AMSC 600 /CMSC 760 Advanced Linear Numerical Analysis

Fall 2007

Arnoldi Methods

Dianne P. O'Leary

Algorithms that use the Arnoldi Basis

Reference: Chapter 6 of Saad

- The Arnoldi Basis
- How to use the Arnoldi Basis
- FOM (usually called the Arnoldi iteration)
- GMRES
- Givens Rotations
- Relation between FOM and GMRES iterates
- Practicalities
- Convergence Results for GMRES

The Arnoldi Basis

Recall:

- The Arnoldi basis $\mathbf{v}_1, \ldots, \mathbf{v}_n$ is obtained column by column from the relation
$$\mathbf{AV} = \mathbf{VH}$$
where $\mathbf{H}$ is upper Hessenberg and $\mathbf{V}^T\mathbf{V} = \mathbf{I}$.
- In particular, $h_{ij} = (\mathbf{A}\mathbf{v}_j, \mathbf{v}_i)$.
- The vectors $\mathbf{v}_j$ can be formed by a (modified) Gram Schmidt process.

For three algorithms for computing the recurrence, see Sec 6.3.1 and 6.3.2.

Facts about the Arnoldi Basis

Let's convince ourselves of the following facts:

- The first m columns of $\mathbf{V}$ form an orthogonal basis for $\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$.
- Let $\mathbf{V}_m$ denote the first m columns of $\mathbf{V}$, and $\mathbf{H}_m$ denote the leading $m \times m$ block of $\mathbf{H}$. Then

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m\mathbf{H}_m + h_{m+1,m}\mathbf{v}_{m+1}\mathbf{e}_m^T.$$

- The algorithm breaks down if $h_{m+1,m} = 0$. In this case, we cannot get a new vector $\mathbf{v}_{m+1}$.
- This breakdown means that $\mathcal{K}_{m+1}(\mathbf{A}, \mathbf{v}_1) = \mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$. In other words, $\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$ is an invariant subspace of $\mathbf{A}$ of dimension m.

How to use the Arnoldi Basis

We want to solve $\mathbf{A}\mathbf{x} = \mathbf{b}$.

We choose $\mathbf{x}_0 = \mathbf{0}$ and $\mathbf{v}_1 = \mathbf{b}/\beta$, where $\beta = \|\mathbf{b}\|$.

Let's choose $\mathbf{x}_m \in \mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$.

Two choices:

- Projection:
  Make the residual orthogonal to $\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$.
  (Section 6.4.1) This is the Full Orthogonalization Method (FOM).
- Minimization:
  Minimize the norm of the residual over all choices of $\mathbf{x}$ in $\mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$.
  (Section 6.5) This is the Generalized Minimum Residual Method (GMRES).

FOM (usually called the Arnoldi iteration)

The Arnoldi relation:

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m\mathbf{H}_m + h_{m+1,m}\mathbf{v}_{m+1}\mathbf{e}_m^T.$$

If $\mathbf{x}_m \in \mathcal{K}_m(\mathbf{A}, \mathbf{v}_1)$, then we can express it as $\mathbf{x}_m = \mathbf{V}_m\mathbf{y}_m$ for some vector $\mathbf{y}_m$.

The residual is

$$\begin{aligned}
\mathbf{r}_m &= \mathbf{b} - \mathbf{A}\mathbf{x}_m \\
&= \beta\mathbf{v}_1 - \mathbf{A}\mathbf{V}_m\mathbf{y}_m \\
&= \beta\mathbf{v}_1 - (\mathbf{V}_m\mathbf{H}_m + h_{m+1,m}\mathbf{v}_{m+1}\mathbf{e}_m^T)\mathbf{y}_m.
\end{aligned}$$

To enforce orthogonality, we want

$$\begin{aligned}\mathbf{0} &= \mathbf{V}_m^T \mathbf{r}_m \\ &= \beta \mathbf{V}_m^T \mathbf{v}_1 - (\mathbf{V}_m^T \mathbf{V}_m \mathbf{H}_m + h_{m+1,m} \mathbf{V}_m^T \mathbf{v}_{m+1} \mathbf{e}_m^T) \mathbf{y}_m. \\ &= \beta \mathbf{e}_1 - \mathbf{H}_m \mathbf{y}_m.\end{aligned}$$

