CME 302: NUMERICAL LINEAR ALGEBRA
FALL 2005/06
LECTURE 12

GENE H. GOLUB

1. Least Squares with Linear Constraints

Suppose that we wish to fit data as in the least squares problem, except that we are using different functions to fit the data on different subintervals. A common example is the process of fitting data using cubic splines, with a different cubic polynomial approximating data on each subinterval.

Typically, it is desired that the functions assigned to each piece form a function that is continuous on the entire interval within which the data lies. This requires that *constraints* be imposed on the functions themselves. It is also not uncommon to require that the function assembled from these pieces also has a continuous first or even second derivative, resulting in additional constraints. The result is a *least squares problem with linear constraints*, as the constraints are applied to coefficients of predetermined functions chosen as a basis for some function space, such as the space of polynomials of a given degree.

The general form of a least squares problem with linear constraints is as follows: we wish to find an n-vector $\mathbf{x}$ that minimizes $\|A\mathbf{x}-\mathbf{b}\|_2$, subject to the constraint $C^\top\mathbf{x}=\mathbf{d}$, where C is a known $n\times p$ matrix and $\mathbf{d}$ is a known p-vector.

This problem is usually solved using *Lagrange multipliers*. We define

$$f(\mathbf{x};\boldsymbol{\lambda}) = \|\mathbf{b}-A\mathbf{x}\|_2^2 + 2\boldsymbol{\lambda}^\top C^\top\mathbf{x}.$$

Then

$$\nabla f = 2(A^\top A\mathbf{x} - A^\top\mathbf{b} + C\boldsymbol{\lambda}).$$

To minimize f, we can solve the system

$$\begin{bmatrix} A^\top A & C \\ C^\top & 0 \end{bmatrix}\begin{bmatrix} \hat{\mathbf{x}} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} A^\top\mathbf{b} \\ \mathbf{d} \end{bmatrix}.$$

From $A^\top A\mathbf{x} = A^\top\mathbf{b} - C\boldsymbol{\lambda}$, we see that we can first compute $\mathbf{x} = \hat{\mathbf{x}} - (A^\top A)^{-1}C\boldsymbol{\lambda}$ where $\hat{\mathbf{x}}$ is the solution to the unconstrained least squares problem. Then, from the equation $C^\top\mathbf{x}=\mathbf{d}$ we obtain the equation $C^\top(A^\top A)^{-1}C\boldsymbol{\lambda} = C^\top\hat{\mathbf{x}} - \mathbf{d}$ which we can now solve for $\boldsymbol{\lambda}$. The algorithm proceeds as follows:

(1) Solve the unconstrained least squares problem $A\mathbf{x}=\mathbf{b}$ for $\hat{\mathbf{x}}$.
(2) Compute $A=QR$.
(3) Form $W=(R^\top)^{-1}C$.
(4) Compute $W=PU$, the QR factorization of W.

Date: November 30, 2005, version 1.0.

Notes originally due to James Lambers. Edited by Lek-Heng Lim.

(5) Solve $U^\top U\boldsymbol{\lambda} = \boldsymbol{\eta} = C^\top\hat{\mathbf{x}} - \mathbf{d}$ for $\boldsymbol{\lambda}$. Note that

$$\begin{aligned} U^\top U &= (P^\top W)^\top (P^\top W) \\ &= W^\top P P^\top W \\ &= C^\top R^{-1}(R^\top)^{-1} C \\ &= C^\top (R^\top R)^{-1} C \\ &= C^\top (R^\top Q^\top Q R)^{-1} C \\ &= C^\top (A^\top A)^{-1} C \end{aligned}$$

(6) Set $\mathbf{x} = \hat{\mathbf{x}} - (A^\top A)^{-1} C\boldsymbol{\lambda}$.

This method is not the most practical since it has more unknowns than the unconstrained least squares problem, which is odd because the constraints should have the effect of eliminating unknowns, not adding them. We now describe an alternate approach.

