CME 302: NUMERICAL LINEAR ALGEBRA
FALL 2005/06
LECTURE 16

GENE H. GOLUB

1. Method of Steepest Descent

An alternative approach is to consider the iteration

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{r}^{(k)}$$

where α_k varies from iteration to iteration. It follows that

$$\mathbf{r}^{(k+1)} = \mathbf{b} - A\mathbf{x}^{(k+1)} = \mathbf{b} - A\mathbf{x}^{(k)} - \alpha_k A\mathbf{r}^{(k)} = \mathbf{r}^{(k)} - \alpha_k A\mathbf{r}^{(k)}.$$

We wish to choose α_k so that $\mathbf{r}^{(k+1)\top} A^{-1}\mathbf{r}^{(k+1)}$ is minimized. Now

$$\begin{aligned}\mathbf{r}^{(k+1)\top} A^{-1}\mathbf{r}^{(k+1)} &= (\mathbf{r}^{(k)\top} - \alpha_k \mathbf{r}^{(k)\top} A)A^{-1}(\mathbf{r}^{(k)} - \alpha_k A\mathbf{r}^{(k)} \\ &= \mathbf{r}^{(k)\top} A^{-1}\mathbf{r}^{(k)} - 2\alpha_k \mathbf{r}^{(k)\top}\mathbf{r}^{(k)} + \alpha_k^2 \mathbf{r}^{(k)\top} A\mathbf{r}^{(k)}. \qquad (1.1)\end{aligned}$$

To find the minimium, we differentiate with respect to α_k and obtain

$$\frac{d}{d\alpha_k}\mathbf{r}^{(k)\top} A^{-1}\mathbf{r}^{(k+1)} = -2\mathbf{r}^{(k)\top}\mathbf{r}^{(k)} + 2\alpha_k \mathbf{r}^{(k)\top} A\mathbf{r}^{(k)}$$

which yields

$$\hat{\alpha}_k = \frac{\mathbf{r}^{(k)\top}\mathbf{r}^{(k)}}{\mathbf{r}^{(k)\top} A\mathbf{r}^{(k)}}$$

which is well-defined since A is symmetric positive definite. This method is known as the *method of steepest descent*.

Note that

$$0 < \lambda_{\min}(A) \le \frac{\mathbf{x}^\top A\mathbf{x}}{\mathbf{x}^\top\mathbf{x}} \le \lambda_{\max}(A)$$

and therefore

$$\frac{1}{\lambda_{\max}(A)} \le \hat{\alpha}_k \le \frac{1}{\lambda_{\min}(A)}.$$

Substituting $\hat{\alpha}_k$ into (1.1) yields

$$\begin{aligned}\mathbf{r}^{(k+1)\top} A^{-1}\mathbf{r}^{(k+1)} &= \mathbf{r}^{(k)\top} A^{-1}\mathbf{r}^{(k)} - 2\mathbf{r}^{(k)\top}\mathbf{r}^{(k)}\frac{\mathbf{r}^{(k)\top}\mathbf{r}^{(k)}}{\mathbf{r}^{(k)\top} A\mathbf{r}^{(k)}} + \left(\frac{\mathbf{r}^{(k)\top}\mathbf{r}^{(k)}}{\mathbf{r}^{(k)\top} A\mathbf{r}^{(k)}}\right)^2 \mathbf{r}^{(k)\top} A\mathbf{r}^{(k)} \\ &= \mathbf{r}^{(k)\top} A^{-1}\mathbf{r}^{(k)} - \frac{(\mathbf{r}^{(k)\top}\mathbf{r}^{(k)})^2}{\mathbf{r}^{(k)\top} A\mathbf{r}^{(k)}}\end{aligned}$$

and therefore

$$\frac{\|\mathbf{r}^{(k+1)}\|_{A^{-1}}^2}{\|\mathbf{r}^{(k)}\|_{A^{-1}}^2} = 1 - \frac{(\mathbf{r}^{(k)\top}\mathbf{r}^{(k)})^2}{(\mathbf{r}^{(k)\top} A^{-1}\mathbf{r}^{(k)})(\mathbf{r}^{(k)\top} A\mathbf{r}^{(k)})}.$$

Date: November 30, 2005, version 1.0.
Notes originally due to James Lambers. Minor editing by Lek-Heng Lim.