Therefore, given the Arnoldi relation at step m, we solve the linear system

$$H_m \mathbf{y}_m = \beta \mathbf{e}_1$$

and set

$$\begin{aligned}\mathbf{x}_m &= \mathbf{V}_m \mathbf{y}_m, \\ \mathbf{r}_m &= \mathbf{b} - \mathbf{A}\mathbf{x}_m \\ &= \mathbf{b} - \mathbf{A}\mathbf{V}_m \mathbf{y}_m \\ &= \beta \mathbf{V}_m \mathbf{e}_1 - (\mathbf{V}_m \mathbf{H}_m + h_{m+1,m} \mathbf{v}_{m+1} \mathbf{e}_m^T) \mathbf{y}_m \\ &= \mathbf{V}_m (\beta \mathbf{e}_1 - \mathbf{H}_m \mathbf{y}_m) - h_{m+1,m} \mathbf{v}_{m+1} \mathbf{e}_m^T \mathbf{y}_m \\ &= -h_{m+1,m} \mathbf{v}_{m+1} \mathbf{e}_m^T \mathbf{y}_m \\ &= -h_{m+1,m} (\mathbf{y}_m)_m \mathbf{v}_{m+1}.\end{aligned}$$

Therefore, the norm of the residual is just $|h_{m+1,m}(\mathbf{y}_m)_m|$, which can be computed without forming $\mathbf{x}_m$.

GMRES

The Arnoldi relation:

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m \mathbf{H}_m + h_{m+1,m} \mathbf{v}_{m+1} \mathbf{e}_m^T.$$

The iterate: $\mathbf{x}_m = \mathbf{V}_m \mathbf{y}_m$ for some vector $\mathbf{y}_m$.

Define $\bar{\mathbf{H}_m}$ to be the leading $(m+1) \times m$ block of $\mathbf{H}$.

The residual:

$$\begin{aligned}\mathbf{r}_m &= \beta \mathbf{v}_1 - (\mathbf{V}_m \mathbf{H}_m + h_{m+1,m} \mathbf{v}_{m+1} \mathbf{e}_m^T) \mathbf{y}_m \\ &= \beta \mathbf{v}_1 - \mathbf{V}_{m+1} \bar{\mathbf{H}_m} \mathbf{y}_m \\ &= \mathbf{V}_{m+1} (\beta \mathbf{e}_1 - \bar{\mathbf{H}_m} \mathbf{y}_m).\end{aligned}$$

The norm of the residual:

$$\|\mathbf{r}_m\| = \|\mathbf{V}_{m+1}(\beta \mathbf{e}_1 - \bar{\mathbf{H}_m} \mathbf{y}_m)\| = \|\beta \mathbf{e}_1 - \bar{\mathbf{H}_m} \mathbf{y}_m\|.$$

We minimize this quantity w.r.t. $\mathbf{y}_m$ to get the GMRES iterate.

Givens Rotations

The linear system involving $\mathbf{H}_m$ is solved using m Givens rotations.

The least squares problems involving $\bar{\mathbf{H}}_m$ is solved using 1 additional Givens rotation.

Relation between FOM and GMRES iterates

FOM: $\mathbf{H}_m \mathbf{y}_m^F = \beta \mathbf{e}_1$, or

$$\mathbf{H}_m^T \mathbf{H}_m \mathbf{y}_m^F = \beta \mathbf{H}_m^T \mathbf{e}_1.$$

GMRES: minimize $\|\beta \mathbf{e}_1 - \bar{\mathbf{H}}_m \mathbf{y}_m^G\|$ by solving the normal equations

$$\bar{\mathbf{H}_m}^T \bar{\mathbf{H}}_m \mathbf{y}_m^G = \beta \bar{\mathbf{H}_m}^T \mathbf{e}_1.$$

Note that

$$\bar{\mathbf{H}}_m = \begin{bmatrix} \mathbf{H}_m \\ h_{m+1,m} \mathbf{e}_m^T \end{bmatrix},$$

so $\mathbf{H}_m^T \mathbf{H}_m$ and $\bar{\mathbf{H}_m}^T \bar{\mathbf{H}}_m$ differ by a matrix of rank 2.

Therefore (Sec 6.5.7 and Sec 6.5.8),

- If we compute the FOM iterate, we can easily get the GMRES iterate, and vice versa.
- It can be shown that the smallest residual computed by FOM in iterations $1, \ldots, m$ is within a factor of $\sqrt{m}$ of the smallest for GMRES.
- There is a recursion for computing the GMRES iterates given the FOM iterates:

  $$\mathbf{x}_m^G = s_m^2 \mathbf{x}_{m-1}^G - c_m^2 \mathbf{x}_m^F,$$

  where $s_m^2 + c_m^2 = 1$ and c_m is the cosine used in the last step of the reduction of $\bar{\mathbf{H}}_m$ to upper triangular form.

