University of Maryland College Park

Institute for Advanced Computer Studies TR–2000-01
Department of Computer Science TR–4099

On the Eigensystems of Graded Matrices*

G. W. Stewart†

January 2000

ABSTRACT

Informally a graded matrix is one whose elements show a systematic decrease or increase as one passes across the matrix. It is well known that graded matrices often have small eigenvalues that are determined to high relative accuracy. Similarly, the eigenvectors can have small components that are nonetheless well determined. In this paper, we give approximations to the eigenvalues and eigenvectors of a graded matrix in terms of a base matrix that show how these phenomena come about. This approach provides condition numbers for eigenvalues and individual components of the eigenvectors. The results are applied to derive related results for the singular value decomposition.

*This report is available by anonymous ftp from thales.cs.umd.edu in the directory pub/reports or on the web at http://www.cs.umd.edu/~stewart/.

†Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (stewart@cs.umd.edu). This work was supported in part by the National Science Foundation under Grant No. 970909-8426.

On the Eigensystems of Graded Matrices

G. W. Stewart

Dedicated to F. L. Bauer

What, then, is time?
If no one asks of me, I know;
If I wish to explain it to him who asks,
I know not.

St. Augustine

I shall not attempt further to define it ...
but I know it when I see it.

Justice Potter Stewart (on pornography)

1. Introduction

Informally a graded matrix is one whose elements show a systematic decrease or increase as one passes across the matrix. Thus we would recognize a matrix whose elements have the magnitudes

$$\begin{pmatrix} 10^{0} & 10^{-4} & 10^{-8} & 10^{-12} \\ 10^{0} & 10^{-4} & 10^{-8} & 10^{-12} \\ 10^{0} & 10^{-4} & 10^{-8} & 10^{-12} \\ 10^{0} & 10^{-4} & 10^{-8} & 10^{-12} \end{pmatrix}$$

as being graded by columns. Similarly, the matrices

$$\begin{pmatrix} 10^{0} & 10^{0} & 10^{0} & 10^{0} \\ 10^{-4} & 10^{-4} & 10^{-4} & 10^{-4} \\ 10^{-8} & 10^{-8} & 10^{-8} & 10^{-8} \\ 10^{-12} & 10^{-12} & 10^{-12} & 10^{-12} \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} 10^{0} & 10^{-2} & 10^{-4} & 10^{-6} \\ 10^{-2} & 10^{-4} & 10^{-6} & 10^{-8} \\ 10^{-4} & 10^{-6} & 10^{-8} & 10^{-10} \\ 10^{-6} & 10^{-8} & 10^{-10} & 10^{-12} \end{pmatrix}$$

are respectively row graded and diagonally graded.

It has long been recognized that the eigenvalues and eigenvectors of graded matrices have special properties. For example, Martin, Reinsch, and Wilkinson writing in 1968 about Householder tridiagonalization [8] assert that a diagonally graded matrix will have small eigenvalues that are insensitive to small relative perturbations in the elements of the matrix. They go on to assert that if the direction of the grading is consonant with their implementation of the algorithm, then the small eigenvalues will be computed accurately.

Informally, then, we know what a graded matrix is — just as St. Augustine knows what time is and Potter Stewart knows pornography when he sees it. Unfortunately, when it comes to a formal definition we encounter difficulties.

1. If we attempt to define grading as a systematic change in the magnitude of the elements, we have to take into account exceptions to the pattern. Does an occasional zero element destroy the grading? What is a graded tridiagonal matrix?

2. Alternatively we can try to base a definition on the properties we observe in graded matrices — e.g., the possession of small, well determined eigenvalues. Unfortunately, matrices that are nicely graded in the informal sense can fail to have these properties.

In this paper we will take an intermediate course. We will define grading as a scaling of a base matrix B. Then we will determine the properties of B that insure that the graded matrix has the properties we want. In particular, we will be concerned with the structure and perturbation theory of the eigensystem of graded matrices.

The paper is organized as follows. We begin with a numerical example that illustrates some of the typical properties of a graded matrix. In the same section we will also give an example in which these properties fail. The heart of the paper is §3, where we establish the structure of the eigensystem of a graded matrix. In §4 we will derive condition numbers for the eigenvectors and eigenvectors of graded matrices. In §5 we will treat positive definite matrices and the singular value decomposition of general matrices. The paper concludes with bibliographical notes surveying previous and related work.

