Dr H. Fawzi

Mathematical Tripos Part II: Michaelmas Term 2022

Numerical Analysis – Lecture 19

Conjugate gradient method The conjugate gradient method is the method of conjugate directions (Theorem 4.22 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

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

for every iteration k, in addition to being pairwise A-orthogonal. This can be achieved by applying the Gram-Schmidt procedure to the residuals $\{\boldsymbol{r}^{(i)}\}$. The conjugate gradient method can thus be expressed using the following iterates. Starting from $\boldsymbol{x}^{(0)} \in \mathbb{R}^n$, and letting $\boldsymbol{r}^{(0)} = \boldsymbol{b} - A\boldsymbol{x}^{(0)}$, iterate, for $k \geq 0$:

$$\begin{cases} \boldsymbol{d}^{(k)} = \boldsymbol{r}^{(k)} - \sum_{i<k} \frac{\langle \boldsymbol{r}^{(k)}, \boldsymbol{d}^{(i)} \rangle_A}{\langle \boldsymbol{d}^{(i)}, \boldsymbol{d}^{(i)} \rangle_A} \boldsymbol{d}^{(i)} \\ \boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + \dfrac{\langle \boldsymbol{r}^{(k)}, \boldsymbol{d}^{(k)} \rangle}{\langle \boldsymbol{d}^{(k)}, A\boldsymbol{d}^{(k)} \rangle} \boldsymbol{d}^{(k)}. \end{cases} \tag{4.9}$$

As written, the iterates above are not attractive from a computational point of view, because they require computing at least $k-1$ inner products at each step k. However it turns out that the equation above defining $\boldsymbol{d}^{(k)}$ can be simplified dramatically and the terms $i < k-1$ in the summation above happen to be zero! This is one of the key points of the CG method. To prove this, we first introduce the following important definition.

Definition 4.26 *Let $\boldsymbol{v} \in \mathbb{R}^n$. The m'th* Krylov *subspace of A with respect to $\boldsymbol{v}$ is*

$$K_m(A, \boldsymbol{v}) = \operatorname{span}\{A^i \boldsymbol{v}\}_{i=0}^{m-1}.$$

The following theorem collects important observations about the iterates (4.9).

Theorem 4.27 (Properties of CGM) *For every $m \geq 0$ such that $\boldsymbol{r}^{(m)} \neq 0$, the conjugate gradient method has the following properties.*

(1) The linear space spanned by the residuals $\{\boldsymbol{r}^{(i)}\}_{i=0}^m$ is the same as the linear space spanned by the conjugate directions $\{\boldsymbol{d}^{(i)}\}_{i=0}^m$ and it coincides with the Krylov subspace $K_{m+1}(A, \boldsymbol{r}^{(0)})$:

$$\{\boldsymbol{r}^{(i)}\}_{i=0}^m \;=\; \{\boldsymbol{d}^{(i)}\}_{i=0}^m \;=\; K_{m+1}(A, \boldsymbol{r}^{(0)})\,.$$

(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 $i < m$.

Proof. (1) For notational convenience, we let $K_m = K_m(A, \boldsymbol{r}^{(0)})$.
By the properties of the Gram-Schmidt procedure we know $\operatorname{span}\{\boldsymbol{d}^{(i)}\}_{i=0}^m = \operatorname{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^m$. Furthermore, from Theorem 4.22 of the previous lecture, we know that $\{\boldsymbol{r}^{(0)}, \dots, \boldsymbol{r}^{(m)}\}$ is an orthogonal basis of this subspace. To show that this subspace is the Krylov subspace K_{m+1}, we proceed by induction. The case $m = 0$ is trivial. For $m \geq 1$, note that since the iterates $\boldsymbol{x}^{(k)}$ satisfy $\boldsymbol{x}^{(m+1)} = \boldsymbol{x}^{(m)} + \alpha_m \boldsymbol{d}^{(m)}$, we get $\boldsymbol{r}^{(m+1)} = \boldsymbol{r}^{(m)} - \alpha_m A\boldsymbol{d}^{(m)}$. By the induction hypothesis, $\boldsymbol{r}^{(m)}, \boldsymbol{d}^{(m)} \in K_{m+1}$, and so $A\boldsymbol{d}^{(m)} \in K_{m+2}$. Thus $\boldsymbol{r}^{(m+1)} \in K_{m+2}$ as desired. This shows the inclusion $\operatorname{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^{m+1} \subseteq K_{m+2}$; and one can show that we have equality by dimension counting. Indeed since $\boldsymbol{r}^{(m+1)} \neq 0$ (by assumption), we know that $\boldsymbol{r}^{(m+1)}$ is orthogonal to $\boldsymbol{r}^{(i)}$ for $i \leq m$, and so we have

$$\dim \operatorname{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^{m+1} = 1 + \dim \operatorname{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^{m} = 1 + \dim K_{m+1}(A, \boldsymbol{r}^{(0)}) \geq \dim K_{m+2}.$$

Combining this with $\operatorname{span}\{\boldsymbol{r}^{(i)}\}_{i=0}^{m+1} \subseteq K_{m+2}$ we get equality of the subspaces.
(2) From Theorem 4.22 of the previous lectures, we know that $\langle \boldsymbol{r}^{(m)}, \boldsymbol{d}^{(i)} \rangle = 0$ for all $i \leq m-1$ which is exactly what we want.
(3) This is by definition, since the $\boldsymbol{d}^{(i)}$ are obtained by A-orthogonalizing the residual vectors. $\square$

The properties above allow us to show that $\langle \boldsymbol{r}^{(k)}, \boldsymbol{d}^{(i)} \rangle_A = 0$ for $i \leq k-2$. Indeed, if we take $i \leq k-2$, then $\boldsymbol{d}^{(i)} \in K_{k-1}(A, \boldsymbol{r}^{(0)})$, which implies $A\boldsymbol{d}^{(i)} \in K_k(A, \boldsymbol{r}^{(0)})$ and so $\langle \boldsymbol{r}^{(k)}, \boldsymbol{d}^{(i)} \rangle_A = \langle \boldsymbol{r}^{(k)}, A\boldsymbol{d}^{(i)} \rangle = 0$.

As a result, the algorithm (4.9) simplifies and reduces to the following: Set $\boldsymbol{d}^{(0)} = \boldsymbol{r}^{(0)} = \boldsymbol{b} - A\boldsymbol{x}^{(0)}$ and iterate, for $k \geq 0$:

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

where $\boldsymbol{r}^{(k)}$ stands for $\boldsymbol{b} - A\boldsymbol{x}^{(k)}$.

We can further simplify the expressions for α_k and β_k using the properties stated in Theorem 4.27. 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:

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

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 4.27(2)), and in (c) we used (4.11).

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

Algorithm 4.28 (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 4.22 in previous lecture) terminates after at most n steps. We restate this result in the special case of the conjugate gradient method.

Corollary 4.29 (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 $(\boldsymbol{r}^{(k)})_{k\geq 0}$ 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 $\operatorname{span}\{A^i\boldsymbol{r}_0\}_{i=0}^{n-1}$, their number is bounded by the dimension of that space. □

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

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

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

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