Right-Hand Side Dependent Bounds for GMRES Applied to Ill-Posed Problems

Jennifer Pestana

Right-hand side dependent bounds for GMRES applied to ill-posed problems[*]

Jennifer Pestana

Mathematical Institute, University of Oxford, Andrew Wiles Building,
Radcliffe Observatory Quarter, Woodstock Road, Oxford, OX2 6GG
`pestana@maths.ox.ac.uk`

Abstract. In this paper we apply simple GMRES bounds to the nearly singular systems that arise in ill-posed problems. Our bounds depend on the eigenvalues of the coefficient matrix, the right-hand side vector and the nonnormality of the system. The bounds show that GMRES residuals initially decrease, as residual components associated with large eigenvalues are reduced, after which semi-convergence can be expected because of the effects of small eigenvalues.

Keywords: GMRES; Convergence; Ill-Posed Problem

1   Introduction

The solution of an ill-posed problem often requires the solution of a large, sparse linear system $A\boldsymbol{x} = \boldsymbol{b}$ where $A \in \mathbb{C}^{n \times n}$ is non-Hermitian and nearly singular, $\boldsymbol{b} \in \mathbb{C}^n$ and $\boldsymbol{b} \in \text{range}(A)$ [1]. We assume throughout that A is diagonalizable since, although possible, analysis using the Jordan canonical form is more complicated. The near-singularity of A is reflected in a number of small eigenvalues.

   In many cases $\boldsymbol{b}$ is unknown and we instead possess a noisy vector $\boldsymbol{b}_\delta$, where $\|\boldsymbol{b} - \boldsymbol{b}_\delta\|_2 = \delta$. This is problematic since the ill-conditioning of A means that $A^{-1}\boldsymbol{b}_\delta$ may be a poor approximation of $\boldsymbol{x}$. Consequently, it is necessary to regularize, i.e., to solve

$$A_\delta \boldsymbol{x}_\delta = \boldsymbol{b}_\delta. \tag{1}$$

   The Generalized Minimal Residual method [2] (GMRES) is an iterative method for solving (1) that, given an initial guess $\boldsymbol{x}_0$ which we assume for simplicity is the zero vector, selects at the kth step the iterate $\boldsymbol{x}_k$ for which the residual $\boldsymbol{r}_k = \boldsymbol{b}_\delta - A_\delta \boldsymbol{x}_k$ satisfies

$$\|\boldsymbol{r}_k\|_2 = \min_{\substack{q \in \Pi_k \\ q(0)=1}} \|q(A_\delta)\boldsymbol{b}_\delta\|_2, \tag{2}$$

where Π_k is the set of polynomials of degree at most k. When GMRES is used to solve (1) it can sometimes give good approximations to $\boldsymbol{x}_\delta$ as long as the method

---

is terminated after the correct number of iterations, i.e., GMRES itself can have a regularizing effect [3, 4]. Alternatively, regularization may be achieved by pre-conditioning [4–6]. In either case it is important to understand the behaviour of GMRES applied to nearly singular systems. Eldén and Simoncini [4] used the Schur decomposition to show that when the right-hand side has leading components in the direction of eigenvectors associated with large eigenvalues, the initial convergence is related to a reduction in the sizes of these components. Here we provide a complementary analysis involving the eigenvalue-eigenvector decomposition and the simple bounds in Titley-Peloquin, Pestana and Wathen [7]. Similarly to Eldén and Simoncini we find that the first phase of convergence is related to large eigenvalues. We additionally observe that the stagnation typically observed in the second phase, known as semi-convergence, is attributable to the remaining small eigenvalues.

2   Structure of nearly singular systems

Let A_δ have diagonalization $A_\delta = Z\Lambda Z^{-1}$, $\Lambda = \mathrm{diag}(\lambda_i)$ and $Z \in \mathbb{C}^{n \times n}$, where without loss of generality $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_n|$. We wish to separate the spectrum of A_δ into p large eigenvalues and the remaining small eigenvalues. The matrix A_δ may have two distinct sets of eigenvalues, for example, when a preconditioner is applied. In other cases, however, there is no obvious separation. In this situation we find that a division on the order of δ is a reasonable choice.

