MATH 5330: Computational Methods of Linear Algebra

Lecture Note 5: The Krylov Subspace Method

Xianyi Zeng
Department of Mathematical Sciences, UTEP

1 The Minimum-Residual Krylov Subspace Methods

We look at more detail into the projection method, where a solution candidate $\tilde{\boldsymbol{x}}$ in the affine space $\boldsymbol{x}_0+\mathcal{K}$ is searched for, so that the residual $\boldsymbol{r}=\boldsymbol{b}-A\boldsymbol{x}$ is orthogonal to a linear subspace $\mathcal{L}$. Let A be any non-singular $n\times n$ matrix (not necessarily symmetric), and let $\{\boldsymbol{v}_1,\cdots,\boldsymbol{v}_m\}$ and $\{\boldsymbol{w}_1,\cdots,\boldsymbol{w}_m\}$ be bases for $\mathcal{K}$ and $\mathcal{L}$, respectively. Denote $V=[\boldsymbol{v}_1\ \boldsymbol{v}_2\ \cdots\ \boldsymbol{v}_m]\in\mathbb{R}^{n\times m}$ and $W=[\boldsymbol{w}_1\ \boldsymbol{w}_2\ \cdots\ \boldsymbol{w}_m]\in\mathbb{R}^{n\times m}$, then we have the first claim that if such a solution $\tilde{\boldsymbol{x}}$ exists, it is given by:

$$\tilde{\boldsymbol{x}}=\boldsymbol{x}_0+V(W^tAV)^{-1}W^t\boldsymbol{r}_0\,. \tag{1.1}$$

To show this, first note that the set $\boldsymbol{x}_0+\mathcal{K}$ is equivalent to:

$$\{\boldsymbol{x}_0+V\boldsymbol{y}\,:\,\boldsymbol{y}\in\mathbb{R}^m\}\,.$$

Thus if we denote $\tilde{\boldsymbol{x}}=\boldsymbol{x}_0+V\tilde{\boldsymbol{y}}$ for some $\tilde{\boldsymbol{y}}\in\mathbb{R}^m$, there is:

$$\tilde{\boldsymbol{r}}\in\mathcal{L}^\perp\quad\Leftrightarrow\quad 0=W^t\tilde{\boldsymbol{r}}=W^t[\boldsymbol{b}-A(\boldsymbol{x}_0+V\tilde{\boldsymbol{y}})]=W^t(\boldsymbol{r}_0-AV\tilde{\boldsymbol{y}})\,, \tag{1.2}$$

or equivalently $\tilde{\boldsymbol{y}}=(W^tAV)^{-1}W^t\boldsymbol{r}_0$[1]; (1.1) follows naturally.

Theorem 1.1. *If A is an arbitrary non-singular $n\times n$ matrix and assume $\mathcal{L}=A\mathcal{K}$, then the $\tilde{\boldsymbol{x}}$ obtained by the previous process minimize the L^2-norm of the residual vector. In addition, let P be the orthogonal projector onto the subspace $A\mathcal{K}$, then the residual $\tilde{r}=\boldsymbol{b}-A\tilde{\boldsymbol{x}}$ satisfies:*

$$\tilde{\boldsymbol{r}}=(I-P)\boldsymbol{r}_0\,. \tag{1.3}$$

Proof. Let $V\in\mathbb{R}^{n\times m}$ be the matrix composed of a basis of $\mathcal{K}$ as before, then we can select a basis of $\mathcal{L}$ as $W=AV$. Define $\phi(\boldsymbol{y})=||\boldsymbol{b}-A(\boldsymbol{x}_0+V\boldsymbol{y})||^2=||\boldsymbol{r}_0-AV\boldsymbol{y}||^2=\boldsymbol{y}^tV^tA^tAV\boldsymbol{y}-2\boldsymbol{y}^tV^tA^t\boldsymbol{r}_0+\boldsymbol{r}_0^t\boldsymbol{r}_0$ then:

$$\nabla\phi(\boldsymbol{y})=V^tA^tAV\boldsymbol{y}-V^tA^t\boldsymbol{r}_0\,.$$

Because A is non-singular and V is full-rank, we have V^tA^tAV is non-singular and thusly symmetric positive-definite. Hence $\tilde{\boldsymbol{y}}$ minimize $\phi(\boldsymbol{y})$ if and only if $\nabla\phi(\tilde{\boldsymbol{y}})=0$ or equivalently:

$$\tilde{\boldsymbol{y}}=(V^tA^tAV)^{-1}V^tA^t\boldsymbol{r}_0=(W^tAV)^{-1}W^t\boldsymbol{r}_0\,,$$

the same as the solution to the projection method before.