Throughout this paper, $\|\cdot\|$ will denote the Euclidean vector norm and its subordinate matrix norm. The conjugate transpose of a matrix A is denoted A^{H}. The reader is assumed to be familiar with the basic matrix decompositions, like Cholesky and the QR decompositions (see, e.g., [7, 13]). In partitioned matrices we will index each block by the indices of the element in the southeast corner. Thus if A is of order n, a partition of A in the form

$$A = \begin{pmatrix} A_{kk} & A_{kn} \\ A_{nk} & A_{nn} \end{pmatrix}$$

implies that A_{kk} is of order k.

2. Examples

Since the structure of the eigensystem of a graded matrix is not widely known, it is appropriate to set the stage with some examples. We begin with a matrix whose eigenvalues and eigenvectors exhibit the properties of a typical graded matrix. We will than present an example in which the properties fail. The computations were performed in MATLAB with rounding unit about 10^{-16}.

2.1. A typical graded matrix

The matrix

$$A = \begin{pmatrix} -6.5\text{e}{-01} & -5.0\text{e}{-05} & 4.4\text{e}{-09} & 4.1\text{e}{-14} & -9.8\text{e}{-17} \\ -1.1\text{e}{+00} & -3.6\text{e}{-06} & 1.3\text{e}{-08} & -7.6\text{e}{-13} & -6.9\text{e}{-17} \\ -4.8\text{e}{-02} & -1.7\text{e}{-05} & -5.0\text{e}{-09} & -8.9\text{e}{-14} & 1.3\text{e}{-16} \\ 3.8\text{e}{-01} & -9.6\text{e}{-05} & -1.1\text{e}{-08} & -2.0\text{e}{-12} & -9.1\text{e}{-17} \\ -3.3\text{e}{-01} & 1.3\text{e}{-04} & 8.1\text{e}{-09} & 1.1\text{e}{-12} & -4.1\text{e}{-17} \end{pmatrix} \tag{2.1}$$

was formed from a matrix B of standard normal deviates by postmultiplying it by

$$D = \text{diag}(1, 10^{-4}, 10^{-8}, 10^{-12}, 10^{-16}).$$

Thus A is column graded with grading ratio of 10^{-4} from column to column. (Here we only display two digits of the double precision representation of A.)

The eigenvalues of A are given by

$$-6.5\text{e}{-01} \quad 7.9\text{e}{-05} \quad -4.3\text{e}{-09} \quad -3.3\text{e}{-12} \quad -3.5\text{e}{-16}$$

It is seen that they share the grading of A, which is typical for such matrices.

The following is the matrix of eigenvectors of A, scaled so that the diagonal elements are one:

$$\begin{pmatrix} 1.0\text{e}{+00} & -7.6\text{e}{-05} & 1.2\text{e}{-08} & -1.4\text{e}{-12} & 3.8\text{e}{-16} \\ 1.7\text{e}{+00} & 1.0\text{e}{+00} & -7.0\text{e}{-05} & 1.4\text{e}{-08} & -3.7\text{e}{-12} \\ 7.3\text{e}{-02} & -1.8\text{e}{-01} & 1.0\text{e}{+00} & -5.5\text{e}{-05} & 3.6\text{e}{-08} \\ -5.8\text{e}{-01} & -1.6\text{e}{+00} & -2.5\text{e}{-02} & 1.0\text{e}{+00} & 9.9\text{e}{-07} \\ 5.0\text{e}{-01} & 2.0\text{e}{+00} & 1.1\text{e}{+00} & -9.0\text{e}{-01} & 1.0\text{e}{+00} \end{pmatrix}. \tag{2.2}$$

The behavior of these vectors is more complicated than the behavior of the eigenvalues. The subdiagonal elements are all of order one in magnitude. As we go upward from the diagonal, the components of the eigenvectors scale downward with ratios of about 10^{-4}. Once again this is typical behavior.

In the next section we will see that there is an intimate relation between the structure of the eigensystem of a graded matrix and the Schur complements of its leading principal submatrices. The following numbers illustrate this connection.

$$\begin{matrix} -6.5479\text{e}{-01} & 7.8905\text{e}{-05} & -4.3292\text{e}{-09} & -3.2932\text{e}{-12} & -3.5208\text{e}{-16} \\ -6.5471\text{e}{-01} & 7.8915\text{e}{-05} & -4.3292\text{e}{-09} & -3.2932\text{e}{-12} & -3.5208\text{e}{-16} \end{matrix}$$

The first row contains the eigenvalues of A, this time displayed to five figures. Below it are the diagonals from the U-factor of an unpivoted LU decomposition of A. The latter approximate the former to four or five figures. Since the U-factor is the triangular matrix computed by Gaussian elimination, the eigenvalues of a graded matrix can typically be approximated by performing Gaussian elimination on the matrix.