Given these two sets of eigenvalues we partition A_δ as

$$A_\delta = \begin{bmatrix} Z_1 & Z_2 \end{bmatrix} \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix} \begin{bmatrix} Y_1^* \\ Y_2^* \end{bmatrix},$$

where $\Lambda_1 \in \mathbb{C}^{p \times p}$, $\Lambda_2 \in \mathbb{C}^{(n-p) \times (n-p)}$, $Z_1, Y_1 \in \mathbb{C}^{n \times p}$ and $Z_2, Y_2 \in \mathbb{C}^{n \times (n-p)}$. We assume that $\|Y_2^* \boldsymbol{b}\|_2 = \epsilon$ is small, i.e., that the true right-hand side vector $\boldsymbol{b}$ is mainly associated with the low-frequency components of A_δ; otherwise the ill-posed problem is intractable.

Integral to our bounds are the co-ordinates of $\boldsymbol{b}_\delta$ in the eigenvector basis

$$\boldsymbol{w} = Z^{-1} \boldsymbol{b}_\delta / \|\boldsymbol{b}_\delta\|_2 = \begin{bmatrix} \boldsymbol{w}_1 \\ \boldsymbol{w}_2 \end{bmatrix} = \frac{1}{\|\boldsymbol{b}_\delta\|_2} \begin{bmatrix} Y_1^* \boldsymbol{b}_\delta \\ Y_2^* \boldsymbol{b}_\delta \end{bmatrix} \tag{3}$$

and, in particular, $\boldsymbol{w}_2 = (Y_2^* \boldsymbol{b} + Y_2^* (\boldsymbol{b}_\delta - \boldsymbol{b})) / \|\boldsymbol{b}_\delta\|_2$, the norm of which is bounded by

$$\|\boldsymbol{w}_2\|_2 \leq (\epsilon + \delta \|Y_2\|_2) / \|\boldsymbol{b}_\delta\|_2. \tag{4}$$

To give some idea of typical spectra, and to show the difference between the components of $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$, we compute these quantities for the baart and wing test problems from the Matlab toolbox Regularization Tools [8, 9]. The problems are described in more detail in Section 4. We add Gaussian noise to the true right-hand side vectors with $\delta = 10^{-7}$, 10^{-5} and 10^{-3}. For baart, $\|\boldsymbol{b}_\delta\|_2 \approx 2.9$, $\|Y_2\|_2 = 64$ and $\epsilon = 10\delta$ when $p = 5$. Thus, (4) gives $\|\boldsymbol{w}_2\|_2 \leq 26\delta$ for baart. For wing, $\|\boldsymbol{b}_\delta\|_2 \approx 0.15$ and $\|Y_2\|_2 = 158$ with $p = 3$. We find that when δ is 10^{-7},

10^{-5} and 10^{-3}, ϵ is 1×10^{-5}, 3.6×10^{-4} and 9×10^{-3}, so that (4) is 2×10^{-4}, 0.01 and 1.

| (a) baart | (b) wing |
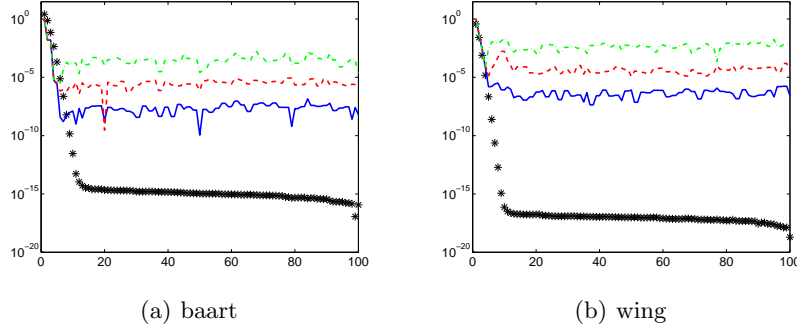
Fig. 1. Magnitudes of eigenvalues ($*$) of A_δ and of corresponding components of $\boldsymbol{w}$ for $\delta = 10^{-7}$ (solid line) $\delta = 10^{-5}$ (dashed line) and $\delta = 10^{-3}$ (dot-dashed line).