Suppose that we compute the QR factorization of C to obtain

$$Q^\top C = \begin{bmatrix} R \\ 0 \end{bmatrix}$$

where R is a $p \times p$ upper triangular matrix. Then the constraint $C^\top \mathbf{x} = \mathbf{d}$ takes the form

$$R^\top \mathbf{u} = \mathbf{d}, \quad Q^\top \mathbf{x} = \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix}.$$

Then

$$\begin{aligned} \|\mathbf{b} - A\mathbf{x}\|_2 &= \|\mathbf{b} - AQQ^\top \mathbf{x}\| \\ &= \left\| \mathbf{b} - \tilde{A} \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \right\|_2, \quad \tilde{A} = AQ \\ &= \left\| \mathbf{b} - \begin{bmatrix} \tilde{A}_1 & \tilde{A}_2 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix} \right\|_2 \\ &= \|\mathbf{b} - \tilde{A}_1 \mathbf{u} - \tilde{A}_2 \mathbf{v}\|_2 \end{aligned}$$

Thus we can obtain $\mathbf{x}$ by the following procedure:

(1) Compute the QR factorization of C
(2) Compute $\tilde{A} = AQ$
(3) Solve $R^\top \mathbf{u} = \mathbf{d}$
(4) Solve the new least squares problem of minimizing $\|(\mathbf{b} - \tilde{A}_1\mathbf{u}) - \tilde{A}_2\mathbf{v}\|_2$
(5) Compute

$$\mathbf{x} = Q \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix}.$$

This approach has the advantage that there are fewer unknowns in each system that needs to be solved, and also that $\kappa(\tilde{A}_2) \leq \kappa(A)$. The drawback is that sparsity can be destroyed.

2. Least Squares with Quadratic Constraints

We wish to solve the problem

$$\|\mathbf{b} - A\mathbf{x}\|_2 = \min, \quad \|\mathbf{x}\|_2 = \alpha, \quad \alpha \leq \|A^+\mathbf{b}\|_2.$$

This problem is known as *least squares with quadratic constraints.* To solve this problem, we define

$$\varphi(\mathbf{x}; \mu) = \|\mathbf{b} - A\mathbf{x}\|_2^2 + \mu(\|\mathbf{x}\|^2 - \alpha^2)$$

and seek to minimize φ. From

$$\nabla\varphi = 2A^\top \mathbf{b} - 2A^\top A\mathbf{x} + 2\mu\mathbf{x}$$

we obtain the system

$$(A^\top A + \mu I)\mathbf{x} = A^\top \mathbf{b}.$$

If we denote the eigenvalues of $A^\top A$ by

$$\lambda_i(A^\top A) = \lambda_1, \dots, \lambda_n, \quad \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq 0$$

then

$$\lambda_i(A^\top A + \mu I) = \lambda_1 + \mu, \cdots, \lambda_n + \mu.$$

If $\mu \geq 0$, then $\kappa(A^\top A + \mu I) \leq \kappa(A^\top A)$, because

$$\frac{\lambda_1 + \mu}{\lambda_n + \mu} \leq \frac{\lambda_1}{\lambda_n},$$

so $A^\top A + \mu I$ is better conditioned.

Solving the least squares problem with quadratic constraints arises in many literatures, including

(1) Statistics: Ridge Regression
(2) Regularization: Tichonov
(3) Generalized cross-validation (GCV)

To solve this problem, we see that we need to compute

$$\mathbf{x} = (A^\top A + \mu I)^{-1} A^\top \mathbf{b}$$

where

$$\mathbf{x}^\top \mathbf{x} = \mathbf{b}^\top A (A^\top A + \mu I)^{-2} A^\top \mathbf{b} = \alpha^2.$$