2.2. An atypical matrix

The matrix

$$\tilde{A} = \begin{pmatrix} -6.5\text{e}{-01} & -5.0\text{e}{-05} & 4.4\text{e}{-09} & 4.1\text{e}{-14} & -9.8\text{e}{-17} \\ -1.1\text{e}{+00} & -8.2\text{e}{-05} & 1.3\text{e}{-08} & -7.6\text{e}{-13} & -6.9\text{e}{-17} \\ -4.8\text{e}{-02} & -1.7\text{e}{-05} & -5.0\text{e}{-09} & -8.9\text{e}{-14} & 1.3\text{e}{-16} \\ 3.8\text{e}{-01} & -9.6\text{e}{-05} & -1.1\text{e}{-08} & -2.0\text{e}{-12} & -9.1\text{e}{-17} \\ -3.3\text{e}{-01} & 1.3\text{e}{-04} & 8.1\text{e}{-09} & 1.1\text{e}{-12} & -4.1\text{e}{-17} \end{pmatrix}$$

was obtained from A by altering its $(2,2)$-element. The eigenvalues of this matrix are

$$-6.5\text{e}{-01} \quad 9.2\text{e}{-07} \quad 7.8\text{e}{-08} \quad 5.2\text{e}{-12} \quad -1.0\text{e}{-15}$$

It is seen that the second and third eigenvalues of A no longer track the original grading.

The matrix of eigenvectors for $\tilde{A}$ is

$$\begin{pmatrix} 1.0\text{e}{+00} & -7.6\text{e}{-05} & 4.6\text{e}{-07} & 6.4\text{e}{-12} & 1.1\text{e}{-15} \\ 1.7\text{e}{+00} & 1.0\text{e}{+00} & -6.0\text{e}{-03} & -6.9\text{e}{-08} & -1.1\text{e}{-11} \\ 7.3\text{e}{-02} & -1.5\text{e}{+01} & 1.0\text{e}{+00} & 1.6\text{e}{-04} & 4.9\text{e}{-08} \\ -5.8\text{e}{-01} & -1.4\text{e}{+02} & 9.5\text{e}{+00} & 1.0\text{e}{+00} & 4.3\text{e}{-04} \\ 5.0\text{e}{-01} & 1.7\text{e}{+02} & -1.2\text{e}{+01} & -1.7\text{e}{+00} & 1.0\text{e}{+00} \end{pmatrix}.$$

The grading in the first, fourth and fifth columns is as above. However the subdiagonal elements in the second column are considerably larger than one and the grading of the superdiagonal elements in the third column is more gentle than above.

Finally, when we compare the the eigenvalues of $\tilde{A}$ with the diagonals of its U-factors we get the following table.

$-6.5479\text{e}{-01}$	$9.1704\text{e}{-07}$	$7.7527\text{e}{-08}$	$5.1728\text{e}{-12}$	$-1.0073\text{e}{-15}$
$-6.5471\text{e}{-01}$	$1.0000\text{e}{-06}$	$7.1214\text{e}{-08}$	$5.1612\text{e}{-12}$	$-1.0080\text{e}{-15}$

The second and third eigenvalues are not well approximated by the diagonals of U.

This last set of numbers has two features well worth noting. First, only the approximations to the second and third eigenvalues are affected. The diagonals of U provide good approximations to the first, fourth and fifth eigenvalues. Somehow the atypical behavior is localized.

Second, the number 1.0000e−06 (already suspect because of the string of zeros) is smaller than one would expect from performing Gaussian elimination on a random matrix scaled by D. This suggests that unusually small elements on the diagonal of U are associated with atypical behavior. We will make the connection clear in the next section.

3. The eigenstructure of a graded matrix

In this section we will describe the structure of the eigenvalues and eigenvectors of a graded matrix. The key idea is that when the grading is sufficiently strong, the matrix can be reduced by a similarity transformation to a block diagonal matrix. Moreover, as the grading increases, the reducing transformation approaches a fixed limit that is independent of the grading. By calculating the eigenvectors of the diagonal blocks of the block tridiagonal matrix we can compute approximations to the eigenvectors of the original matrix that amount to scaling certain essentially constant vectors.