Figure 1 shows that for both problems, as expected, the eigenvalues decay and there are a number of very small eigenvalues present. Associated with large eigenvalues are relatively large components of $\boldsymbol{w}$ in magnitude. Once the eigenvalues decrease to around the level of the noise, the components of $\boldsymbol{w}$ stay constant in magnitude at a level that depends on $\|\boldsymbol{b}_\delta\|_2$, the amount of noise and the conditioning of the eigenvectors associated with small eigenvalues. This level is, consequently, higher for wing than for baart. The structure of these two systems is typical of ill-posed linear systems and is exploited in the next section to analyse the convergence of GMRES.

3   GMRES bounds

Our interest is in explaining the behaviour of GMRES applied to (1). To this end, we apply the bounds in Section 2 in Titley-Peloquin *et al.* [7], the first of which is cast in terms of a weighted least squares problem.

Theorem 1. *Suppose that A_δ has diagonalization $A_\delta = Z \Lambda Z^{-1}$, $\Lambda = \mathrm{diag}(\lambda_i)$, and let $\boldsymbol{w}_1 = W_1 \boldsymbol{e}$ and $\boldsymbol{w}_2 = W_2 \boldsymbol{e}$, where $\boldsymbol{w}_1$ and $\boldsymbol{w}_2$ are as in (3), $W = \mathrm{diag}(w_i)$ and $\boldsymbol{e} = [1, \ldots, 1]^T$. Then the GMRES residuals satisfy*

$$\frac{\|\boldsymbol{r}_k\|_2}{\|\boldsymbol{b}\|_2} \leq \|Z\|_2 \min_{\substack{q \in \Pi_k \\ q(0)=1}} \left\| \begin{bmatrix} W_1 q(\Lambda_1) & \\ & W_2 q(\Lambda_2) \end{bmatrix} \boldsymbol{e} \right\|_2. \tag{5}$$

For our ill-posed problem, the weights in W_1 are larger in magnitude than those in W_2 and the eigenvalues in Λ_1 are all larger in magnitude than the eigenvalues in Λ_2. Thus, GMRES will initially choose polynomials that primarily

reduce the size of $W_1q(\Lambda_1)$ to the size of $W_2q(\Lambda_2)$. In particular, when $\|\boldsymbol{w}_1\|_2 \gg \|\boldsymbol{w}_2\|_2$ we would expect that for the first p steps GMRES would mainly work on reducing the components of the residual associated with Λ_1 and Z_1.

When $\|W_1q(\Lambda_1)\|_2$ is on the order of $\|W_2q(\Lambda_2)\|_2$ it is common for convergence to stagnate, after which residuals may increase in norm; this is known as semi-convergence. The following theorem can help to explain why semi-convergence occurs by explicitly separating the effects of large and small eigenvalues [7].

Theorem 2. *Let A_δ have diagonalization $A_\delta = Z\Lambda Z^{-1}$. For any subset of indices $\mathcal{J}$ with $|\mathcal{J}| = p$, GMRES residuals with $k > p$ satisfy*

$$\frac{\|\boldsymbol{r}_k\|_2}{\|\boldsymbol{b}_\delta\|_2} \leq \|Z\|_2 \min_{\substack{q \in \Pi_{k-p} \\ q(0)=1}} \left(\sum_{\substack{i=1 \\ i \notin \mathcal{J}}}^{n} |\tilde{w}_i|^2 |q(\lambda_i)|^2 \right)^{1/2}, \tag{6}$$

where

$$\tilde{w}_i = w_i \prod_{j \in \mathcal{J}} \left(1 - \frac{\lambda_i}{\lambda_j} \right).$$

To examine the semi-convergence phase, we choose $\mathcal{J} = [1, p]$. Then for any $i \in [p+1, n]$, we have that $|\lambda_i| \leq |\lambda_j|$ and $|\tilde{w}_i| \leq \alpha^p |w_i|$, where $\alpha \leq 2$ and α is around 1 or smaller when, say, there is a decent gap between the large and small eigenvalues or when all eigenvalues have the same sign. Thus, for any $k > p$

$$\frac{\|\boldsymbol{r}_k\|_2}{\|\boldsymbol{b}_\delta\|_2} \leq \alpha^p \|Z\|_2 \|\boldsymbol{w}_2\|_2 \min_{\substack{q \in \Pi_{k-p} \\ q(0)=1}} \|q(\Lambda_2)\|_2.$$

Now, let us consider $\|q(\Lambda_2)\|_2$. Since $|\lambda_i| \ll 1$, $i = p+1, \ldots, n$ and $q(0) = 1$ it will be difficult to reduce $\|q(\Lambda_2)\|_2$ significantly below 1. Consequently, we expect the residuals to stagnate at a level bounded by

$$\frac{\|\boldsymbol{r}_k\|_2}{\|\boldsymbol{b}_\delta\|_2} \leq \alpha^p \|Z\|_2 \|\boldsymbol{w}_2\|_2. \tag{7}$$

This, in conjunction with (4), indicates that the level of semi-convergence depends on the sizes of the large and small eigenvalues, the noise level δ, the norm of $\boldsymbol{b}_\delta$ and the conditioning of the eigenvectors associated with small eigenvalues.