[1] That W^tAV is invertible comes classical results in linear algebra; for example we can use the Sylvester's rank inequality to show that rankW^tAm, and then use the fact that the rank of a matrix equals the largest order of any non-zero minor to show that rank$(W^tA)V=m$.

For the second part, we note that:

$$\tilde{r} = \boldsymbol{r}_0 - AV\tilde{\boldsymbol{y}} = (I - AV(V^tA^tAV)^{-1}V^tA^t)\tilde{\boldsymbol{r}}_0\,,$$

where $P = AV(V^tA^tAV)^{-1}V^tA^t = W(W^tW)^{-1}W^t$ is exactly the orthogonal projector on to $\mathcal{L} = A\mathcal{K}$ (see the exercise from the last lecture). □

Usually we hope that if the true solution is in the affine space $\boldsymbol{x}_0 + \mathcal{K}$, the projection method will provides the exact solution, i.e. $\tilde{\boldsymbol{r}} = 0$. Recalling that in the case of conjugate gradient method, when the solution is in the affine space $\boldsymbol{x}_0 + \mathcal{K}_k$ we have $A^{k+1}\boldsymbol{r}_0 \in \mathcal{K}_k$ or $\mathcal{K}_k = A\mathcal{K}_k$. This turns out to be a very general situation, as seen in Exercise 1.

A *Krylov subspace method* (or Krylov method for simplicity) is a special case of the project methods such that:

$$\mathcal{K}_m = \text{span}(\boldsymbol{r}_0, A\boldsymbol{r}_0, \cdots, A^{m-1}\boldsymbol{r}_0)^2\,, \tag{1.4}$$

where $\boldsymbol{r}_0 = \boldsymbol{b} - A\boldsymbol{x}_0$. Two widely used versions of the Krylov methods are (1) $\mathcal{L}_m = A\mathcal{K}_m$ and (2) $\mathcal{L}_m = A^t\mathcal{K}_m$. Due to Theorem 1.1, the first variation is also called the *minimum-residual Krylov subspace methods*. We shall discuss the minimum-residual Krylov methods in more detail in this lecture, whereas the other one in the next lecture.

2 Arnoldi's Algorithm

Using the basis (1.4) is generally not computational efficient, especially when m is large. A common strategy is to construct a basis $\{\boldsymbol{v}_1, \boldsymbol{v}_2, \cdots, \boldsymbol{v}_m\}$ of $\mathcal{K}_m$ gradually, so that the basis is composed of mutually orthogonal unit vectors. For example, in the conjugate gradient method it turns out that $\mathcal{K}_m = \text{span}(\boldsymbol{r}_0, \cdots, \boldsymbol{r}_{m-1})$ and all the residual vectors are orthogonal to each other.

The method that constructs the basis $V_m = [\boldsymbol{v}_1\ \boldsymbol{v}_2\ \cdots\ \boldsymbol{v}_m]$ from (1.4) is called the *Arnoldi's algorithm*. This is in essence an orthonormalization problem that will be discussed in more detail in later lectures. Here we just adopt one of the simplest method for this purpose (the Gram-Schmidt process), and other variants will be discussed when the orthogonal reduction is formally introduced. The Gram-Schmidt-based Arnoldi's algorithm is given below (we suppose $\boldsymbol{r}_0 \neq 0$):

$$\begin{aligned}
&\boldsymbol{v}_1 = \boldsymbol{r}_0/||\boldsymbol{r}_0||\,; \\
&\text{for } j = 1, \cdots, m: \\
&\quad \boldsymbol{w}_j = A\boldsymbol{v}_j\,; \\
&\quad \text{for } i = 1, \cdots, j: \\
&\quad\quad h_{ij} = \boldsymbol{w}_j \cdot \boldsymbol{v}_i\,; \\
&\quad\quad \boldsymbol{w}_j = \boldsymbol{w}_j - h_{ij}\boldsymbol{v}_i\,; \\
&\quad \text{end} \\
&\quad h_{j+1,j} = ||\boldsymbol{w}_j||\,; \\
&\quad \text{if } h_{j+1,j} = 0,\ \text{stop}\,; \\
&\quad \boldsymbol{v}_{j+1} = \boldsymbol{w}_j/h_{j+1,j}\,; \\
&\text{end}
\end{aligned} \tag{2.1}$$

[2] Note that we change the last power comparing to the conjugate gradient method.

