

14 Arnoldi Iteration and GMRES

14.1 Arnoldi Iteration

The classical iterative solvers we have discussed up to this point were of the form

$$\mathbf{x}^{(k)} = G\mathbf{x}^{(k-1)} + \mathbf{c}$$

with constant G and \mathbf{c} . Such methods are also known as *stationary methods*. We will now study a different class of iterative solvers *based on optimization*.

All methods will require a “black box” implementation of a matrix-vector product. In most library implementations of such solvers the user can therefore provide a custom function which computes the matrix-vector product as efficiently as possible for the specific system matrix at hand.

One of the main ingredients in all of the following methods are *Krylov subspaces*. Given $A \in \mathbb{C}^{m \times m}$ and $\mathbf{b} \in \mathbb{C}^m$ one generates

$$\{\mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, A^3\mathbf{b}, \dots\},$$

which is referred to as a *Krylov sequence*. Clearly, (fast) matrix-vector products play a crucial role in generating this sequence since each subsequent vector in the sequence is obtained from the previous one by multiplication by A .

With the Krylov sequence at hand one defines

$$\mathcal{K}_n = \text{span}\{\mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, \dots, A^{n-1}\mathbf{b}\}$$

as the n -th order *Krylov subspace*.

The Arnoldi iteration method to be derived will be applicable to both linear systems and eigenvalue problems, and therefore we are interested in re-examining similarity transformations of the form

$$A = QHQ^*,$$

where H is an upper Hessenberg matrix.

In our earlier work we used Householder reflectors to transform A to upper Hessenberg form. This had its advantages since the resulting algorithm is a stable one. When studying the QR factorization we also looked at the modified Gram-Schmidt algorithm. That algorithm was less stable. However, it has the advantage that one gets one column of the unitary matrix Q one column at a time, i.e., the modified Gram-Schmidt algorithm can be stopped at any time and yields a partial set of orthonormal column vectors. On the other hand, with Householder reflectors we always have to perform the entire QR factorization before we get (all) orthonormal vectors.

Thus, Arnoldi iteration can be seen as the use of the modified Gram-Schmidt algorithm in the context of Hessenberg reduction.

14.2 Derivation of Arnoldi Iteration

We start with the similarity transformation $A = QHQ^*$ with $m \times m$ matrices A , Q , and H . Clearly, this is equivalent to

$$AQ = QH.$$

Now we take $n < m$, so that the eigenvalue equations above can be written as

$$A[\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n, \mathbf{q}_{n+1}, \dots, \mathbf{q}_m] = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n, \mathbf{q}_{n+1}, \dots, \mathbf{q}_m] \times$$

$$\times \begin{bmatrix} h_{11} & h_{12} & \dots & & h_{1n} & \dots & h_{1m} \\ h_{21} & h_{22} & \dots & & h_{2n} & \dots & h_{2m} \\ 0 & h_{32} & h_{33} & & h_{3n} & \dots & h_{3m} \\ & 0 & h_{43} & \ddots & & & \\ & & 0 & \ddots & h_{n-1, n-2} & \vdots & \vdots \\ \vdots & & & \ddots & h_{n, n-1} & h_{nn} & \\ & & & & 0 & h_{n+1, n} & \\ & & & & & 0 & \ddots & \ddots \\ 0 & & \dots & & & & 0 & h_{m, m-1} & h_{mm} \end{bmatrix}.$$

Next, we consider only part of this system. Namely, we let

$$Q_n = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n]$$

$$Q_{n+1} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n, \mathbf{q}_{n+1}]$$

$$\tilde{H}_n = \begin{bmatrix} h_{11} & h_{12} & \dots & & h_{1n} \\ h_{21} & h_{22} & \dots & & h_{2n} \\ 0 & h_{32} & h_{33} & & h_{3n} \\ & 0 & h_{43} & \ddots & \\ & & 0 & \ddots & h_{n-1, n-2} & \vdots \\ \vdots & & & \ddots & h_{n, n-1} & h_{nn} \\ & & & & 0 & h_{n+1, n} \end{bmatrix}$$

and then take

$$AQ_n = Q_{n+1}\tilde{H}_n.$$

Note that here $A \in \mathbb{C}^{m \times m}$, $Q_n \in \mathbb{C}^{m \times n}$, $Q_{n+1} \in \mathbb{C}^{m \times n+1}$, and $\tilde{H}_n \in \mathbb{C}^{n+1 \times n}$ so that both sides of the equation result in an $m \times n$ matrix.

