Dr H. Fawzi

Mathematical Tripos Part II: Michaelmas Term 2021

Numerical Analysis – Lecture 19

Conjugate gradient method The conjugate gradient method is the method of conjugate directions (Theorem 4.23 from previous lecture) where the directions $\boldsymbol{d}^{(i)}$ are chosen so that they A-orthogonalize the residuals, i.e., the $\boldsymbol{d}^{(i)}$ satisfy

$$\mathrm{span}(\boldsymbol{d}^{(0)},\dots,\boldsymbol{d}^{(k-1)}) = \mathrm{span}(\boldsymbol{r}^{(0)},\dots,\boldsymbol{r}^{(k-1)}) \tag{4.8}$$

for every iteration k, in addition to being pairwise A-orthogonal. This can be achieved by setting $\boldsymbol{d}^{(0)} = \boldsymbol{r}^{(0)}$ and applying the Gram-Schmidt step at each iteration

$$\boldsymbol{d}^{(k+1)} = \boldsymbol{r}^{(k+1)} - \sum_{i\le k} \frac{\langle \boldsymbol{r}^{(k+1)}, \boldsymbol{d}^{(i)}\rangle_A}{\langle \boldsymbol{d}^{(i)}, \boldsymbol{d}^{(i)}\rangle_A}\boldsymbol{d}^{(i)}. \tag{4.9}$$

Because of our particular choice of $\boldsymbol{d}^{(k)}$, the equation above simplifies dramatically and the terms $i \le k-1$ in the summation above happen to be zero! This is the *key point* of the CG method. Let's prove this. Recall that the iterates are defined by $\boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + \alpha_k \boldsymbol{d}^{(k)}$ so that the residuals satisfy $\boldsymbol{r}^{(k+1)} = \boldsymbol{r}^{(k)} - \alpha_k A\boldsymbol{d}^{(k)}$. It is easy then to see by induction, using the property (4.8) that

$$\mathrm{span}\{\boldsymbol{d}^{(i)}\}_{i=0}^{k-1} = \mathrm{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^{k-1} = \mathrm{span}\{\boldsymbol{r}^{(0)}, A\boldsymbol{r}^{(0)},\dots,A^{k-1}\boldsymbol{r}^{(0)}\} =: K_k(A,\boldsymbol{r}^{(0)}),$$

where $K_m(A,\boldsymbol{v}) = \mathrm{span}\{A^i\boldsymbol{v}\}_{i=0}^{m-1}$ is the m'th *Krylov* subspace of A wrt $\boldsymbol{v}$. The result of Theorem 4.23 tells us that $\boldsymbol{r}^{(k+1)}$ is orthogonal to $K_{k+1}(A,\boldsymbol{r}^{(0)})$. Now for $i<k$, we have $\boldsymbol{d}^{(i)} \in K_k(A,\boldsymbol{r}^{(0)})$ and so $A\boldsymbol{d}^{(i)} \in K_{k+1}(A,\boldsymbol{r}^{(0)})$. This implies that $\langle \boldsymbol{r}^{(k+1)}, A\boldsymbol{d}^{(i)}\rangle = 0$ for $i<k$, and shows that the terms $i<k$ in Equation (4.9) are equal to zero.

The conjugate gradient algorithm can thus be summarized in the following: Set $\boldsymbol{d}^{(0)} = \boldsymbol{r}^{(0)} = \boldsymbol{b} - A\boldsymbol{x}^{(0)}$ and iterate, for $k \ge 0$:

$$\begin{cases} \boldsymbol{x}^{(k+1)} &= \boldsymbol{x}^{(k)} + \alpha_k \boldsymbol{d}^{(k)} \qquad \alpha_k = \dfrac{\langle \boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)}\rangle}{\langle \boldsymbol{d}^{(k)}, A\boldsymbol{d}^{(k)}\rangle} \\[2ex] \boldsymbol{d}^{(k+1)} &= \boldsymbol{r}^{(k+1)} + \beta_k \boldsymbol{d}^{(k)} \qquad \beta_k = -\dfrac{\langle \boldsymbol{r}^{(k+1)}, A\boldsymbol{d}^{(k)}\rangle}{\langle \boldsymbol{d}^{(k)}, A\boldsymbol{d}^{(k)}\rangle} \end{cases} \tag{4.10}$$

where $\boldsymbol{r}^{(k)}$ stands for $\boldsymbol{b} - A\boldsymbol{x}^{(k)}$. We can summarize the properties of the Conjugate Gradient Method in the following theorem.