We will begin this section with some definitions and observations. We will then show how to block-triangularize a graded matrix. We will then use the block triangular matrix to compute approximations to the eigensystem of the matrix. We conclude with an example of a gently graded matrix.

3.1. Definitions and observations

As we mentioned in the introduction, our approach to graded matrices amounts to grading a base matrix B and then determining what properties of B yield a tractable graded matrix. This approach leads to the following definition.

Definition 3.1. *Let B be a given base matrix of order n and let*

$$D = \operatorname{diag}(\delta_1, \delta_2, \ldots, \delta_n), \qquad \delta_1 \geq \delta_2 \geq \cdots \delta_n > 0.$$

Then

1. *$A = BD$ is column graded with respect to B and D,*
2. *$A = DB$ is row graded with respect to B and D,*
3. *$A = D^{\frac{1}{2}}BD^{\frac{1}{2}}$ is diagonally graded with respect to B and D,*

The numbers δ_k are called the grading factors. The numbers

$$\rho_k = \frac{\delta_{k+1}}{\delta_k}$$

are called the grading ratios.

There are four comments to make about this definition.

1. Because the grading factors decrease, we say the grading is downward. It is also possible to grade upward. Our results, derived here for downward grading, also apply with obvious modifications to upward grading.

2. Although we formally regard our graded matrices as coming from a base matrix B and a diagonal grading matrix D, in practice it will be the other way around. For example, given a column graded matrix A one might defined B by normalizing the columns of A and define D to consist of the reciprocals of the normalizing factors.

3. No particular assumption is made about the structure of B, although it is natural to think of it as being in some sense balanced. In particular, the results we are going to establish hold if B is a band matrix or a Hessenberg matrix.

4. The grading ratios are never greater than one, but they are allowed to be equal to one. Thus our theory applies to block-graded matrices, as well as the more conventional grading appearing in the first two sections.

An important observation is that the three types of grading in Definition 3.1 can be obtained from one another by diagonal similarities. For example if $A = BD$ is column graded, then $DAD^{-1} = DB$ is row graded. This means that we are free to chose a style of grading and stick to it through our analysis, after which the results can be transfered to the other styles. It turns out that column grading gives the cleanest derivations.

Our main result will be cast in terms of partitioned matrices and certain numbers obtained from these partitions. In particular, we introduce the following notation and terminology.

Definition 3.2. *Let $A = BD$ be partitioned in the form*

$$\begin{pmatrix} A_{kk} & A_{kn} \\ A_{nk} & A_{nn} \end{pmatrix} = \begin{pmatrix} B_{kk}D_k & B_{kn}D_n \\ B_{nk}D_k & B_{nn}D_n \end{pmatrix},$$

where A_{kk} is of order k. Then the number

$$\kappa_k \stackrel{\text{def}}{=} \|B_{kk}^{-1}\|\|B\|$$

is called the kth grading impediment. *The number*

$$\gamma_k \stackrel{\text{def}}{=} \kappa_k \rho_k$$

is called the kth grading coefficient.

The grading impediments get their name as follows. It will turn out that the behavior of graded matrices is controled by the size of the grading coefficients γ_k — the smaller the better. These coefficients can be made arbitrarily small by making the grading ratios sufficiently small. But if κ_k (which is never less than one) is large, the grading ratios will have to be correspondingly small for the the grading coefficients to be small. For this reason we call the numbers κ_k grading impediments.

3.2. Block triangularization

In this subsection we will be concerned with reducing A to block triangular form by a triangular similarity transformation. Specifically, we will try to find a matrix P_{nk} such that

$$\begin{pmatrix} I & 0 \\ -P_{nk} & I \end{pmatrix} \begin{pmatrix} A_{kk} & A_{kn} \\ A_{nk} & A_{nn} \end{pmatrix} \begin{pmatrix} I & 0 \\ P_{nk} & I \end{pmatrix} = \begin{pmatrix} A_{kk} + A_{kn}P_{nk} & A_{kn} \\ 0 & A_{nn} - P_{nk}A_{kn} \end{pmatrix}. \tag{3.1}$$

From elementary linear algebra, we know that the eigenvalues of A are, counting multiplicities, the union of the eigenvalues of $A_{kk} + A_{kn}P_{nk}$ and those of $A_{nn} - P_{nk}A_{kn}$. Moreover, it is easily verified that if y is an eigenvector of $A_{kk} + A_{kn}P_{nk}$ then

$$\begin{pmatrix} y \\ P_{nk}y \end{pmatrix} \tag{3.2}$$

is an eigenvector of A and conversely.

