Mathematical Tripos Part II: Michaelmas Term 2020

Numerical Analysis – Lecture 20

Algorithm 4.26 (The conjugate gradient method) Here it is.

(A) For any initial vector $\boldsymbol{x}^{(0)}$, set $\boldsymbol{d}^{(0)} = \boldsymbol{r}^{(0)} = \boldsymbol{b} - A\boldsymbol{x}^{(0)}$;

(B) For $k \geq 0$, calculate $\boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + \alpha_k \boldsymbol{d}^{(k)}$ and the residual

$$\boldsymbol{r}^{(k+1)} = \boldsymbol{r}^{(k)} - \alpha_k A\boldsymbol{d}^{(k)}, \quad \text{with} \quad \alpha_k := \{\, \boldsymbol{r}^{(k+1)} \perp \boldsymbol{d}^{(k)} \} = \frac{(\boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)})}{(A\boldsymbol{d}^{(k)}, \boldsymbol{d}^{(k)})}\,, \quad k \geq 0\,. \tag{4.8}$$

(C) For the same k, the next conjugate direction is the vector

$$\boldsymbol{d}^{(k+1)} = \boldsymbol{r}^{(k+1)} + \beta_k \boldsymbol{d}^{(k)}, \quad \text{with} \quad \beta_k := \{\, \boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(k)} \} = -\frac{(\boldsymbol{r}^{(k+1)}, A\boldsymbol{d}^{(k)})}{(\boldsymbol{d}^{(k)}, A\boldsymbol{d}^{(k)})}, \quad k \geq 0\,. \tag{4.9}$$

Theorem 4.27 (Properties of CGM) *For every $m \geq 0$, the conjugate gradient method has the following properties.*

(1) *The linear space spanned by the residuals $\{\boldsymbol{r}^{(i)}\}$ is the same as the linear space spanned by the conjugate directions $\{\boldsymbol{d}^{(i)}\}$ and it coincides with the space spanned by $\{A^i\boldsymbol{r}^{(0)}\}$:*

$$\operatorname{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^m = \operatorname{span}\{\boldsymbol{d}^{(i)}\}_{i=0}^m = \operatorname{span}\{A^i\boldsymbol{r}^{(0)}\}_{i=0}^m\,.$$

(2) *The residuals satisfy the orthogonality conditions: $(\boldsymbol{r}^{(m)}, \boldsymbol{r}^{(i)}) = (\boldsymbol{r}^{(m)}, \boldsymbol{d}^{(i)}) = 0$ for $i < m$.*

(3) *The directions are conjugate (A-orthogonal): $(\boldsymbol{d}^{(m)}, \boldsymbol{d}^{(i)})_A = (\boldsymbol{d}^{(m)}, A\boldsymbol{d}^{(i)}) = 0$ for $i < m$.*

Proof. We use induction on $m \geq 0$, the assertions being trivial for $m = 0$, since $\boldsymbol{d}^{(0)} = \boldsymbol{r}^{(0)}$ and (2)-(3) are void. Therefore, assuming that the assertions are true for some $m = k$, we ask if they remain true when $m = k + 1$.

(1) Formula (4.9)

$$\boldsymbol{d}^{(k+1)} = \boldsymbol{r}^{(k+1)} + \beta_k \boldsymbol{d}^{(k)}$$

readily implies that equivalence of the spaces spanned by $(\boldsymbol{r}^{(i)})_0^k$ and $(\boldsymbol{d}^{(i)})_0^k$, is preserved when k is increased to $k+1$. Similarly, from $\boldsymbol{r}^{(k+1)} = \boldsymbol{r}^{(k)} - \alpha_k A\boldsymbol{d}^{(k)}$ in (4.8), and from the inductive assumption $\boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)} \in \operatorname{span}\{A^i\boldsymbol{r}^{(0)}\}_{i=0}^k$, it follows that $\boldsymbol{r}^{(k+1)} \in \operatorname{span}\{A^i\boldsymbol{r}^{(0)}\}_{i=0}^{k+1}$.

(2) Turning to assertion (2), we need $\boldsymbol{r}^{(k+1)} \perp \boldsymbol{r}^{(i)}$ for $i \leq k$, which by (1) is equivalent to

$$\boldsymbol{r}^{(k+1)} \perp \boldsymbol{d}^{(i)} \quad \text{for} \quad i \leq k\,.$$

We have $\boldsymbol{r}^{(k+1)} \perp \boldsymbol{d}^{(k)}$ by the definition of α_k in (4.8), so we need

$$\boldsymbol{r}^{(k+1)} \stackrel{(4.8)}{=} \boldsymbol{r}^{(k)} - \alpha_k A\boldsymbol{d}^{(k)} \perp \boldsymbol{d}^{(i)} \quad \text{for} \quad i < k\,,$$

and this follow from the inductive assumptions $\boldsymbol{r}^{(k)} \perp \boldsymbol{d}^{(i)}$ and $A\boldsymbol{d}^{(k)} \perp \boldsymbol{d}^{(i)}$.

