Dr J Lellmann

Mathematical Tripos Part II: Michaelmas Term 2014

Numerical Analysis – Lecture 21

Lemma 4.29 (Properties of Krylov subspaces) *Given A and nonzero $\boldsymbol{v}$, let δ_m be the dimension of the Krylov subspace $K_m(A, \boldsymbol{v})$. Then the sequence $\{\delta_m\}_1^n$ increases monotonically and has the following properties.*

1) *There exists a positive integer $s \le n$ such that $\delta_m = m$ for $m \le s$ and $\delta_m = s$ for $m > s$.*

2) *If we can express $\boldsymbol{v}$ as $\boldsymbol{v} = \sum_{i=1}^{s'} c_i \boldsymbol{w}_i$, where $(\boldsymbol{w}_i)$ are eigenvectors of A corresponding to distinct eigenvalues and all (c_i) are nonzero, then $s = s'$.*

Remark 4.30 Assumption in the second part (regarding $\boldsymbol{v}$ and $\boldsymbol{w}_i$) does not require that all the eigenvalues of A should be distinct. It is sufficient to have n linearly independent eigenvectors.

Proof. 1) Clearly, $K_m(A, \boldsymbol{v}) \subseteq K_{m+1}(A, \boldsymbol{v}) \subseteq \mathbb{R}^n$, therefore $\delta_m \le \delta_{m+1} \le n$. We further note that $\delta_1 = 1$ (since $A^0\boldsymbol{v} = \boldsymbol{v} \neq 0$) and $\delta_m \le m$, because each subspace $K_m(A, \boldsymbol{v})$ is spanned by m vectors. Let s be the *greatest* integer such that $\delta_s = s$. Then $s = \delta_s \le \delta_{s+1} \le s$, therefore $\delta_{s+1} = \delta_s$ and the spaces $K_s(A, \boldsymbol{v})$ and $K_{s+1}(A, \boldsymbol{v})$ are the same. This implies that $A^s\boldsymbol{v}$ belongs to $K_s(A, \boldsymbol{v})$, i.e., $A^s\boldsymbol{v} = \sum_{j=0}^{s-1} a_j A^j \boldsymbol{v}$. But then

$$A^{s+r}\boldsymbol{v} = \sum_{j=0}^{s-1} a_j A^{j+r}\boldsymbol{v}, \qquad r \ge 0,$$

and that shows that the spaces $K_{s+r+1}(A, \boldsymbol{v})$ and $K_{s+r}(A, \boldsymbol{v})$ are the same for every $r \ge 0$. Therefore, for every $m > s$, we have $K_m(A, \boldsymbol{v}) = K_s(A, \boldsymbol{v})$ and respectively $\delta_m = \delta_s = s$.

2) Suppose now that $\boldsymbol{v} = \sum_{i=1}^{s'} c_i \boldsymbol{w}_i$, where $(\boldsymbol{w}_i)$ are eigenvectors of A with the corresponding distinct eigenvalues λ_i. Then $A^j\boldsymbol{v} = \sum_{i=1}^{s'} c_i \lambda_i^j \boldsymbol{w}_i$, and we deduce that

$$K_s(A, \boldsymbol{v}) \subseteq \mathrm{Sp}\{\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_{s'}\}.$$

Since eigenvectors are linearly independent, it follows that $\delta_s = s \le s'$.

Assume next that $s < s'$. We have already proved that in this case $\delta_{s'} = \delta_s = s$, therefore the vectors $(A^j\boldsymbol{v})_0^{s'-1}$ are linearly dependent. In other words, there exist $a_0, a_1, \ldots, a_{s'-1}$, not all zero, so that $p(A)\boldsymbol{v} := \sum_{j=0}^{s'-1} a_j A^j \boldsymbol{v} = 0$, where $p(x) := \sum_{j=0}^{s'-1} a_j x^j$ is a polynomial of degree $\le s' - 1$. Therefore,

$$0 = p(A)\boldsymbol{v} = p(A)\sum_{i=1}^{s'} c_i \boldsymbol{w}_i = \sum_{i=1}^{s'} p(\lambda_i) c_i \boldsymbol{w}_i.$$

Since the eigenvectors are linearly independent and all c_i are nonzero, we deduce from the above that $p(\lambda_i) = 0$ for $i = 1...s'$, i.e. that the polynomial p has s' different roots $x = \lambda_i$. But this is a contradiction because p is of degree $\le s' - 1$. Hence the assumption $s < s'$ is false, therefore $s = s'$, and the proof is complete. $\square$

Application 4.31 (Number of iterations in CGM) It follows from the previous lemma that the number of iterations of the CGM for solving $A\boldsymbol{x} = \boldsymbol{b}$ is at most the number of distinct eigenvalues of A. Further, if $\boldsymbol{b}$ is expressed as a linear combination of eigenvectors of A with distinct eigenvalues, then the number of iterations is bounded from above by the number of nonzero terms in the linear combination.

