J. Austral. Math. Soc. Ser. B 36(1995), 381–388

RANK-1 PERTURBATIONS AND THE LANCZOS METHOD, INVERSE ITERATION, AND KRYLOV SUBSPACES

CHRISTOPHER T. LENARD[1]

(Received 3 July 1993; revised 16 September 1993)

Abstract

The heart of the Lanczos algorithm is the systematic generation of orthonormal bases of invariant subspaces of a perturbed matrix. The perturbations involved are special since they are always rank-1 and are the smallest possible in certain senses. These minimal perturbation properties are extended here to more general cases.

Rank-1 perturbations are also shown to be closely connected to inverse iteration, and thus provide a novel explanation of the global convergence phenomenon of Rayleigh quotient iteration.

Finally, we show that the restriction to a Krylov subspace of a matrix differs from the restriction of its inverse by a rank-1 matrix.

1. Introduction

Rank-1 matrices are an integral part of numerical linear algebra; to give just one example, the reduction of a matrix to tridiagonal form by Householder transformations. More recently, rank-1 matrices have found other uses in eigenvalue problems. For example, the divide-and-conquer strategy uses a rank-1 perturbation to decompose a tridiagonal matrix into a decoupled tridiagonal form (see [1]). In [2] a rank-1 perturbation is also used as a starting point for homotopy methods.

To motivate what is to follow, we begin with some elementary facts concerning rank-1 perturbations. These are in the spirit of backward error analysis and show how small a perturbation is required to specify an eigenpair.

Fact 1 Any scalar μ and unit vector z is an eigenpair of a perturbation of A:

$$[A - (A - \mu)zz^{\mathsf{T}}]z = \mu z,$$
$$\|(A - \mu)zz^{\mathsf{T}}\| = \|(A - \mu)z\|.$$

[1]Department of Mathematics, LaTrobe University, Bendigo, PO Box 199, Bendigo VIC 3550, Australia.
© Australian Mathematical Society, 1995, Serial-fee code 0334-2700/95

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

Fact 2 The above perturbation is the smallest possible: if E is any other matrix for which $(A - E)z = \mu z$, then

$$\|(A - \mu)z\| = \|Ez\| \leq \|E\| .$$

Fact 3 The norm of the residual $\|(A - \mu)z\|$ is minimized by the Rayleigh quotient

$$\rho(z) = \langle Az, z\rangle / \langle z, z\rangle ,$$

that is,

$$\|(A - \rho(z))z\| \leq \|(A - \mu)z\| \quad \forall \mu.$$

Thus $D = (A - \rho(z))zz^{\mathsf{T}}$ is the smallest matrix for which the unit vector z is an eigenvector of $A - D$.

2. The Lanczos method

In this section, we will see how the preceding results apply to the Lanczos method.

The Lanczos method for approximating outer eigenvalues and eigenvectors of a symmetric $n \times n$ matrix A can be motivated in several different ways; for details see [4]. Algebraically, it is the successive formation of $m \times m$ tridiagonal matrices T_m defined by

$$T_m = Q_m^{\mathsf{T}} A Q_m. \tag{1}$$

The $n \times m$ orthonormal matrices Q_m are built up column by column,

$$Q_{m+1} = [q_1, \ldots, q_{m+1}]$$

from vectors defined by

$$\begin{aligned} AQ_m - Q_m T_m &= \hat{q}_{m+1} e_m^{\mathsf{T}} , \\ q_{m+1} &= \hat{q}_{m+1} / \|\hat{q}_{m+1}\| , \\ e_m &= (0, \ldots, 1)^{\mathsf{T}} \in \mathcal{R}^m . \end{aligned} \tag{2}$$

The matrix T_m is the restriction of A to the subspace spanned by the columns of Q_m and as such its eigenvalues ρ_i are approximations to outer eigenvalues of A. The corresponding eigenvectors, y_i, of T_m need to be transformed as $x_i = Q_m y_i$ to yield approximations to the corresponding eigenvectors of A.