If $A = U\Sigma V^\top$ is the SVD of A, then we have

$$\begin{aligned}
\alpha^2 &= \mathbf{b}^\top U\Sigma V^\top (V\Sigma^\top \Sigma V^\top + \mu I)^{-2} V\Sigma^\top U^\top \mathbf{b} \\
&= \mathbf{c}^\top \Sigma (\Sigma^\top \Sigma + \mu I)^{-2} \Sigma^\top \mathbf{c}, \quad U^\top \mathbf{b} = \mathbf{c} \\
&= \sum_{i=1}^{r} \frac{c_i^2 \sigma_i^2}{(\sigma_i^2 + \mu)^2} \\
&= \chi(\mu)
\end{aligned}$$

The function $\chi(\mu)$ has poles at $-\sigma_i^2$ for $i = 1, \dots, n$. Furthermore, $\lim_{\mu \to \infty} \chi(\mu) = 0$.

We now have the following procedure for solving this problem, given A, $\mathbf{b}$, and α^2:

(1) Compute the SVD of A to obtain $A = U\Sigma V^\top$.
(2) Compute $\mathbf{c} = U^\top \mathbf{b}$.
(3) Solve $\chi(\mu^*) = \alpha^2$ where $\mu^* \geq 0$. Don't use Newton's method on this equation directly; solving $1/\chi(\mu) = 1/\alpha^2$ is much better.
(4) Use the SVD to compute

$$\mathbf{x} = (A^\top A + \mu I)^{-1} A^\top \mathbf{b} = V(\Sigma^\top \Sigma + \mu I)^{-1} \Sigma^\top U^\top \mathbf{b}.$$

3. Applications of the SVD

3.1. **Minimum-norm least squares solution.** One of the most well-known applications of the SVD is that it can be used to obtain the solution to the problem

$$\|\mathbf{b} - A\mathbf{x}\|_2 = \min, \quad \|\mathbf{x}\|_2 = \min.$$

The solution is

$$\hat{\mathbf{x}} = A^+ \mathbf{b} = V\Sigma^+ U^\top \mathbf{b}$$

where A^+ is the *pseudo-inverse* of A.

3.2. **Closest Orthogonal Matrix.** Let $\mathcal{Q}_n$ be the set of all $n \times n$ orthogonal matrices. Given an $n \times n$ matrix A, we wish to find the matrix Q that satisfies

$$\|A - Q\|_F = \min\,, \quad Q \in \mathcal{Q}_n, \quad \sigma_i(Q) = 1.$$

Given $A = U\Sigma V^\top$, if we compute $\hat{Q} = UIV^\top$, then

$$\begin{aligned}\|A - \hat{Q}\|_F^2 &= \|U(\Sigma - I)V^\top\|_F^2 \\ &= \|\Sigma - I\|_F^2 \\ &= (\sigma_1 - 1)^2 + \cdots + (\sigma_n - 1)^2\end{aligned}$$

It can be shown that this is in fact the minimum.

A more general problem is to find $Q \in \mathcal{Q}_n$ such that

$$\|A - BQ\|_F = \min$$

for given matrices A and B. The solution is

$$\hat{Q} = UV^\top, \quad B^\top A = U\Sigma V^\top.$$

3.3. **Low-Rank Approximations.** Let $\mathcal{M}_{m,n}^{(r)}$ be the set of all $m \times n$ matrices of rank r, and let $A \in \mathcal{M}_{m,n}^{(r)}$. We wish to find $B \in \mathcal{M}_{m,n}^{(k)}$, where $k < r$, such that $\|A - B\|_F = \min$.