If we compare the n -th columns on both sides, then we get

$$A\mathbf{q}_n = h_{1n}\mathbf{q}_1 + h_{2n}\mathbf{q}_2 + \dots + h_{nn}\mathbf{q}_n + h_{n+1, n}\mathbf{q}_{n+1} \quad (40)$$

which constitutes an $(n+1)$ term recursion for the vector \mathbf{q}_{n+1} . Equation (40) can be re-written as

$$\mathbf{q}_{n+1} = \frac{A\mathbf{q}_n - \sum_{i=1}^n h_{i, n}\mathbf{q}_i}{h_{n+1, n}}. \quad (41)$$

The recursive computation of the columns of the unitary matrix Q in this manner is known as *Arnoldi iteration*.

Example The first step of Arnoldi iteration proceeds as follows. We start with the matrix A and an arbitrary normalized vector \mathbf{q}_1 . Then, according to (41),

$$\mathbf{q}_2 = \frac{A\mathbf{q}_1 - h_{11}\mathbf{q}_1}{h_{21}}.$$

Note that this step involves the matrix-vector product $A\mathbf{q}_1$ (which has to be computed efficiently with a problem specific subroutine).

Since we want $\mathbf{q}_1^*\mathbf{q}_2 = 0$ in order to have orthogonality of the columns of Q we get

$$0 = \mathbf{q}_1^*A\mathbf{q}_1 - h_{11}\mathbf{q}_1^*\mathbf{q}_1.$$

Taking advantage of the normalization of \mathbf{q}_1 this is equivalent to

$$h_{11} = \mathbf{q}_1^*A\mathbf{q}_1$$

– a Rayleigh quotient.

Finally, we let $\mathbf{v} = A\mathbf{q}_1 - h_{11}\mathbf{q}_1$, compute $h_{21} = \|\mathbf{v}\|$, and normalize

$$\mathbf{q}_2 = \frac{\mathbf{v}}{h_{21}}.$$

If we compare (15) with the formula

$$\mathbf{q}_n = \frac{\mathbf{a}_n - \sum_{i=1}^{n-1} r_{i,n}\mathbf{q}_i}{r_{nn}}$$

that we used earlier for the Gram-Schmidt method (cf. (17)), then it is clear that an algorithm for Arnoldi iteration will be similar to that for the modified Gram-Schmidt algorithm:

Algorithm (Arnoldi Iteration)

Let \mathbf{b} be an arbitrary initial vector

$$\mathbf{q}_1 = \mathbf{b}/\|\mathbf{b}\|_2$$

for $n = 1, 2, 3, \dots$

$$\mathbf{v} = A\mathbf{q}_n$$

for $j = 1 : n$

$$h_{jn} = \mathbf{q}_j^*\mathbf{v}$$

$$\mathbf{v} = \mathbf{v} - h_{jn}\mathbf{q}_j$$

end

$$h_{n+1,n} = \|\mathbf{v}\|_2$$

$$\mathbf{q}_{n+1} = \mathbf{v}/h_{n+1,n}$$

end

Remark The most expensive operation in the algorithm is the matrix-vector product $A\mathbf{q}_n$. The rest of the operations are on the order of $\mathcal{O}(mn)$ (so they get a little more expensive in each iteration). Therefore, in addition to the basic Arnoldi algorithm we need an efficient implementation of the matrix-vector product this should be tailored to the problem. Moreover, the algorithm above treats the matrix-vector product as a “black box” and the algorithm does not need to know or store the matrix A . The only quantity of interest is the product $A\mathbf{q}_n$, i.e., the action of A on \mathbf{q}_n .