From now on we will assume m is fixed and understand that the ρ_i, x_i, y_i depend on m. We shall also assume the x_i are normalized : $\|x_i\| = 1$.

Of course, the Lanczos *algorithm* does not form the matrices explicitly in (1) and (2), but rather is a clever way of recursively generating the vectors q_m and using

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

this information to obtain T_{m+1} by bordering T_m. For our purposes, though, (2) highlights an essential but often overlooked feature, namely, the rank-1 right hand side. It is this aspect which we focus on in this section.

From (2) we get immediately

$$(A - \rho_i)x_i = \hat{q}_{m+1}e_m^{\mathsf{T}}y_i.$$

That is, each residual is a multiple of q_{m+1}.

Parallel residuals is essentially a geometric property. The elegance of the Lanczos algorithm turns on the fact that the orthonormal columns of Q_m span the Krylov subspace

$$\mathcal{K}_m = \operatorname{span}\{q_1, \ldots, A^{m-1}q_1\}.$$

The scalars ρ_i and vectors x_i are the eigenpairs of the projection of A onto $\mathcal{K}_m$:

$$Q_m Q_m^{\mathsf{T}} A Q_m Q_m^{\mathsf{T}} x_i = \rho_i x_i, \qquad i = 1, \ldots, m,$$

which is just the Galerkin approximation on $\mathcal{K}_m$. Therefore each residual $r_i = (A - \rho_i)x_i$ is orthogonal to $\mathcal{K}_m$ (which is the motivation and definition of the Galerkin approximation). At the same time, it is clear that $r_i \in \mathcal{K}_{m+1}$. However, $\mathcal{K}_{m+1}$ is at most one dimension more than $\mathcal{K}_m$, and therefore all the residuals are contained in a 1-dimensional subspace. So we conclude as before that the residuals are scalar multiples of a unit vector r : $r_i = \varepsilon_i r$. And of course $r = q_{m+1}$.

Since the x_i are mutually orthogonal, we have, similar to Fact 1,

$$\left[A - \sum r_j x_j{}^{\mathsf{T}}\right]x_i = \rho_i x_i.$$

On setting $s = \sum \varepsilon_j x_j{}^{\mathsf{T}}$, this is $[A - rs^{\mathsf{T}}]x_i = \rho_i x_i$ or

$$[A - rs^{\mathsf{T}}]X = X\Delta\,, \tag{3}$$

$$X = [x_1, \ldots, x_m], \quad \Delta = \operatorname{diag}(\rho_1, \ldots, \rho_m).$$

In this form, the Lanczos method is interpreted as producing an invariant subspace of a perturbation of A.

In [1], it is shown how this can be carried one step further by constructing a symmetric rank-1 matrix E for which the subspace spanned by the columns of X is an invariant subspace of $A + E$. This can then be used as another way of approximating the eigenvalues of A.

That the perturbation which occurs in the Lanczos method is necessarily rank-1 rests on the fact that all the residuals are parallel. This last property is shared by all rank-1 changes, that is, if $(A - ab^{\mathsf{T}})z = \mu z$, then $(A - \mu)z$ is a multiple of a.

It also turns out that the perturbation rs^{T} is the smallest possible in the same sense as Fact 2.

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

THEOREM 1. *If E is such that $(A - E)x_i = \rho_i x_i$ for each i, then*

$$\|rs^{\mathsf{T}}\| \le \|E\|.$$

PROOF. Since the x_i are orthonormal, and $\|r\| = 1$, we have

$$\|rs^{\mathsf{T}}\|^2 = \|s\|^2 = \sum \varepsilon_j{}^2 = \left\|\sum \varepsilon_j{}^2 r\right\|.$$

But

$$\begin{aligned}\left\|\sum \varepsilon_j{}^2 r\right\| &= \left\|\sum \varepsilon_j r_j\right\| \\ &= \left\|\sum \varepsilon_j E x_j\right\| \\ &= \left\|E \sum \varepsilon_j x_j\right\| \\ &= \|Es\| \\ &\le \|E\|\,\|s\|.\end{aligned}$$

That is, $\|rs^{\mathsf{T}}\| \le \|E\|$.