We can derive a matrix representation for the previous algorithm:

$$A\boldsymbol{v}_j = \sum_{i=1}^{j+1} h_{ij}\boldsymbol{v}_i \quad \Rightarrow \quad h_{j+1,j}\boldsymbol{v}_{j+1} = A\boldsymbol{v}_j - \sum_{i=1}^{j} h_{ij}\boldsymbol{v}_i\,, \qquad 0 \le j \le m\,.$$

This is equivalent to:

$$AV_m = V_m H_m + \boldsymbol{w}_m \boldsymbol{e}_m^t = V_{m+1}\overline{H}_m\,, \tag{2.2}$$

where $H_m \in \mathbb{R}^{m\times m}$ is an Hessenberg matrix:

$$H_m = \begin{bmatrix} h_{11} & h_{12} & h_{13} & \cdots & h_{1,m-2} & h_{1,m-1} & h_{1m} \\ h_{21} & h_{22} & h_{23} & \cdots & h_{2,m-2} & h_{2,m-1} & h_{2m} \\ 0 & h_{32} & h_{33} & \cdots & h_{3,m-2} & h_{3,m-1} & h_{3m} \\ 0 & 0 & h_{43} & \cdots & h_{4,m-2} & h_{4,m-1} & h_{4m} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & h_{m,m-1} & h_{mm} \end{bmatrix}, \tag{2.3}$$

$\boldsymbol{e}_m$ is the last unit vector of $\mathbb{R}^m$, and $\overline{H}_m \in \mathbb{R}^{(m+1)\times m}$ is given by $\overline{H}_m = \begin{bmatrix} H_m \\ h_{m+1,m}\boldsymbol{e}_m^t \end{bmatrix}$. Furthermore, because $V_m^t V_m = I \in \mathbb{R}^{m\times m}$ and $V_m^t \boldsymbol{w}_m = 0$, we have:

$$V_m^t A V_m = H_m\,. \tag{2.4}$$

Finally, let us look at the stopping condition for the Arnoldi's algorithm: The method stops at the j-th iteration if and only if $h_{j+1,j} = 0$ and in this case $\boldsymbol{w}_j = 0$. Hence $A\boldsymbol{v}_j = \sum_{i=1}^{j} h_{ij}\boldsymbol{v}_j \in \mathcal{K}_j$, i.e., $\mathcal{K}_j$ is invariant under A. Later on we will see that in the GMRES algorithm, this condition indicates that the residual is zero after the j-th iteration.

3 GMRES

The generalized minimum residual method (GMRES) is a Krylov subspace method where $\mathcal{L}_m = A\mathcal{K}_m$ and $\boldsymbol{v}_1 = \boldsymbol{r}_0/||\boldsymbol{r}_0||$, so that an orthonormal unit basis for $\mathcal{K}_m$ can be constructed according to the Arnoldi's algorithm. From Theorem 1.1, the iteration $\boldsymbol{x}_m$ is found in the affine space $\boldsymbol{x}_0 + \mathcal{K}_m$ so that the L^2-norm of the residual $||\boldsymbol{b} - A\boldsymbol{x}||$ is minimized. Writing $\boldsymbol{x} = \boldsymbol{x}_0 + V_m\boldsymbol{y}$, $\boldsymbol{y} \in \mathbb{R}^m$ we have:

$$\boldsymbol{r} = \boldsymbol{b} - A\boldsymbol{x} = \boldsymbol{r}_0 - AV_m\boldsymbol{y} = \beta\boldsymbol{v}_1 - V_{m+1}\overline{H}_m\boldsymbol{y} = V_{m+1}(\beta\boldsymbol{e}_1 - \overline{H}_m\boldsymbol{y})\,, \tag{3.1}$$

where $\beta = ||\boldsymbol{r}_0||$ and $\boldsymbol{e}_1$ is the first unit basis vector of $\mathbb{R}^{m+1}$. Clearly, $||\boldsymbol{r}|| = \big|\big|\beta\boldsymbol{e}_1 - \overline{H}_m\boldsymbol{y}\big|\big|$, thus $\boldsymbol{x}_m = \boldsymbol{x}_0 + V_m\boldsymbol{y}_m$ can be constructed by solving the smaller least-squares problem:

$$\boldsymbol{y}_m = \arg\min_{\boldsymbol{y}\in\mathbb{R}^m} \big|\big|\beta\boldsymbol{e}_1 - \overline{H}_m\boldsymbol{y}\big|\big|\,. \tag{3.2}$$

To this end, we convert the problem of solving a linear system into least-squares problems. Thus whenever a robust least-squares solver is available, we can combine it with the Arnoldi's algorithm before to construct a GMRES solver for linear systems.

