Dr J Lellmann

Mathematical Tripos Part II: Michaelmas Term 2014

Numerical Analysis – Lecture 20

Algorithm 4.23 (The conjugate gradient method) Here it is.

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

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

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

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

$$\boldsymbol{d}^{(k+1)} = -\boldsymbol{g}^{(k+1)} + \beta^{(k)}\boldsymbol{d}^{(k)}, \quad \text{with} \quad \beta^{(k)} = \{\, \boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(k)} \} = \frac{\boldsymbol{g}^{(k+1)T}A\boldsymbol{d}^{(k)}}{\boldsymbol{d}^{(k)T}A\boldsymbol{d}^{(k)}}, \quad k \geq 0. \tag{4.10}$$

Theorem 4.24 (Properties of Algorithm 4.23) *For every integer $m \geq 0$, the conjugate gradient method enjoys the following properties.*

(1) *The linear space spanned by the gradients $\{\boldsymbol{g}^{(i)} : i = 0...m\}$*

(a) *is the same as the linear space spanned by the search directions $\{\boldsymbol{d}^{(i)} : i = 0...m\}$*

(b) *it coincides with the space $K_{m+1} = \operatorname{span}\{A^i\boldsymbol{g}^{(0)} : i = 0..m\}$.*

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

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

Proof. We use induction on $m \geq 0$, the assertions being trivial for $m = 0$, since $\boldsymbol{d}^{(0)} = -\boldsymbol{g}^{(0)}$, 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 $\boldsymbol{d}^{(k+1)} = -\boldsymbol{g}^{(k+1)} + \beta^{(k)}\boldsymbol{d}^{(k)}$ in (4.10) readily implies that (1a), i.e. equivalence of the spaces spanned by $(\boldsymbol{g}^{(i)})_0^k$ and $(\boldsymbol{d}^{(i)})_0^k$, is preserved when k is increased to $k + 1$. Similarly, it follows from $\boldsymbol{g}^{(k+1)} = \boldsymbol{g}^{(k)} + \omega^{(k)}A\boldsymbol{d}^{(k)}$ in (4.9), that (1b) holds for $m = k + 1$ as well.

(2) Turning to assertion (2), we need $\boldsymbol{g}^{(k+1)} \perp \boldsymbol{g}^{(i)}$ for $i \leq k$, which is equivalent to $\boldsymbol{g}^{(k+1)} \perp \boldsymbol{d}^{(i)}$ for $i \leq k$ because of (1a). The latter follows from (4.9): for $i = k$ by definition of $\omega^{(k)}$, and for $i < k$ by the inductive assumptions $\boldsymbol{g}^{(k)} \perp \boldsymbol{d}^{(i)}$ and $A\boldsymbol{d}^{(k)} \perp \boldsymbol{d}^{(i)}$.

(3) It remains to justify (3), namely that $\boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(i)}$ in (4.10). The value of $\beta^{(k)}$ in (4.10) is defined to give $\boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(k)}$, so we need $\boldsymbol{d}^{(k+1)} \perp A\boldsymbol{d}^{(i)}$ for $i < k$. By the inductive hypothesis $\boldsymbol{d}^{(k)} \perp A\boldsymbol{d}^{(i)}$, hence it is sufficient to establish that $\boldsymbol{g}^{(k+1)} \perp A\boldsymbol{d}^{(i)}$ for $i < k$. Now, the formula (4.9) yields $A\boldsymbol{d}^{(i)} = (\boldsymbol{g}^{(i+1)} - \boldsymbol{g}^{(i)})/\omega^{(i)}$, therefore we require the conditions $\boldsymbol{g}^{(k+1)} \perp (\boldsymbol{g}^{(i+1)} - \boldsymbol{g}^{(i)})$ for $i < k$, and they are a consequence of the assertion (2) for $m = k + 1$ obtained previously. □

Corollary 4.25 (A termination property) *If Algorithm 4.23 is applied in exact arithmetic, then, for any $\boldsymbol{x}^{(0)} \in \mathbb{R}^n$, termination occurs after at most n iterations.*

Proof. Assertion (2) of Theorem 4.24 states that residuals $(\boldsymbol{g}^{(k)})_{k\geq 0}$ form a sequence of mutually orthogonal vectors in $\mathbb{R}^n$. Therefore at most n of them can be nonzero. □