ALTERNATIVE PROOF. First write $(A - E)X = X\Delta$, which, with (3), gives $EX = rs^{\mathsf{T}}X$. This, together with the fact that X is orthonormal and $\|r\| = 1$, yields,

$$\|s\| = \|rs^{\mathsf{T}}X\| = \|EX\| \le \|E\|\,\|X\| = \|E\|.$$

More generally, for any subspace spanned by the columns of an orthonormal Z, there is a perturbation of A for which span(Z) is an invariant subspace.

THEOREM 2. *Suppose $Z = [z_1, \ldots, z_m]$ is orthonormal, M is $m \times m$, and $R = [AZ - ZM]Z^{\mathsf{T}}$, then*

$$(A - R)Z = ZM,$$
$$\|R\| = \|AZ - ZM\|$$

and $\|R\| \le \|E\|$ for every E such that $(A - E)Z = ZM$.

PROOF. Since Z is orthonormal, it follows that

$$\|R\| = \|(AZ - ZM)Z^{\mathsf{T}}\| = \|AZ - ZM\| = \|RZ\|$$

and

$$\|RZ\| = \|EZ\| \le \|E\|\,\|Z\| = \|E\|.$$

Just as the norm of the residual is minimized by the Rayleigh quotient (Fact 3), so it is that $\|AZ - ZM\|$, and consequently $\|R\|$, is minimized by $M = Z^{\mathsf{T}}AZ$ (see [4]).

As a final comment, we note that as the Lanczos algorithm proceeds, it is not necessarily the case that the norms of the perturbations rs^{T} decrease monotonically.

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

3. Inverse iteration and iterated Galerkin approximations

3.1. Inverse iteration As noted before, $[A - (A - \mu)zz^{\mathsf{T}}]z = \mu z$, for an arbitrary A, scalar μ and unit vector z. Now consider the corresponding left eigenvector y given by

$$y^{\mathsf{T}}[A - (A - \mu)zz^{\mathsf{T}}] = y^{\mathsf{T}}\mu,$$

that is,

$$(A - \mu)^{\mathsf{T}} y = zz^{\mathsf{T}}(A - \mu)^{\mathsf{T}} y.$$

When A is symmetric, the left eigenvector y is the inverse iterate of z, and, in general, if μ is near an eigenvalue, y will be closer than z to the corresponding eigenvector.

The idea of smallest perturbation occurs in this context too, and is closely connected to the global convergence property of Rayleigh quotient iteration.

THEOREM 3. *If D, μ, and unit vectors x and y satisfy*

$$\begin{aligned} (A - D)x &= \mu x\,, \\ y^{\mathsf{T}}(A - D) &= \mu y^{\mathsf{T}} \quad \textit{and} \\ \|D\| &\le \|E\| \quad \forall E : (A - E)x = \mu x\,, \end{aligned}$$

then

$$\|(A - \mu)^{\mathsf{T}} y\| \le \|(A - \mu)x\|\,.$$

PROOF.

$$\|(A - \mu)^{\mathsf{T}} y\| = \|D^{\mathsf{T}} y\| \le \|D^{\mathsf{T}}\| = \|D\| = \|(A - \mu)x\|\,.$$

The final equality is just Fact 2.

That is, the residual of the left eigenvector y (with respect to A^{T}) is less than the residual of the right eigenvector (with respect to A). Basically this means that y is a better approximation to a left eigenvector than x is to a right eigenvector. The implication of this is that iterating with the adjoint is the best way to improve an approximation to a right eigenvector. It also strongly suggests that inverse iteration should really be performed with $(A - \mu)^{-1}(A - \mu)^{-\mathsf{T}}$ when A is nonsymmetric.