The *Kantorovich inequality*, which comes up very often in applications such as optimization and statistics, states that

$$\frac{\mathbf{x}^\top A\mathbf{x}\cdot\mathbf{x}^\top A^{-1}\mathbf{x}}{(\mathbf{x}^\top\mathbf{x})^2} \le \left(\frac{\sqrt{\kappa}+\sqrt{\kappa}^{-1}}{2}\right)^2, \quad \kappa = \frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}.$$

It follows that

$$\frac{\|\mathbf{r}^{(k+1)}\|_{A^{-1}}^2}{\|\mathbf{r}^{(k)}\|_{A^{-1}}^2} \le \left(\frac{\kappa-1}{\kappa+1}\right)^2.$$

Thus,

$$\frac{\|\mathbf{r}^{(1)}\|_{A^{-1}}}{\|\mathbf{r}^{(0)}\|_{A^{-1}}}\cdot\frac{\|\mathbf{r}^{(2)}\|_{A^{-1}}}{\|\mathbf{r}^{(1)}\|_{A^{-1}}}\cdot\cdots\cdot\frac{\|\mathbf{r}^{(k)}\|_{A^{-1}}}{\|\mathbf{r}^{(k-1)}\|_{A^{-1}}} \le \left(\frac{\kappa-1}{\kappa+1}\right)^k$$

which yields

$$\frac{\|\mathbf{r}^{(k)}\|_{A^{-1}}}{\|\mathbf{r}^{(0)}\|_{A^{-1}}} \le \left(\frac{\kappa-1}{\kappa+1}\right)^k.$$

In other words, the rate of convergence is the same as when the parameter α_k is chosen a priori to be

$$\hat{\alpha} = \frac{2}{\mu_1+\mu_n}.$$

Which method is preferable? For the first approach, the problem is that we must know μ_1 and μ_n. For the second approach, we must compute α_k at each step, which is worse for computation, but in practice works better for certain problems.

Now consider the iteration

$$\mathbf{x}^{(k+1)} = (I-\alpha_k A)\mathbf{x}^{(k)} + \alpha_k\mathbf{b}.$$

Since the exact solution $\mathbf{x}$ satisfies

$$\mathbf{x} = (I-\alpha_k A)\mathbf{x} + \alpha_k\mathbf{b},$$

it follows that

$$\mathbf{e}^{(k+1)} = (I-\alpha_k A)\mathbf{e}^{(k)}.$$

So, we have

$$\begin{aligned}\mathbf{e}^{(1)} &= (I-\alpha_0 A)\mathbf{e}^{(0)}\\ &\vdots\\ \mathbf{e}^{(k)} &= (I-\alpha_{k-1}A)(I-\alpha_{k-2}A)\cdots(I-\alpha_0 A)\mathbf{e}^{(0)}.\end{aligned}$$

In other words,

$$\mathbf{e}^{(k)} = P_k(A)\mathbf{e}^{(0)}$$

where P_k is a polynomial of degree k.

By the Cayley-Hamilton theorem,

$$\psi(A) = \prod_{i=0}^{d-1}(A-\mu_i I) = 0$$

where d is the number of distinct eigenvalues μ_i of A, when $A = A^\top$. In other words

$$\psi(A) = \prod_{i=0}^{d-1}\left(I - \frac{1}{\mu_i}A\right) = 0$$

so we could choose $\alpha_i = 1/\mu_i$, but this choice is nonsense because one almost never knows the eigenvalues of A and even so, this choice is unstable because μ_i can vary immensely in magnitude. However, we have

$$\frac{\|\mathbf{e}^{(k)}\|}{\|\mathbf{e}^{(0)}\|} \le \|P_k(A)\|_2,$$

so we will now use approximation theory to find a suitable P_k.