4 Numerical results

We now compare the bounds (5) and (7) to the GMRES residuals for the baart and wing problems mentioned above, both of which are discretizations of Fredholm integral equations of the first kind. The integral equation for baart is

$$\int_0^\pi e^{s \cos(t)} f(t) dt = 2 \frac{\sinh(s)}{s}, \quad 0 \leq s \leq \frac{\pi}{2},$$

which has the continuous solution $f(t) = \sin(t)$. For the wing problem we solve

$$\int_0^1 te^{-st^2} f(t)dt = \frac{e^{-st_1^2} - e^{-st_2^2}}{2s}, \quad 0 \leq s \leq 1,$$

with $t_1 = 1/3$ and $t_2 = 2/3$. The discontinuous solution is

$$f(t) = \begin{cases} 1 & t_1 < t < t_2, \\ 0 & \text{elsewhere} . \end{cases}$$
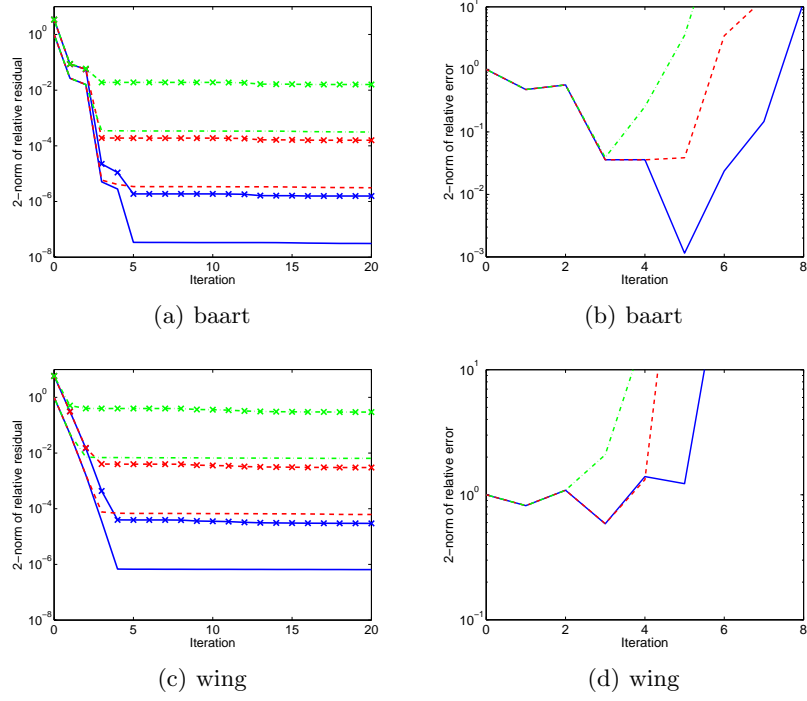
Figure 2 shows the relative GMRES residuals and the relative errors. For both baart and wing the relative residuals decrease before stagnating at a level related to the noise level δ. Note that the staircase-like convergence behaviour for baart is particular to this problem. It appears to be related to the harmonic Ritz values, which at the kth step of GMRES are the eigenvalues of a certain $k \times k$ matrix, and which define the GMRES polynomial q in (2) [10]. For fast convergence it is desirable that these harmonic Ritz values are good approximations of eigenvalues of A. For baart, however, at the second and fourth steps there is a harmonic Ritz value that lies between two consecutive eigenvalues of A; these are precisely the steps at which there is little reduction in the relative residual norm.