14.3 Arnoldi Iteration as Projection onto Krylov Subspaces

An alternative derivation of Arnoldi iteration starts with the *Krylov matrix*

$$K_n = [\mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, \dots, A^{n-1}\mathbf{b}].$$

Then

$$AK_n = [A\mathbf{b}, A^2\mathbf{b}, A^3\mathbf{b}, \dots, A^n\mathbf{b}]. \quad (42)$$

Since the first $n - 1$ columns of the matrix on the right-hand side are the last $n - 1$ columns of K_n we can also write

$$AK_n = K_n[\mathbf{e}_2, \mathbf{e}_3, \dots, \mathbf{e}_n, -\mathbf{c}],$$

where

$$\mathbf{c} = -K_n^{-1}A^n\mathbf{b}$$

and we assume that K_n is invertible. Equivalently,

$$AK_n = K_n C_n$$

with the upper Hessenberg matrix

$$C_n = [\mathbf{e}_2, \mathbf{e}_3, \dots, \mathbf{e}_n, -\mathbf{c}].$$

This A and C_n are similar via $K_n^{-1}AK_n = C_n$. The problem with this formulation is that K_n is usually ill-conditioned (since all of its columns converge to the dominant eigenvector of A , cf. our earlier discussion of simultaneous power iteration).

Remark As a side remark we mention that the matrix C_n above is known as a *companion matrix*. The matrix has characteristic polynomial $p(z) = z^n + \sum_{i=1}^n c_i z^{i-1}$, where the c_i are the components of \mathbf{c} . In other words, the eigenvalues of C_n are the roots of p . This goes also in the other direction. Given a monic polynomial p , we can form its companion matrix C_n and then know that the roots of p are the same as the eigenvalues of C_n .

Returning to the derivation of Arnoldi iteration, we still need to show how the above Krylov subspace formulation is related to the earlier one based on the Gram-Schmidt method.

Denote the QR factorization of the Krylov matrix K_n by

$$K_n = Q_n R_n.$$

Then

$$\begin{aligned} & K_n^{-1}AK_n = C_n \\ \iff & R_n^{-1}Q_n^*AQ_nR_n = C_n \\ \iff & Q_n^*AQ_n = \underbrace{R_n C_n R_n^{-1}}_{=H_n}. \end{aligned}$$

It needs to be pointed out that this approach is computationally not a good one. It is both too expensive and unstable. On the one hand, finding \mathbf{c} involves solution of the (ill-conditioned) linear system $K_n \mathbf{c} = A^n \mathbf{b}$. On the other hand, we would also be required to provide the inverse of R_n .

However, we can get some theoretical insight from this approach. The formula

$$Q_n^* A Q_n = H_n$$

can be interpreted as an *orthogonal projection of A onto K_n* with the column of Q_n as basis, i.e., Arnoldi iteration is nothing but orthogonal projection onto Krylov subspaces.

Remark If A is Hermitian, then everything above simplifies (e.g., Hessenberg matrices turn into tridiagonal), and we get what is known in the literature separately as *Lanczos iteration*.

14.4 GMRES

The method of *generalized minimum residuals* (or GMRES) was suggested in 1986 by Saad and Schultz.

While application of the classical iterative solvers was limited to either diagonally dominant or positive definite matrices, the GMRES method can be used for linear systems $A\mathbf{x} = \mathbf{b}$ with *arbitrary (nonsingular) square matrices A* . The essential ingredient in this general iterative solver is Arnoldi iteration.

The main idea of the GMRES method is to solve a *least squares problem* at each step of the iteration. More precisely, at step n we approximate the exact solution $\mathbf{x}^* = A^{-1}\mathbf{b}$ by a vector $\mathbf{x}_n \in \mathcal{K}_n$ (the n -th order Krylov subspace) such that the residual

$$\|\mathbf{r}_n\|_2 = \|A\mathbf{x}_n - \mathbf{b}\|_2$$

is minimized. We now describe how to solve this least squares problem.