Combining Theorem 3 with Fact 3 results in a novel interpretation of some of the properties of Rayleigh quotient iteration as defined by

$$(A - \rho(x_k))x_{k+1} = \alpha_k x_k \quad , \quad \|x_{k+1}\| = 1.$$

Using the equivalence between the inverse iterate and left eigenvector this is recast as

$$x_{k+1}^{\mathsf{T}}(A - R_k) = \rho(x_k)x_{k+1}^{\mathsf{T}}, \quad R_k = (A - \rho(x_k))x_k x_k^{\mathsf{T}}\,.$$

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

It is evident from Theorem 3 that the residuals $\|(A-\rho(x_k))x_k\|$ decrease monotonically when A is normal. In fact, this results in the global convergence property of Rayleigh quotient iteration (see [3, 4]).

Since $\|R_k\| = \|(A-\rho(x_k))x_k\|$, the vectors x_k are eigenvectors of successively smaller perturbations of A, while the scalars $\rho(x_k)$ are such that the perturbation at each step is the smallest possible given x_k.

Bearing in mind the remark about iterating with $(A-\mu)^{-1}(A-\mu)^{-\mathsf{T}}$, we see that the appropriate strategy in Rayleigh quotient iteration is

$$(A-\rho(x_k))^{\mathsf{T}}(A-\rho(x_k))x_{k+1} = \alpha_k x_k,$$

which is just the alternative version of the iteration examined in [3].

Recall that Lanczos, at step m, produces vectors x_i with $(A-rs^{\mathsf{T}})x_i = \rho(x_i)x_i$. In view of the above, it is natural to ask if the left eigenvectors y_i of $A-rs^{\mathsf{T}}$ corresponding to ρ_i have smaller residuals than the right (Lanczos) eigenvectors x_i. In fact this is not necessarily the case; however, we may expect the residuals of y_i to be less than those of the x_i in a collective sense. Such a result is contained in the following more general situation.

Suppose Z is orthonormal, $R = [AZ - ZM]Z^{\mathsf{T}}$ for some M, and $Y^{\mathsf{T}}(A-R) = MY^{\mathsf{T}}$, then

$$\|Y^{\mathsf{T}}A - MY^{\mathsf{T}}\| = \|Y^{\mathsf{T}}R\| \le \|R\|\|Y^{\mathsf{T}}\|,$$

that is,

$$\|A^{\mathsf{T}}Y - YM^{\mathsf{T}}\|\|Y\|^{-1} \le \|R\|.$$

In the Lanczos case, $Z = X$, $M = \Delta$, $R = rs^{\mathsf{T}}$ and $\|R\| = \|s\|$, but the columns of Y are not necessarily mutually orthogonal, even though they can be normalized. We therefore cannot conclude that

$$\sum \|(A-\rho(y_i)I)y_i\|^2 \le \|s\|^2$$

and counterexamples to this inequality are easily found. However, we do have

$$\begin{aligned}\|(A-\rho(y_i))y_i\| &\le \|(A-\rho(x_i))y_i\| \\ &= \|sr^{\mathsf{T}}y_i\| \\ &= \|s\||r^{\mathsf{T}}y_i|,\end{aligned}$$

so

$$\sum \|(A-\rho(y_i))y_i\|^2 \le \|s\|^2 \sum |r^{\mathsf{T}}y_i|^2.$$

Again, it is not necessarily true that $\sum |r^{\mathsf{T}}y_i|^2 \le 1$. But $\sum |r^{\mathsf{T}}y_i|^2 = \|Vr\|^2 \le \|V\|^2$, where $V = \sum x_i y_i{}^{\mathsf{T}}$, so $\sum \|(A-\rho(y_i))y_i\|^2 \le \|V\|^2\|s\|^2$.

So far it still an open question as to how Theorem 3 can be extended to cover the Lanczos method.

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

3.2. Iterated Galerkin approximation As an aside, we comment on the way left eigenvectors also occur as refinements in Sloan's iteration of Galerkin approximations examined in [5].