Unlike the relative residuals, for both problems the norm of the error initially decreases but then starts to increase. This increase occurs during the semi-convergence phase for baart but for the wing problem the errors increase before semi-convergence and exhibit a sawtooth-like behaviour. This highlights the importance of applying a sensible stopping criterion and the potential unsuitability of standard (unpreconditioned) GMRES for some ill-posed problems. Interestingly, (5) seems to provide a better indication of when the iterations should be stopped than the onset of semi-convergence for the wing problem for noisy right-hand side vectors, although we have not investigated this further.

It is clear from Figure 2 that the bound (5) is very descriptive during the first phase of convergence. Although the bound is not quantitatively descriptive in the second phase of convergence, it accurately predicts the onset of semi-convergence. The approximation (7) is an upper bound on the relative residuals during the semi-convergence phase for both problems (see Figure 3). Note that for both problems $\alpha \approx 1$. Since (6) is an upper bound on (5), we cannot expect (7) to be quantitatively accurate. Nevertheless, it provides an analysis of semi-convergence and the factors that can affect the level at which residual norms stagnate.

Conclusions

In this paper we have applied simple bounds on GMRES convergence to the nearly singular systems that arise from ill-posed problems. We have shown that GMRES initially reduces the residual components associated with large eigenvalues. Once these components are commensurate with those associated with small eigenvalues semi-convergence sets in, with the level at which residuals stagnate determined by the sizes of small eigenvalues, the noise in the right-hand side vector, the size of b and the eigenvectors associated with small eigenvalues.

Fig. 2. Plots of the relative GMRES residuals and (5) (×) (left) and relative errors (right) for $\delta = 10^{-7}$ (solid line) $\delta = 10^{-5}$ (dashed line) and $\delta = 10^{-3}$ (dot-dashed line).
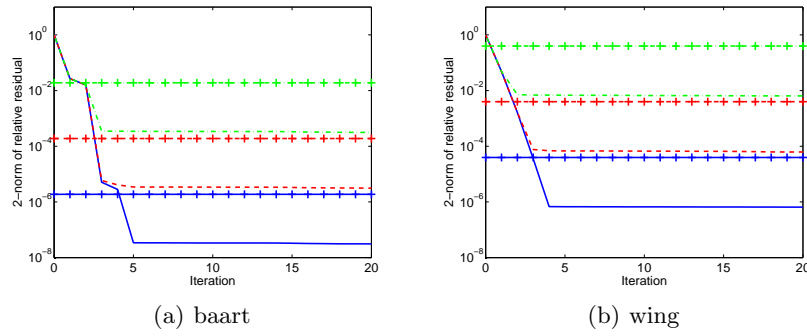
Fig. 3. Plots of the relative GMRES residuals and (7) (+) for $\delta = 10^{-7}$ (solid line) $\delta = 10^{-5}$ (dashed line) and $\delta = 10^{-3}$ (dot-dashed line).

References

1. Hansen, P.C.: Rank-Deficient and Discrete Ill-Posed Problems: Numerical Aspects of Linear Inversion. SIAM, Philadelphia, PA (1998)
2. Saad, Y., Schultz, M.H.: GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. SIAM J. Sci. Stat. Comput. **7** (1986) 856–869
3. Calvetti, D., Lewis, B., Reichel, L.: On the regularizing properties of the GMRES method. Numer. Math. **91** (2002) 605–625
4. Eldén, L., Simoncini, V.: Solving ill-posed linear systems with GMRES and a singular preconditioner. SIAM J. Matrix Anal. Appl. **33** (2012) 1369–1394
5. Hanke, M., Nagy, J.: Inverse Toeplitz preconditioners for ill-posed problems. Linear Algebra Appl. **284** (1998) 137–156
6. Hansen, P.C., Jensen, T.K.: Smoothing-norm preconditioning for regularizing minimum-residual methods. SIAM J. Matrix Anal. Appl. **29** (2006) 1–14
7. Titley-Peloquin, D., Pestana, J., Wathen, A.J.: GMRES convergence bounds that depend on the right-hand-side vector. IMA Journal of Numerical Analysis (2013)
8. Hansen, P.C.: Regularization Tools Version 4.0 for Matlab 7.3. Numer. Algorithms **46** (2007) 189–194
9. Hansen, P.C.: Regularization Tools Version 4.1 (for Matlab Version 7.3)
10. Goossens, S., Roose, D.: Ritz and harmonic Ritz values and the convergence of FOM and GMRES. Numer. Linear Algebra Appl. **6** (1997) 281–293