Theorem 4.26 (Properties of CGM) *For every $m \ge 0$, 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 $\{A^i\boldsymbol{r}^{(0)}\}$:*

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

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

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

Using these properties we can simplify the expressions for α_k and β_k. Indeed, using the second equation in (4.10), and the fact that $\boldsymbol{r}^{(k)} \perp \boldsymbol{d}^{(k-1)}$, we have

$$\langle \boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)}\rangle = \langle \boldsymbol{r}^{(k)}, \boldsymbol{r}^{(k)}\rangle = \|\boldsymbol{r}^{(k)}\|_2^2 \tag{4.11}$$

which shows that

$$\alpha_k = \frac{\|\boldsymbol{r}^{(k)}\|_2^2}{\langle \boldsymbol{d}^{(k)}, A\boldsymbol{d}^{(k)}\rangle} > 0.$$

Also, we can write:

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

where we used in (a) the fact that $A\boldsymbol{d}^{(k)}$ is a multiple of $\boldsymbol{r}^{(k+1)} - \boldsymbol{r}^{(k)}$, and in (b) orthogonality of $\boldsymbol{r}^{(k+1)}$ to both $\boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)}$ (Theorem 4.26(2)), and in (c) we used (4.11).

The complete conjugate gradient method can thus be written as follows:

Algorithm 4.27 (Standard form of the conjugate gradient method) –

(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 $\alpha_k = \|\boldsymbol{r}^{(k)}\|^2 / \langle \boldsymbol{d}^{(k)}, \boldsymbol{v}^{(k)}\rangle > 0$;

(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*.

Finite termination We have already seen that the method of conjugate directions (Theorem 4.23 in previous lecture) terminates after at most n steps. We restate this result in the special case of the conjugate gradient method.

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 $s = \dim \mathrm{span}\{A^i \boldsymbol{r}_0\}_{i=0}^{n-1}$ (which can be smaller than n).*

Proof. Assertion (2) of Theorem 4.26 states that residuals $(\boldsymbol{r}^{(k)})_{k\geq 0}$ 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 $\mathrm{span}\{A^i \boldsymbol{r}_0\}_{i=0}^{n-1}$, their number is bounded by the dimension of that space. □

We can bound the dimension of the Krylov subspace $\mathrm{span}\{A^i \boldsymbol{r}_0\}_{i=0}^{n-1}$ using the number of distinct eigenvalues of A.

Theorem 4.29 (Number of iterations in CGM) *Let $A > 0$, and let s be the number of its distinct eigenvalues. Then, for any $\boldsymbol{v}$,*

$$\dim K_m(A, \boldsymbol{v}) \leq s \quad \forall\, m\,. \tag{4.12}$$

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 (4.12) 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 $K_m(A, \boldsymbol{v}) \subseteq \mathrm{span}\{\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_s\}$, and that proves (4.12). 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 4.30 Theorem 4.29 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.

Convergence One can prove a more quantitative version of Theorem (4.29).

Theorem 4.31 *Let A be symmetric positive definite. After k iterations of the conjugate gradient method, the error $\boldsymbol{e}^{(k)} = \boldsymbol{x}^* - \boldsymbol{x}^{(k)}$ satisfies*

$$\|\boldsymbol{e}^{(k)}\|_A = \min_{P_k} \|P_k(A)\boldsymbol{e}^{(0)}\|_A$$

where the minimization is over all polynomials P_k of degree $\leq k$ that satisfy $P_k(0) = 1$.

Proof. We know from Lecture 18 (Equation (4.7) applied recursively) that $\boldsymbol{e}^{(k)}$ is obtained from $\boldsymbol{e}^{(0)}$ by projecting out (in the inner product $\langle \cdot, \cdot \rangle_A$) the components $\boldsymbol{d}^{(0)}, \ldots, \boldsymbol{d}^{(k-1)}$. This means that

$$\|\boldsymbol{e}^{(k)}\|_A = \min_{\boldsymbol{v}} \|\boldsymbol{e}^{(0)} - \boldsymbol{v}\|_A$$

where the minimization is over all $\boldsymbol{v} \in \operatorname{span}(\boldsymbol{d}^{(0)}, \ldots, \boldsymbol{d}^{(k-1)})$. For the conjugate gradient method, this subspace is the same as $\operatorname{span}(\boldsymbol{r}^{(0)}, \ldots, A^{k-1}\boldsymbol{r}^{(0)})$, and since $\boldsymbol{r}^{(0)} = A\boldsymbol{e}^{(0)}$, this means that any such $\boldsymbol{v}$ can be written as $\boldsymbol{v} = \sum_{i=1}^k c_i A^i \boldsymbol{e}^{(0)}$. Let $P_k(t) = 1 - \sum_{i=1}^k c_i t^i$ we get the desired equality.

Remark 4.32 *If A has s distinct eigenvalues $\lambda_1, \ldots, \lambda_s > 0$, then with $P_s(t) = \prod_{i=1}^s (1 - t/\lambda_i)$ we have $\deg P_s = s$, $P_s(0) = 1$, and and $P_s(A) = 0$. Thus this shows that the CG method terminates after s iterations.*