If $A = Q\Lambda Q^\top$, then $P_k(A) = QP_k(\Lambda)Q^\top$, and therefore

$$\frac{\|\mathbf{e}^{(k)}\|}{\|\mathbf{e}^{(0)}\|} \leq \|P_k(\Lambda)\|_2.$$

And since

$$P_k(\Lambda) = \begin{bmatrix} P_k(\lambda_1) & & \\ & \ddots & \\ & & P_k(\lambda_n) \end{bmatrix},$$

it follows that

$$\|P_k(\Lambda)\|_2 \leq \max_{1\leq i\leq n} |P_k(\lambda_i)|.$$

So, because $P_k(0) = I$, we want to find a polynomial $\hat{p}_k(\lambda)$ such that $\hat{p}_k(0) = 1$ and

$$\max_{1\leq i\leq n} |\hat{p}_k(\lambda_i)| = \min_{p_k(0)=1} \max_{1\leq i\leq n} |p_k(\lambda_i)|.$$

But clearly,

$$\min_{p_k(0)=1} \max_{1\leq i\leq n} |p_k(\lambda_i)| \leq \min_{p_k(0)=1} \max_{\lambda_n\leq\lambda\leq\lambda_1} |p_k(\lambda)|.$$

Therefore, we will try to find the polynomial $\hat{p}_k$ that satisfies $\hat{p}(0) = 1$ and is of minimum absolute value on the interval $[\lambda_n, \lambda_1]$. The solution to this problem is given by the *Chebyshev polynomials*.

The Chebyshev polynomial of degree k is defined to be

$$C_k(x) = \begin{cases} \cos(k\cos^{-1}(x)) & \text{if } |x| \leq 1, \\ \cosh(k\cosh^{-1}(x)) & \text{if } |x| > 1. \end{cases}$$

For example,

$$C_0(x) = 1, \quad C_1(x) = x, \quad C_2(x) = 2x^2 - 1.$$

These polynomials are designed to be bounded by 1 in absolute value on the interval $|x| \leq 1$.

If $\theta = \cos^{-1} x$ then, using the trigonometric identities

$$\cos(k+1)\theta = \cos k\theta\cos\theta - \sin k\theta\sin\theta$$
$$\cos(k-1)\theta = \cos k\theta\cos\theta + \sin k\theta\sin\theta$$

we obtain

$$\cos(k+1)\theta = 2\cos k\theta\cos\theta - \cos(k-1)\theta$$

which yields the three-term recurrence relation of the Chebyshev polynomials

$$C_{k+1}(x) = 2xC_k(x) - C_{k-1}(x).$$

Since this relation leads to a leading coefficient of 2^{k-1} for $C_k(x)$ when $k \geq 1$, it is customary to normalize, defining

$$T_k(x) = \frac{C_k(x)}{2^{k-1}}, \quad k \geq 1.$$

We now claim that for $k = 2$, $\hat{p}_2(x)$ is

$$T_2(x) = x^2 - \frac{1}{2},$$

scaled and translated appropriately so as to be small on the interval $[\lambda_n, \lambda_1]$ and satisfy $\hat{p}_2(0) = 1$.

Note that on $[-1, 1]$, $T_2(x)$ has a maximum at $x = -1$ and $x = 1$, and a local minimum at $x = 0$. Now, suppose that there is another polynomial $p_2(x) = x^2 + bx + c$ such that $p_2(-1) < T_2(-1)$, $p_2(1) < T_2(1)$, and $p_2(0) > T_2(0)$. Then the polynomial $q_1(x) = T_2(x) - p_2(x)$ has three sign changes in the interval $[-1, 1]$, but since $T_2(x)$ and $p_2(x)$ have the same leading coefficient, $q_1(x)$ can have degree at most 1, so it must be identically zero.

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