More generally, it follows from (3.1) that

$$A \begin{pmatrix} I \\ P_{nk} \end{pmatrix} = \begin{pmatrix} I \\ P_{nk} \end{pmatrix} (A_{kk} + A_{kn}P_{nk}). \tag{3.3}$$

We say that the columns of $(I\ \ P_{nk}^{\mathrm{H}})^{\mathrm{H}}$ span an eigenspace (or invariant subspace) of A and that $A_{kk} + A_{kn}P_{nk}$ is the representation of A on that subspace. In particular, if $k = 1$, then $(1\ \ p_{n1}^{\mathrm{H}})^{\mathrm{H}}$ is an eigenvector of A corresponding the the eigenvalue $a_{11}+a_{1n}^{\mathrm{H}}p_{n1}$.

Turning now to the existence of P_{nk}, if we write out the $(2,1)$-block of the right-hand side of (3.1) and set the result to zero, we get the equation

$$P_{nk}A_{kk} - A_{nn}P_{nk} = A_{nk} - P_{nk}A_{kn}P_{nk}.$$

Assuming that A_{kk} is nonsingular, we can write this equation in the form

$$P_{nk} - A_{nn}P_{nk}A_{kk}^{-1} = A_{nk}A_{kk}^{-1} - P_{nk}A_{kn}P_{nk}A_{kk}^{-1},$$

or in terms of B and D

$$P_{nk} - B_{nn}D_nP_{nk}D_k^{-1}B_{kk}^{-1} = B_{nk}B_{kk}^{-1} - P_{nk}B_{kn}D_nP_{nk}D_k^{-1}B_{kk}^{-1}. \tag{3.4}$$

This equation already exhibits the asymptotic form of P_{nk} as the granding ratio ρ_k approaches zero. Specifically, we have

$$\|B_{nn}D_nP_{nk}D_k^{-1}B_{kk}^{-1}\| \le \|B_{kk}^{-1}\|\|B_{nn}\|\|D_k^{-1}\|\|D_n\|\|P_{nk}\| \le \kappa_k\rho_k\|P_{nk}\| = \gamma_k\|P_{nk}\|.$$

Thus the second term on the left-hand side of (3.4) vanishes as $\rho_k \to 0$. Similarly,

$$\|P_{nk}B_{kn}D_nP_{nk}D_k^{-1}B_{kk}^{-1}\| \le \gamma_k\|P\|^2,$$

so that the second term on the right-hand side vanishes as $\rho_k \to 0$. We are left with the equation

$$P_{nk} \cong B_{nk}B_{kk}^{-1}.$$

This remarkable approximation says that as the kth grading ratio approaches zero the block diagonalizing similarity transformation in (3.1) effectively depends only on the base matrix and not the grading. It also says that the norm of P_{nk} is asymptotically bounded by the kth grading impediment.

Regarding the existence of P, the equation (3.4) is nonlinear and cannot be solved in closed form. Fortunately, similar equations appear in the perturbation theory of eigenspaces, and the analyses contained in that literature can be adapted to prove the following theorem.[1]

Theorem 3.3. *If*

$$\frac{\gamma_k\|B_{nk}B_{kk}^{-1}\|}{(1-\gamma_k)^2} < \frac{1}{4},$$

then (3.4) *has a unique solution satisfying*

$$\|P_{nk}\| \le \frac{2\|B_{nk}B_{kk}^{-1}\|}{1-\gamma_k} \le \frac{2\kappa_k}{1-\gamma_k}.$$

Moreover,

$$\frac{\|B_{nk}B_{kk}^{-1} - P\|}{\|P\|} \le \gamma_k(1+\|P\|). \tag{3.5}$$

Here are some observations on this theorem.