Sloan's iteration arises in connection with a sequence of Galerkin approximations

$$P_i A x_i = \mu_i x_i, \quad P_i \to I, \quad (\mu_i, x_i) \to (\lambda, u).$$

The P_i are projections onto subspaces in a nested sequence. The idea behind Sloan's iteration is that the sequence $\{Ax_i/\mu_i\}$ converges more rapidly than does $\{x_i\}$. What interests us here is the fact that $AP_iAx_i = \mu_i Ax_i$. When A is symmetric this is just

$$(P_i A)^{\mathrm{T}} A x_i = \mu_i A x_i .$$

In other words, the Sloan iterates are simply the left eigenvectors of the Galerkin approximations $P_i A$.

4. Krylov subspaces

For various reasons it is of some interest to examine the Galerkin-Krylov method for the matrix A^{-1}

$$PA^{-1} y_i = \mu_i y_i,$$

where P is the orthogonal projection onto the Krylov subspace

$$\mathcal{K}_m = \operatorname{span}\{q,\ Aq, \ldots,\ A^{m-1}q\}.$$

This leads to a surprising result concerning the 'Gram' matrices

$$G(A) = [\langle Ax_i, x_j\rangle] = X^{\mathrm{T}}AX \quad \text{and} \quad G(A^{-1}) = [\langle A^{-1}x_i, x_j\rangle] = X^{\mathrm{T}}A^{-1}X,$$

where $\{x_i\}$ is an orthonormal basis for $\mathcal{K}_m$ and $X = [x_1, \ldots, x_m]$.

THEOREM 4. *$G(A^{-1}) - G(A)^{-1}$ is rank-1.*

PROOF. We use the same notation as in Section 1. Without losing any generality, we may take the x_i to be the unit eigenvectors of PA.

Now $G(A) = \operatorname{diag}(\rho_1, \ldots, \rho_m) = \Delta$, by properties of the vectors x_i, and

$$G(A^{-1}) - G(A)^{-1} = X^{\mathrm{T}}A^{-1}X - \Delta^{-1}.$$

But, recalling the results of Section 1,

$$[A - rs^{\mathrm{T}}]X = X\Delta,$$

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press

that is,

$$\Delta^{-1} = X^{\mathsf{T}}[A - rs^{\mathsf{T}}]^{-1}X$$
$$= X^{\mathsf{T}}[A^{-1} + \frac{A^{-1}rs^{\mathsf{T}}A^{-1}}{1 - s^{\mathsf{T}}A^{-1}r}]X\,.$$

Therefore

$$G(A^{-1}) - G(A)^{-1} = \frac{-X^{\mathsf{T}}A^{-1}r\,s^{\mathsf{T}}A^{-1}X}{1 - s^{\mathsf{T}}A^{-1}r}\,,$$

which has rank 1 as claimed.

This relies crucially on the x_i being orthonormal vectors which span a Krylov subspace for some q. Neither condition can be relaxed.

Acknowledgements

This work was completed at the California Institute of Technology and supported in part by the NSF under Cooperative Agreement No. CCR-9120008. The U. S. Government has certain rights in this material. Thank you also to Bob Anderssen and DMS, CSIRO Australia for facilities to prepare this report.

References

[1] G. H. Golub and C. F. Van Loan, *Matrix computations*, 2 ed. (Johns Hopkins University Press, Baltimore, 1989).

[2] C. T. Lenard, "Rank-1 homotopy methods for eigenproblems", submitted for publication.

[3] B. N. Parlett, "The Rayleigh quotient iteration and some generalizations for non-normal matrices", *Math. Comp.* **13** (1974) 679–693.

[4] B. N. Parlett, *The symmetric eigenvalue problem* (Prentice-Hall, Englewood Cliffs, New Jersey, 1980).

[5] I. H. Sloan, "Iterated Galerkin method for eigenvalue problems", *SIAM J. Numer. Anal.* **13** (1976) 753–760.

https://doi.org/10.1017/S0334270000007438 Published online by Cambridge University Press