(3) It remains to justify (3), namely that $\boldsymbol{d}^{(k+1)}$ defined in (4.9) satisfies

$$\boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(i)} \quad \text{for} \quad i \leq k\,.$$

The value of β_k in (4.9) is defined to give $\boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(k)}$, so we need

$$\boldsymbol{d}^{(k+1)} \stackrel{(4.9)}{=} \boldsymbol{r}^{(k+1)} + \beta_k \boldsymbol{d}^{(k)} \perp A\boldsymbol{d}^{(i)} \quad \text{for} \quad i < k\,.$$

By the inductive hypothesis $\boldsymbol{d}^{(k)} \perp A\boldsymbol{d}^{(i)}$, hence it remains to establish that $\boldsymbol{r}^{(k+1)} \perp A\boldsymbol{d}^{(i)}$ for $i < k$. Now, the formula (4.8) yields $A\boldsymbol{d}^{(i)} = (\boldsymbol{r}^{(i)} - \boldsymbol{r}^{(i+1)})/\alpha_i$, therefore we require the conditions $\boldsymbol{r}^{(k+1)} \perp (\boldsymbol{r}^{(i)} - \boldsymbol{r}^{(i+1)})$ for $i < k$, and they are a consequence of the assertion (2) for $m = k+1$ obtained previously. □

Corollary 4.28 (A termination property) *If the conjugate gradient method is applied in exact arithmetic, then, for any $\boldsymbol{x}^{(0)} \in \mathbb{R}^n$, termination occurs after at most n iterations. More precisely, termination occurs after at most s iterations, where $s = \dim\operatorname{span}\{A^i\boldsymbol{r}_0\}_{i=0}^{n-1}$ (which can be smaller than n).*

Proof. Assertion (2) of Theorem 4.27 states that residuals $(\boldsymbol{r}^{(k)})_{k\ge 0}$ form a sequence of mutually orthogonal vectors in $\mathbb{R}^n$, therefore at most n of them can be nonzero. Since they also belong to the space $\operatorname{span}\{A^i\boldsymbol{r}_0\}_{i=0}^{n-1}$, their number is bounded by the dimension of that space. □

Definition 4.29 (The Krylov subspaces) Let A be an $n\times n$ matrix, $\boldsymbol{v}\in\mathbb{R}^n$ nonzero, and $m\in\mathbb{N}$. The linear space $K_m(A,\boldsymbol{v}) := \operatorname{span}\{A^i\boldsymbol{v}\}_{i=0}^{m-1}$ is called the *m-th Krylov subspace* of $\mathbb{R}^n$.

Theorem 4.30 (Number of iterations in CGM) *Let $A > 0$, and let s be the number of its distinct eigenvalues. Then, for any $\boldsymbol{v}$,*

$$\dim K_m(A,\boldsymbol{v}) \le s \quad \forall\, m\,. \tag{4.10}$$

Hence, for any $A>0$, the number of iterations of the CGM for solving $A\boldsymbol{x}=\boldsymbol{b}$ is bounded by the number of distinct eigenvalues of A.

Proof. Inequality (4.10) is true not just for positive definite $A>0$, but for any A with n linearly independent eigenvectors $(\boldsymbol{u}_i)$. Indeed, in that case one can expand $\boldsymbol{v} = \sum_{i=1}^n a_i\boldsymbol{u}_i$, and then group together eigenvectors with the same eigenvalues: for each λ_ν we set $\boldsymbol{w}_\nu = \sum_{k=1}^{m_\nu} a_{i_k}\boldsymbol{u}_{i_k}$ if $A\boldsymbol{u}_{i_k} = \lambda_\nu \boldsymbol{u}_{i_k}$. Then

$$\boldsymbol{v} = \textstyle\sum_{\nu=1}^s c_\nu \boldsymbol{w}_\nu\,, \qquad c_\nu \in \{0,1\}\,,$$

hence $A^i\boldsymbol{v} = \sum_{\nu=1}^s c_\nu\lambda_\nu^i\boldsymbol{w}_\nu$, thus for any m we get $K_m(A,\boldsymbol{v}) \subseteq \operatorname{span}\{\boldsymbol{w}_1,\boldsymbol{w}_2,\ldots,\boldsymbol{w}_s\}$, and that proves (4.10). By Corollary 4.28, the number of iteration in CGM is bounded by $\dim K_m(A,\boldsymbol{r}^{(0)})$, hence the final conclusion. □

Remark 4.31 Theorem 4.30 shows that, unlike other iterative schemes, the conjugate gradient method is both iterative and direct: each iteration produces a reasonable approximation to the exact solution, and the exact solution itself will be recovered after n iterations at most.