Practicalities

Practicality 1: GMRES and FOM can stagnate.

It is possible (but unlikely) that GMRES and FOM make no progress at all until the last iteration.

Example: Let

$$\mathbf{A} = \begin{bmatrix} \mathbf{e}^T & 1 \\ \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad \mathbf{x}_{true} = \mathbf{e}_n, \quad \mathbf{b} = \mathbf{e}_1.$$

($\mathbf{e}$ is the vector with every entry equal to 1.) Then
$\mathcal{K}_m(A, b) = span\{\mathbf{e}_1, \ldots, \mathbf{e}_m\}$, so the solution vector is orthogonal to $\mathcal{K}_m(A, b)$ for $m < n$.

This is called complete stagnation. Partial stagnation can also occur.

The usual remedy is preconditioning (discussed later).

Practicality 2: m must be kept relatively small.

As m gets big,

- The work per iteration becomes overwhelming, since the upper Hessenberg part of $\mathbf{H}_m$ is generally dense. In fact, the work for the orthogonalization grows as $O(m^2 n)$
- The orthogonality relations are increasingly contaminated by round-off error, and this can cause the computation to compute bad approximations to the true GMRES/FOM iterates.

Two remedies:

- restarting (Sec 6.4.1 and Sec 6.5.5)
- truncating the orthogonalization (Sec 6.4.2 and Sec 6.5.6)

Restarting

In the restarting algorithm, we run the Arnoldi iteration for a fixed number of steps, perhaps $m = 100$.

If the resulting solution is not good enough, we repeat, by setting v_1 to be the normalized residual.

This limits the work per iteration, but removes the finite-termination property.

In fact, due to stagnation, it is possible (but unlikely) that the iteration never makes any progress at all!

Truncating the orthogonalization

In this algorithm, we modify the Arnoldi iteration so that we only orthogonalize against the latest k vectors, rather than orthogonalizing against all of the old vectors.

The matrix $\bar{\mathbf{H}}_m^{approx}$ then has only $k-1$ nonzeros above the main diagonal, and the work for orthogonalization is reduced from $O(m^2 n)$ to $O(kmn)$.

We are still working with the same Krylov subspace, but

- The basis vectors $\mathbf{v}_1, \ldots, \mathbf{v}_m$ are not exactly orthogonal.
- We are approximating the matrix $\bar{\mathbf{H}}_m$, so we don't minimize the residual (exactly) or make the residual (exactly) orthogonal to the Krylov subspace.

All is not lost, though.

In Section 6.5.6, Saad works out the relation between the GMRES residual norm and the one produced from the truncated method (quasi-GMRES).

Note:

- Quasi-GMRES is not often used.
- Restarted-GMRES is the standard algorithm.

Practicality 3: Breakdown of the Arnoldi iteration is a good thing.

If the Arnoldi iteration breaks down, with $h_{m+1,m} = 0$, so that there is no new vector $\mathbf{v}_{m+1}$, then

- $\mathbf{H}_m = \bar{\mathbf{H}}_m$, so both GMRES and FOM compute the same iterate.
- More importantly, Arnoldi breaks down if and only if $\mathbf{x}_m^G = \mathbf{x}_m^F = \mathbf{x}_{true}$. This follows from the fact that the norm of the GMRES residual is $|h_{m+1,m}(\mathbf{y}_m)_m|$.

Practicality 4: These algorithms also work in complex arithmetic.

If $\mathbf{A}$ or $\mathbf{b}$ is complex rather than real, the algorithms still work, as long as we remember to

- use conjugate transpose instead of transpose.

- use complex Givens rotations (Section 6.5.9):

$$\begin{bmatrix} \bar{c} & \bar{s} \\ -s & c \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \gamma \\ 0 \end{bmatrix}$$

  if

$$s = \frac{y}{\sqrt{|x|^2 + y^2}}, \quad c = \frac{x}{\sqrt{|x|^2 + y^2}}.$$

Convergence Results for GMRES

(Sec 6.11.4) Two definitions:

- We say that a matrix $\mathbf{A}$ is positive definite if $\mathbf{x}^T\mathbf{A}\mathbf{x} > 0$. for all $\mathbf{x} \neq \mathbf{0}$.
  - If $\mathbf{A}$ is symmetric, this means that all of its eigenvalues are positive.
  - For a nonsymmetric matrix, this means that the symmetric part of $\mathbf{A}$, which is $\frac{1}{2}(\mathbf{A} + \mathbf{A}^T)$, is positive definite.
- The algorithm that restarts GMRES every m iterations is called GMRES(m).