Especially, by exploring the fact that $\overline{H}_m$ is an Hessenberg matrix one can construct a very efficient least-squares solver for (3.2) that is based on Householder transformation. The Householder

transformation is very robust because it involves only multiplication rather than divisions, c.f. in the Gaussian elimination process we inevitably need to divide by the numbers $a_{kk}^{(k)}$ that could be very small and cannot be predicted from the initial data.

Concerning the robustness, one may observe that in the Arnoldi's algorithm (2.1), we may also run into the issue with $h_{j+1,j} \neq 0$ is very small; it turns out that this issue can also be resolved by incorporating the Householder transformation in the construction of a orthonormal unit basis for $\mathcal{K}_m$. The Householder transformation as well as the enhanced GMRES will be discussed in more detail in the later lecture on the least-squares problems.

Lastly, we look at the breakdown of the GMRES algorithm. Particularly, if $h_{m+1,m} = 0$ at the m-th iteration, as discussed in the Arnoldi's algorithm there is $\mathcal{K}_m = A\mathcal{K}_m$ and the iterations stuck at $\boldsymbol{x}_m$. The important question we want to answer is that in this case, $\boldsymbol{x}_m$ is the exact solution, or $\boldsymbol{r}_m = 0$. This is left as Exercise 2.

4 Practical Variants of GMRES

Other than enhancing the robustness of the GMRES method with Householder transformation, people also concerns about other practical issues. For example, larger m will demand more memory storage for the basis of $\mathcal{K}_m$, i.e., V_m. A first effort to improve the memory usage results in the so-called restarted GMRES, in which case people set a small $0 < m_0 \ll n$, and whenever the m_0-th iteration is reached, one set $\boldsymbol{x}_0 = \boldsymbol{x}_{m_0}$ and start from scratch the GMRES algorithm with this new "initial guess". The major argument against the restarted GMRES is that: Despite the fact that GMRES will converge in no more than n iterations, the restarted GMRES may fail to converge. Thus people need to find a balance between choosing the number m_0, where a smaller m_0 requires less memory but increases the chance of divergent solutions.

A more reliable method that works much better in practice is the direct quasi-GMRES method (DQGMRES) or truncated GMRES. In this case, people still needs to select a small number $m_0 > 0$; but instead of restart after m_0 iterations, DQGMRES will keep up to m_0 most recent orthonormal basis vectors of V_m. The major benefit of doing so is due to a theorem of Nachtigal, that if V_{m+1}, the Arnoldi's basis associated with the DQGMRES is of full rank (note that only the adjacent m_0 column vectors of V_{m+1} are guaranteed to be orthogonal to each other), then we have:

$$\left|\left|\boldsymbol{r}_m^Q\right|\right| \leq \kappa(V_{m+1})\left|\left|\boldsymbol{r}_m^G\right|\right|,$$

where $\boldsymbol{r}_m^Q$ is the residual of DQGMRES at the m-th iteration, $\boldsymbol{r}_m^G$ is that of the full GMRES, and $\kappa()$ is the condition number computed using the induced L^2-norm for matrices (pseudo-inverse is used here for non-square matrices).

Exercises

Exercise 1. *Under the assumption of Theorem 1.1, suppose $\boldsymbol{x}_0 = 0$, $\boldsymbol{b} \in \mathcal{K}$, and $\mathcal{K} = A\mathcal{K}$, then $\tilde{\boldsymbol{x}}$ is the exact solution to $A\boldsymbol{x} = \boldsymbol{b}$.*
***Hint**: That is, to show $\tilde{\boldsymbol{r}} = 0$ or $\tilde{\boldsymbol{r}}_0 \in A\mathcal{K}$.*

Exercise 2. *Suppose H_m (and hence also $\overline{H}_m$, with possible $h_{m+1,m}=0$) is constructed, then by $V_m^t A V_m = H_m$ we know that H_m is non-singular. Define $\tilde{\boldsymbol{y}}_m = H_m^{-1}(\beta \boldsymbol{e}_1)$ and let $\tilde{\boldsymbol{x}}_m = \boldsymbol{x}_0 + V_m \tilde{\boldsymbol{y}}_m$ and $\tilde{\boldsymbol{r}}_m = \boldsymbol{b} - A\tilde{\boldsymbol{x}}_m$. Show that:*

$$||\tilde{\boldsymbol{r}}_m|| = h_{m+1,m} \left| \boldsymbol{e}_m^t \tilde{\boldsymbol{y}}_m \right|,$$

and deduce that if $h_{m+1,m}=0$ for some m, the solution $\boldsymbol{x}_m$ to the GMRES method is exact (i.e., $\boldsymbol{r}_m = 0$).

***Hint**: For the second part, use the minimum-residual property of $\boldsymbol{x}_m$.*