Standard Form 4.26 (Reformulation of the conjugate gradient method) We now simplify and reformulate Algorithm 4.23. Specifically, we write the parameters $\omega^{(k)}$ and $\beta^{(k)}$ in (4.9)-(4.10) as

$$\omega^{(k)} = -\frac{\boldsymbol{d}^{(k)T}\boldsymbol{g}^{(k)}}{\boldsymbol{d}^{(k)T}A\boldsymbol{d}^{(k)}} = \frac{\|\boldsymbol{g}^{(k)}\|^2}{\boldsymbol{d}^{(k)T}A\boldsymbol{d}^{(k)}} > 0\,, \qquad \beta^{(k)} = \frac{\boldsymbol{g}^{(k+1)T}(\boldsymbol{g}^{(k+1)} - \boldsymbol{g}^{(k)})}{\boldsymbol{d}^{(k)T}(\boldsymbol{g}^{(k+1)} - \boldsymbol{g}^{(k)})} = \frac{\|\boldsymbol{g}^{(k+1)}\|^2}{\|\boldsymbol{g}^{(k)}\|^2} > 0\,.$$

Here we used (for β) the fact that $A\boldsymbol{d}^{(k)}$ is a multiple of $\boldsymbol{g}^{(k+1)} - \boldsymbol{g}^{(k)}$ and orthogonality of $\boldsymbol{g}^{(k+1)}$ to both $\boldsymbol{g}^{(k)}, \boldsymbol{d}^{(k)}$ proved above, and (for both β and ω) the property $\boldsymbol{d}^{(k)T}\boldsymbol{g}^{(k)} = -\|\boldsymbol{g}^{(k)}\|^2$ which follows from (4.10) with index $k+1$. Furthermore, we let $\boldsymbol{x}^{(0)}$ be the zero vector and we write $-\boldsymbol{r}^{(k)}$ instead of $\boldsymbol{g}^{(k)}$, where $\boldsymbol{r}^{(k)}$ is the (sign reversed) residual $\boldsymbol{b} - A\boldsymbol{x}^{(k)}$.

Thus, Algorithm 4.23 takes the following form.

(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 $\omega^{(k)} = \|\boldsymbol{r}^{(k)}\|^2/\boldsymbol{d}^{(k)T}\boldsymbol{v}^{(k)} > 0$;

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

(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 by one, and then go back to (2).

The total work is usually dominated by the number of iterations, multiplied by the time it takes to compute $\boldsymbol{v}^{(k)} = A\boldsymbol{d}^{(k)}$. It follows from Corollary 4.25 that the conjugate gradient algorithm is highly suitable when most of the elements of A are zero, i.e. when A is *sparse*.

Definition 4.27 (Krylov subspace) Let A be an $n \times n$ matrix, $\boldsymbol{v} \in \mathbb{R}^n$ nonzero, and $m \in \mathbb{N}$. The linear space $K_m(A, \boldsymbol{v}) = \mathrm{Sp}\{A^j\boldsymbol{v} : j = 0...m-1\}$ is said to be the *mth Krylov subspace* of $\mathbb{R}^n$.

Remark 4.28 (The Krylov subspaces of the conjugate gradient method) In the standard form of the method, we set $\boldsymbol{d}^{(0)} = -\boldsymbol{g}^{(0)} = \boldsymbol{b} \in K_1(A, \boldsymbol{b})$, and from the formulas

$$\boldsymbol{g}^{(k+1)} = \boldsymbol{g}^{(k)} + \omega^{(k)}A\boldsymbol{d}^{(k)}, \qquad \boldsymbol{d}^{(k+1)} = -\boldsymbol{g}^{(k+1)} + \beta^{(k)}\boldsymbol{d}^{(k)}$$

we deduced by induction that

$$\mathrm{Sp}\,\{\boldsymbol{g}^{(0)}, \boldsymbol{g}^{(1)}, \ldots, \boldsymbol{g}^{(m)}\} = \mathrm{Sp}\,\{\boldsymbol{g}^{(0)}, A\boldsymbol{g}^{(0)}, \ldots, A^m\boldsymbol{g}^{(0)}\} = K_{m+1}(A, \boldsymbol{b})\,.$$

By Theorem 4.24, the residuals $\boldsymbol{g}^{(i)}$ are orthogonal to each other, thus, the number of nonzero residuals (and hence the number of iterations) is bounded from above by the largest dimension of the subspaces $K_m(A, \boldsymbol{b})$. The latter is n at most, but it can be smaller as the following consideration shows.

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 $s = s'$.*