)||_2^2$$
$$= ||\Sigma_1 V_1^T x - U_1^T b||_2^2 + ||U_2^T b||_2^2 + ||\tilde{U}^T b||_2^2$$

Thus, x is multiplied with V_1, anything with V_2 will add zero.

- the solution x has minimal norm $||x||_2$ precisely when $z = 0$, in which case $x = V_1\Sigma_1^{-1}U_1^T b$ and $||x||_2 \leq ||b||_2/\sigma$.

Proof: Since V_1 and V_2 are mutually orthogonal by Pythagoras

$$||x||_2^2 = ||V_1\Sigma_1^{-1}U_1^T b||_2^2 + ||V_2 z||_2^2$$

which is minimized when $z = 0$.

- Changing b into $b + \delta b$ can change the minimal norm solution x by at most $||\delta b||_2/\sigma$

  Proof:

  $$||V_1 \Sigma_1^{-1} U_1^T \delta b||_2 \leq ||\Sigma_1^{-1}||_2 ||\delta b||_2 = ||\delta b||_2/\sigma$$

- The norm and condition number of the unique minimal norm solution x depends on the smalleest non-zero singular value of A.
- This is the key to a practical algorithm!

Pseudoinverse for Rank Deficient matrix

- Let $A = U\Sigma V^T = U_1 \Sigma_1 V_1^T$ Then $A^+ = V_1 \Sigma_1^{-1} U_1^T$ or $A^+ = V^T \Sigma^+ U$ where $\Sigma^+ = \begin{bmatrix} \Sigma_1 & 0 \\ 0 & 0 \end{bmatrix}^+ = \begin{bmatrix} \Sigma_1^{-1} & 0 \\ 0 & 0 \end{bmatrix}$
- So the least squares solution is always $x = A^+ b$. When A is rank deficient, x has minimum norm.
- So we need to know the rank of A and the smallest singular value.

Example: Demmel p128

$A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$ has smallest nonzero eigenvalue 1. With $b = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ we get least square solution $x = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ with condition number $1/\sigma = 1$.

But if we have $A = \begin{pmatrix} 1 & 0 \\ 0 & \varepsilon \end{pmatrix}$ we have smallest nonzero eigenvalue ε and $x = \begin{pmatrix} 1 \\ 1/\varepsilon \end{pmatrix}$ and condition number $1/\varepsilon$.

- The practical solution is to treat all σ_i smaller than a tolerance (normally $O(\varepsilon) \cdot ||A||_2$) as zero.
- This is called *truncated SVD*
- A similar idea can be used in QR-decomposition, but it is less reliable.

Some review questions:

- **Q45.** What is the range $R(A)$ of a matrix A? How do you find a basis for it by means of SVD?
- **Q50.** What is meant by a rank deficient matrix?
- **Q51.** How can we determine $A^{(k)}$, the matrix of rank k closest to a given matrix A using SVD?
- **Q53.** What are the advantages and disadvantages of replacing the matrix A by a lower rank approximation $A^{(k)}$ when solving a least squares problem?