Technique 4.32 (Preconditioning) We change variables, $\boldsymbol{x} = P^T\widehat{\boldsymbol{x}}$, where P is a nonsingular $n \times n$ matrix. Thus, instead of $A\boldsymbol{x} = \boldsymbol{b}$, we are solving the linear system

$$PAP^T\widehat{\boldsymbol{x}} = P\boldsymbol{b}\,.$$

Note that symmetry and positive definiteness of A imply that PAP^T is also symmetric and positive definite. Therefore, we can apply conjugate gradients to the new system. This results in the solution $\widehat{\boldsymbol{x}}$, hence $\boldsymbol{x} = P^T\widehat{\boldsymbol{x}}$. This procedure is called the *preconditioned conjugate gradient method* and P is called the *preconditioner*.

The *condition number* $\kappa(A)$ of a symmetric positive-definite matrix A is the ratio $\lambda_{\max}/\lambda_{\min}$ between the magnitude of its largest and the least eigenvalue. The closer is this number to 1, the faster is convergence. The main idea of preconditioning is to pick P so that $\kappa(P^TAP)$ is much smaller than $\kappa(A)$, thus accelerating convergence.

The identity $(PAP^T)^jP = P(AP^TP)^j$ implies that

$$\dim K_m(PAP^T, P\boldsymbol{b}) = \dim K_m(AP^TP, \boldsymbol{b}),$$

i.e. that the dimension of the Krylov subspace for the preconditioned CGM, is equal to the dimension of $K_m(AP^TP, \boldsymbol{b})$. If we set

$$S^{-1} := P^TP =: (QQ^T)^{-1},$$

then it is suggestive to choose $S = QQ^T$ as an approximation to A which is easy to invert, so that AS^{-1} is close to identity, thus

$$\dim K_m(AP^TP, \boldsymbol{b}) = \dim K_m(AS^{-1}, \boldsymbol{b}) \approx \dim K_m(I, \boldsymbol{b}) \ll n.$$

1) The simplest choice of S is $D = \operatorname{diag} A$.

2) Another possibility is to choose S as a band matrix with small bandwidth. For example, solving the Poisson equation with the five-point formula, we may take S to be the tridiagonal part of A. In that case we commence with the Cholesky factorization of $S = QQ^T$, so that $S^{-1} = Q^{-T}Q$, hence $P = Q^{-1}$. The main expense in each step of the method is the computation of

$$\boldsymbol{z} = P\boldsymbol{y} = Q^{-1}\boldsymbol{y}$$

for some $\boldsymbol{y} \in \mathbb{R}^n$, but note that computing $Q^{-1}\boldsymbol{y}$ is the same as solving the linear system $Q\boldsymbol{z} = \boldsymbol{y}$, which is cheap as Q is a triangular matrix.

3) One can also take $P = L^{-1}$, where L is the lower triangular part of A (maybe imposing some changes). For example, for the Poisson equation, with $m = 20$ hence dealing with 400×400 system, we take P^{-1} as the lower triangular part of A, but change the diagonal elements from 4 to $\frac{5}{2}$. Then we get a computer precision after just 30 iterations.

Example 4.33 For the tridiagonal system $A\boldsymbol{x} = \boldsymbol{b}$, we choose the preconditioner as follows.

$$A = \begin{bmatrix} 2 & -1 & & \\ -1 & 2 & \ddots & \\ & \ddots & \ddots & -1 \\ & & -1 & 2 \end{bmatrix}, \qquad Q = \begin{bmatrix} 1 & & & \\ -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix}, \qquad S = QQ^T = \begin{bmatrix} 1 & -1 & & \\ -1 & 2 & \ddots & \\ & \ddots & \ddots & -1 \\ & & -1 & 2 \end{bmatrix}.$$

The matrix S coincides with A except at the $(1,1)$ entry. The matrix $C = Q^{-T}AQ^{-1}$ for the preconditioned CGM has just two distinct eigenvalues, and we recover the exact solution just in two steps.

Matlab demo: Download the Matlab GUI for *Preconditioning of Conjugate Gradient* from `http://www.maths.cam.ac.uk/undergrad/course/na/ii/precond/precond.php`. Run the GUI to solve different systems of linear equations, trying different preconditioners P. You can select from some preset preconditioners but can propose your own customised preconditioners as well. What does preconditioning do to the spectrum of the system matrix?