We now simplify and reformulate Algorithm 4.26.

Firstly, we rewrite expressions for the parameters α_k and β_k in (4.8)-(4.9) as follows:

$$\alpha_k = \frac{(\boldsymbol{r}^{(k)},\boldsymbol{d}^{(k)})}{(\boldsymbol{d}^{(k)},A\boldsymbol{d}^{(k)})} \overset{(c)}{=} \frac{\|\boldsymbol{r}^{(k)}\|^2}{(\boldsymbol{d}^{(k)},A\boldsymbol{d}^{(k)})} > 0\,,$$

$$\beta_k = -\frac{(\boldsymbol{r}^{(k+1)},A\boldsymbol{d}^{(k)})}{(\boldsymbol{d}^{(k)},A\boldsymbol{d}^{(k)})} \overset{(a)}{=} -\frac{(\boldsymbol{r}^{(k+1)},\boldsymbol{r}^{(k+1)}-\boldsymbol{r}^{(k)})}{(\boldsymbol{d}^{(k)},\boldsymbol{r}^{(k+1)}-\boldsymbol{r}^{(k)})} \overset{(b)}{=} \frac{\|\boldsymbol{r}^{(k+1)}\|^2}{(\boldsymbol{d}^{(k)},\boldsymbol{r}^{(k)})} \overset{(c)}{=} \frac{\|\boldsymbol{r}^{(k+1)}\|^2}{\|\boldsymbol{r}^{(k)}\|^2} > 0\,.$$

Here, for β, we used in (a) the fact that $A\boldsymbol{d}^{(k)}$ is a multiple of $\boldsymbol{r}^{(k+1)} - \boldsymbol{r}^{(k)}$ by (4.8), and in (b) orthogonality of $\boldsymbol{r}^{(k+1)}$ to both $\boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)}$ proved in Theorem 4.27(2). Then, for both β and α, we used in (c) the property $(\boldsymbol{d}^{(k)},\boldsymbol{r}^{(k)}) = \|\boldsymbol{r}^{(k)}\|^2$ which follows from (4.9) with index $k+1$, taking in account orthogonality $\boldsymbol{r}^{(k+1)}\perp\boldsymbol{d}^{(k)}$.

Secondly, we let $\boldsymbol{x}^{(0)}$ be the zero vector.

Algorithm 4.32 (Standard form of the conjugate gradient method) Here it is.

(1) Set $k=0$, $\boldsymbol{x}^{(0)}=0$, $\boldsymbol{r}^{(0)}=\boldsymbol{b}$, and $\boldsymbol{d}^{(0)}=\boldsymbol{r}^{(0)}$;

(2) Calculate the matrix-vector product $\boldsymbol{v}^{(k)} = A\boldsymbol{d}^{(k)}$ and $\alpha_k = \|\boldsymbol{r}^{(k)}\|^2/(\boldsymbol{d}^{(k)},\boldsymbol{v}^{(k)}) > 0$;

(3) Apply the formulae $\boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + \alpha_k\boldsymbol{d}^{(k)}$ and $\boldsymbol{r}^{(k+1)} = \boldsymbol{r}^{(k)} - \alpha_k\boldsymbol{v}^{(k)}$;

(4) Stop if $\|\boldsymbol{r}^{(k+1)}\|$ is acceptably small;

(5) Set $\boldsymbol{d}^{(k+1)} = \boldsymbol{r}^{(k+1)} + \beta_k\boldsymbol{d}^{(k)}$, where $\beta_k = \|\boldsymbol{r}^{(k+1)}\|^2/\|\boldsymbol{r}^{(k)}\|^2 > 0$;

(6) Increase $k\to k+1$ and go back to (2).

The total work is dominated by the number of iterations, multiplied by the time it takes to compute $\boldsymbol{v}^{(k)} = A\boldsymbol{d}^{(k)}$. Thus the conjugate gradient algorithm is highly suitable when most of the elements of A are zero, i.e. when A is *sparse*.