We start with the Krylov matrix

$$K_n = [\mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, \dots, A^{n-1}\mathbf{b}] \in \mathbb{C}^{m \times n}.$$

Thus the column space of AK_n is AK_n .

Now the desired vector $\mathbf{x}_n \in \mathcal{K}_n$ can be written as

$$\mathbf{x}_n = K_n \mathbf{c}$$

for some appropriate vector $\mathbf{c} \in \mathbb{C}^n$. Therefore the residual minimization becomes

$$\|\mathbf{r}_n\|_2 = \|A\mathbf{x}_n - \mathbf{b}\|_2 = \|AK_n \mathbf{c} - \mathbf{b}\|_2 \rightarrow \min.$$

The obvious way to find the least squares solution to this problem would be to compute the QR factorization of the matrix AK_n . However, this is both unstable and too expensive.

Instead, we look for an orthonormal basis for the Krylov subspace \mathcal{K}_n . We will denote this by $\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n\}$, the columns of the matrix Q_n used in Arnoldi iteration. With this new basis the approximate solution $\mathbf{x}_n \in \mathcal{K}_n$ can be written as

$$\mathbf{x}_n = Q_n \mathbf{y}$$

for some appropriate vector $\mathbf{y} \in \mathbb{C}^n$. The residual minimization is then

$$\|A\mathbf{x}_n - \mathbf{b}\|_2 = \|AQ_n\mathbf{y} - \mathbf{b}\|_2 \rightarrow \min. \quad (43)$$

It is now time to recall the principle behind the Arnoldi iteration. That algorithm is based on the partial similarity transform

$$AQ_n = Q_{n+1}\tilde{H}_n.$$

Thus (43) turns into

$$\|Q_{n+1}\tilde{H}_n\mathbf{y} - \mathbf{b}\|_2 \rightarrow \min.$$

Next we take advantage of the fact that multiplication by a unitary matrix does not change the 2-norm. Thus, we arrive at

$$\begin{aligned} & \|Q_{n+1}^*Q_{n+1}\tilde{H}_n\mathbf{y} - Q_{n+1}^*\mathbf{b}\|_2 \rightarrow \min \\ \iff & \|\tilde{H}_n\mathbf{y} - Q_{n+1}^*\mathbf{b}\|_2 \rightarrow \min \end{aligned}$$

Note that the system matrix AQ_n in (43) is an $m \times n$ matrix, while the new matrix \tilde{H}_n is an $(n+1) \times n$ matrix which is smaller, and therefore will permit a more efficient solution.

The final simplification we can make is for the vector $Q_{n+1}^*\mathbf{b}$. In detail, this vector is given by

$$Q_{n+1}^*\mathbf{b} = \begin{bmatrix} \mathbf{q}_1^*\mathbf{b} \\ \mathbf{q}_2^*\mathbf{b} \\ \vdots \\ \mathbf{q}_{n+1}^*\mathbf{b} \end{bmatrix}.$$

Recall that the Krylov subspaces are given by

$$\begin{aligned} \mathcal{K}_1 &= \text{span}\{\mathbf{b}\}, \\ \mathcal{K}_2 &= \text{span}\{\mathbf{b}, A\mathbf{b}\}, \\ &\vdots, \end{aligned}$$

and that the columns \mathbf{q}_j of Q_n form an orthonormal basis for \mathcal{K}_n . Thus

$$\mathbf{q}_1 = \frac{\mathbf{b}}{\|\mathbf{b}\|}$$

and $\mathbf{q}_j^*\mathbf{b} = 0$ for any $j > 1$. Therefore we actually have

$$Q_{n+1}^*\mathbf{b} = \|\mathbf{b}\|\mathbf{e}_1.$$

Combining all of this work we arrive at the final least squares formulation

$$\begin{aligned} & \left\| \tilde{H}_n\mathbf{y} - \|\mathbf{b}\|\mathbf{e}_1 \right\|_2 \rightarrow \min \\ & \text{with } \mathbf{x}_n = Q_n\mathbf{y}. \end{aligned}$$

This leads to

Algorithm (GMRES)

Let $\mathbf{q}_1 = \mathbf{b}/\|\mathbf{b}\|$

for $n = 1, 2, 3, \dots$

 Perform step n of Arnoldi iteration, i.e., compute new entries for \tilde{H}_n and Q_n .