To solve this problem, let $A = U\Sigma V^\top$ be the SVD of A, and let $\hat{B} = U\Omega_k V^\top$ where

$$\Omega_k = \begin{bmatrix} \sigma_1 & & & & & \\ & \ddots & & & & \\ & & \sigma_k & & & \\ & & & 0 & & \\ & & & & \ddots & \\ & & & & & 0 \end{bmatrix}.$$

Then

$$\begin{aligned}\|A - \hat{B}\|_F^2 &= \|U(\Sigma - \Omega_k)V^\top\|_F^2 \\ &= \|\Sigma - \Omega_k\|_F^2 \\ &= \sigma_{k+1}^2 + \cdots + \sigma_r^2.\end{aligned}$$

We now consider a variation of this problem. Suppose that B is a perturbation of A such that $A = B + E$, where $\|E\|_F^2 \le \epsilon^2$. We wish to find $\hat{B}$ such that $\|A - \hat{B}\|_F^2 \le \epsilon^2$, where the rank of $\hat{B}$ is minimized. We know that if $B_k = U\Omega_k V^\top$ then

$$\|A - B_K\|_F^2 = \sigma_{k+1}^2 + \cdots + \sigma_r^2.$$

It follows that $\hat{B} = B_k$ is the solution if

$$\sigma_{k+1} + \cdots + \sigma_r^2 \le \epsilon^2, \quad \sigma_k^2 + \cdots + \sigma_r^2 > \epsilon^2.$$

Note that

$$\|A^+ - \hat{B}^+\|_F^2 = \left(\frac{1}{\sigma_{k+1}^2} + \cdots + \frac{1}{\sigma_r^2}\right).$$

4. Total Least Squares

In the ordinary least squares problem, we are solving

$$A\mathbf{x} = \mathbf{b} + \mathbf{r}, \quad \|\mathbf{r}\|_2 = \min .$$

In the *total least squares* problem, we wish to solve

$$(A + E)\mathbf{x} = \mathbf{b} + \mathbf{r}, \quad \|E\|_F^2 + \lambda^2 \|\mathbf{r}\|_2^2 = \min .$$

From $A\mathbf{x} - \mathbf{b} + E\mathbf{x} - \mathbf{r}$ we obtain the system

$$\begin{bmatrix} A & \mathbf{b} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ -1 \end{bmatrix} + \begin{bmatrix} E & \mathbf{r} \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ -1 \end{bmatrix} = \mathbf{0},$$

or

$$(C + F)\mathbf{z} = \mathbf{0}.$$

We need the matrix $C + F$ to have rank $\leq n + 1$, and we want to minimize $\|F\|$.

To solve this problem, we compute the SVD of $C = \begin{bmatrix} A & \mathbf{b} \end{bmatrix} = U\Sigma V^\top$. Let $\hat{C} = U\Omega_n V^\top$. Then, if $\mathbf{v}_i$ is the ith column of V, we have

$$\hat{C}\mathbf{v}_{n+1} = U\Omega_n V^\top \mathbf{v}_{n+1} = \mathbf{0}.$$

Our solution is

$$\begin{bmatrix} \hat{\mathbf{x}} \\ -1 \end{bmatrix} = -\frac{1}{v_{n+1,n+1}} \mathbf{v}_{n+1}$$

provided that $v_{n+1,n+1} \neq 0$.

Now, suppose that only some of the data is contaminated, i.e. $E = \begin{bmatrix} 0 & E_1 \end{bmatrix}$ where the first p columns of E are zero. Then, in solving $(C + F)\mathbf{z} = \mathbf{0}$, we use Householder transformations to compute $Q^\top(C+F)$ where the first p columns are zero below the diagonal. Since $\|F\|_F = \|Q^\top F\|_F$, we then have a block upper triangular system

$$\begin{bmatrix} R_{11} & R_{12} + F_{12} \\ 0 & R_{22} + F_{22} \end{bmatrix} \mathbf{z} = \mathbf{0}, \quad \mathbf{z} = \begin{bmatrix} \mathbf{u} \\ \mathbf{v} \end{bmatrix}.$$

We can find the total least squares solution of

$$(R_{22} + F_{22})\mathbf{v} = \mathbf{0},$$

and then set $F_{12} = 0$ and solve

$$R_{11}\mathbf{u} + R_{12}\mathbf{v} = \mathbf{0}.$$

Department of Computer Science, Gates Building 2B, Room 280, Stanford, CA 94305-9025
E-mail address: golub@stanford.edu