GMRES(m) converges for positive definite matrices.

Proof: The iteration can't stagnate, because it makes progress at the first iteration. In particular, if we seek $\mathbf{x} = \alpha\mathbf{v}_1$ to minimize

$$\begin{aligned} \|\mathbf{r}\|^2 &= \|\mathbf{b} - \alpha\mathbf{A}\mathbf{v}_1\|^2 \\ &= \mathbf{b}^T\mathbf{b} - 2\alpha\mathbf{b}^T\mathbf{A}\mathbf{v}_1 + \alpha^2\mathbf{v}_1^T\mathbf{A}^T\mathbf{A}\mathbf{v}_1, \end{aligned}$$

where $\mathbf{b} = \beta\mathbf{v}_1$ and $\beta > 0$ then

$$\alpha = \frac{\mathbf{b}^T\mathbf{A}\mathbf{v}_1}{\mathbf{v}_1^T\mathbf{A}^T\mathbf{A}\mathbf{v}_1} > 0,$$

so the residual norm is reduced.

By working a little harder, we can show that it reduced enough to guarantee convergence. []

GMRES(m) is an optimal polynomial method.

- The iterates take the form

$$\mathbf{x}_m = \mathbf{x}_0 + q_{m-1}(\mathbf{A})\mathbf{r}_0,$$

  where q_{m-1} is a polynomial of degree at most $m - 1$.

- Saying that $\mathbf{x}$ is in $\mathcal{K}_m(\mathbf{A}, \mathbf{r}_0)$ is equivalent to saying $\mathbf{x} = q_{m-1}(\mathbf{A})\mathbf{r}_0$ for some polynomial of degree at most $m-1$.
- The resulting residual is

$$\begin{aligned} \mathbf{r}_m &= \mathbf{r}_0 - q_{m-1}(\mathbf{A})\mathbf{A}\mathbf{r}_0 \\ &= (\mathbf{I} - q_{m-1}(\mathbf{A})\mathbf{A})\mathbf{r}_0 \\ &= p_m(\mathbf{A})\mathbf{r}_0, \end{aligned}$$

  where p_m is a polynomial of degree m satisfying $p_m(0) = 1$.
- since GMRES(m) minimizes the norm of the residual, it must pick the optimal polynomial p_m.

- Let $\mathbf{A} = \mathbf{X}\mathbf{\Lambda}\mathbf{X}^{-1}$, where $\mathbf{\Lambda}$ is diagonal and contains the eigenvalues of $\mathbf{A}$. (In other words, we are assuming that $\mathbf{A}$ has a complete set of linearly independent eigenvectors, the columns of $\mathbf{X}$.) Then

$$\begin{aligned} \|\mathbf{r}_m\| &= \min_{p_m(0)=1} \|p_m(\mathbf{A})\mathbf{r}_0\| \\ &= \min_{p_m(0)=1} \|\mathbf{X}p_m(\mathbf{\Lambda})\mathbf{X}^{-1}\mathbf{r}_0\| \\ &\leq \|\mathbf{X}\|\ \|\mathbf{X}^{-1}\| \min_{p_m(0)=1} \|p_m(\mathbf{\Lambda})\|\|\mathbf{r}_0\| \\ &\leq \|\mathbf{X}\|\ \|\mathbf{X}^{-1}\| \min_{p_m(0)=1} \max_{j=1,\ldots,n} |p_m(\lambda_j)|\|\mathbf{r}_0\|. \end{aligned}$$

  Therefore, any polynomial that is small on the eigenvalues can give us a bound on the GMRES(m) residual norm, since GMRES(m) picks the optimal polynomial.

- By choosing a Chebyshev polynomial, Saad derives the bound

$$\|\mathbf{r}_m\| \leq \left(\frac{a + \sqrt{a^2 - d^2}}{c + \sqrt{c^2 - d^2}} \right)^k \|\mathbf{X}\|\ \|\mathbf{X}^{-1}\|\ \|\mathbf{r}_0\|.$$

  whenever the ellipse centered at c with focal distance d and semi-major axis a contains all of the eigenvalues and excludes the origin (See the picture on p. 207).