 Find \mathbf{y} that minimizes $\left\| \tilde{H}_n \mathbf{y} - \|\mathbf{b}\| \mathbf{e}_1 \right\|_2$ (e.g., with QR factorization)

 Set $\mathbf{x}_n = Q_n \mathbf{y}$

end

The computational cost for the GMRES algorithm depends on the cost of the method used for the least squares problem and that of the Arnoldi iteration. Since the least squares system matrix is of size $(n+1) \times n$ it can be done in $\mathcal{O}(n^2)$ floating point operations. If an updating QR factorization (which we did not discuss) is used, then even $\mathcal{O}(n)$ is possible. Arnoldi iteration takes $\mathcal{O}(mn)$ operations plus the cost for matrix-vector multiplications. These can usually be accomplished at somewhere between $\mathcal{O}(m)$ and $\mathcal{O}(m^2)$ flops.

14.4.1 Convergence of GMRES

In *exact* arithmetic (which we of course do not have in a numerical computing environment) GMRES iteration will always converge in at most m steps (since $\mathcal{K}_m = \mathbb{C}^m = \text{range}(A)$). Moreover, convergence is monotonic since $\|\mathbf{r}_{n+1}\| \leq \|\mathbf{r}_n\|$. This latter fact is true since \mathbf{r}_n is minimized over \mathcal{K}_n , and we have $\mathcal{K}_{n+1} \supset \mathcal{K}_n$. Thus, minimization over a larger subspace will allow us to achieve a smaller residual norm.

For practical computations the only case of interest is when the algorithm converges (to within a specified tolerance) in n iterations with $n \ll m$. Usually the *relative residual* is used to control the convergence, i.e., we check whether

$$\frac{\|\mathbf{r}_n\|}{\|\mathbf{b}\|} < \text{tol},$$

where $\text{tol} = 10^{-6}$ or 10^{-8} . The MATLAB code `GMRESDemo.m` illustrates this.

The test problems are given by 200×200 systems with random matrices and random right-hand side vectors \mathbf{b} . The different test matrices are a matrix A whose entries are independent samples from the real normal distribution of mean 2 and standard deviation $0.5/\sqrt{200}$. Its eigenvalues are clustered in a disk in the complex plane of radius $1/2$ centered at $z = 2$. The other test matrices have different eigenvalue distributions. They are obtained as

- $B = A + D$, where the entries of the diagonal matrix D are the complex numbers

$$d_k = (-2 + 2 \sin \theta_k) + i \cos \theta_k, \quad \theta_k = \frac{k}{m-1}, \quad 0 \leq k \leq m-1.$$

- C , a random matrix whose eigenvalues are loosely clustered in the unit disk.

- $D = C^T C$, a random symmetric positive definite matrix (with real positive eigenvalues).
- $E = \frac{1}{2}C$, another random matrix whose eigenvalues are more densely clustered than those of C .

The rate of convergence of the GMRES method depends on the distribution of the eigenvalues of A in the complex plane. For fast convergence the eigenvalues need to be *clustered away from the origin*. Note that the eigenvalue distribution is much more important than the condition number of A , which is the main criterion for rapid convergence of the conjugate gradient method (see next section).

If the GMRES method is applied to a matrix with a “good” eigenvalue distribution (or is appropriately *preconditioned*) then it may solve even a system with an unstructured dense matrix faster than standard LU factorization.