1. The theorem is local, depending only on the grading from k to $k+1$—i.e., γ_k.

[1]Specifically, in Theorem V.2.1 in [14] take $T = P_{nk} \mapsto P_{nk} - B_{nn}D_nP_{nk}D_k^{-1}B_{kk}^{-1}$, $g = B_{nk}B_{kk}^{-1}$, and $\varphi(P_{nk}) = P_{nk}B_{kn}D_nP_{nk}D_k^{-1}B_{kk}^{-1}$.

2. The bound (3.5) quantifies the fact that P_{nk} is approximated by $B_{nk}B_{kk}^{-1}$. Specifically, the bound on the normwise relative error in P_{nk} is proportional to the grading coefficient γ_k.

3. The matrix

$$A_{kk} + A_{kn}P_{nk} = B_{kk}D_k + B_{kn}D_nP_{kn}$$

contains the eigenvalues of A corresponding to the eigenspace spanned by $(I\ P_{nk}^{\mathrm{H}})^{\mathrm{H}}$. As γ_k decreases, the second term on the right becomes insignificant compared to the first. In other words, the eigenvalues of A_{kk} provide approximations to the largest k eigenvalues of A.

4. The matrix

$$A_{nn} - P_{nk}A_{kn} = (B_{nn} - P_{nk}B_{kn})D_n.$$

contains the remaining eigenvalues of A. As γ_k decreases, this matrix approaches

$$(B_{nn} - B_{nk}B_{kk}^{-1}B_{kn})D_n = A_{nn} - A_{nk}A_{kk}^{-1}A_{kn}.$$

The right hand side is the Schur complement of A_{kk} in A, which therefore contains approximations to the $n-k$ smallest eigenvalues of A.

5. Since the subspace spanned by the columns of a matrix does not change when the matrix is postmultiplied by a nonsingular matrix, the matrix

$$\begin{pmatrix} I \\ P_{nk} \end{pmatrix} B_{kk}D_k$$

spans an eigenspace of A. As γ_k decreases, this matrix approaches

$$\begin{pmatrix} A_{kk} \\ A_{nk} \end{pmatrix}.$$

Thus the span of the first k columns of A approximates an eigenspace of A.

In the special case where $k = 1$, and γ_1 is small, it follows from the above results that a_{11} is an approximate eigenvalue of A whose eigenvector is approximately the first column of A. To say something about the other eigenvalues and eigenvectors we must perform a further reduction, to which we now turn.

3.3. Eigenvalues and eigenvectors

If γ_k is sufficiently small, we can compute the first k eigenvalues of A from the matrix

$$C_{kk} = A_{kk} + A_{kn}P_{nk}$$

Moreover, any eigenvector y of C_{kk} can be converted into an eigenvector of A via the formula (3.2). We will be interested in the eigenvector corresponding to the kth eigenvalue λ_k.

To calculate y we first note that

$$C_{kk} = (B_{kk} + B_{kn}D_nP_{nk}D_k^{-1})D_k \tag{3.6}$$

is graded by columns with respect to $B_{kk} + B_{kn}D_nP_{nk}D_k^{-1}$ and D_k. Partition

$$C_{kk} = \begin{pmatrix} C_{k-1,k-1} & c_{k-1,k} \\ c_{k,k-1}^{\mathrm{H}} & c_{kk} \end{pmatrix},$$

and let $\tilde{\gamma}_{k-1}$ be the $(k-1)$th grading coefficient of C_{kk}. Then if $\tilde{\gamma}_{k-1}$ is sufficiently small, there is a vector $q_{k,k-1}^{\mathrm{H}}$ such that

$$\begin{aligned} &\begin{pmatrix} I & 0 \\ -q_{k,k-1}^{\mathrm{H}} & 1 \end{pmatrix} \begin{pmatrix} C_{k-1,k-1} & c_{k-1,k} \\ c_{k,k-1}^{\mathrm{H}} & c_{kk} \end{pmatrix} \begin{pmatrix} I & 0 \\ q_{k,k-1}^{\mathrm{H}} & 1 \end{pmatrix} \\ &\qquad = \begin{pmatrix} C_{k-1,k-1} + c_{k-1,k}q_{k,k-1}^{\mathrm{H}} & c_{k-1,k} \\ 0 & c_{kk} - q_{k,k-1}^{\mathrm{H}}c_{k-1,k} \end{pmatrix}. \end{aligned}$$

The quantity

$$\lambda_k = c_{kk} - q_{k,k-1}^{\mathrm{H}}c_{k-1,k}$$

is the kth eigenvalue of A. The corresponding eigenvector of C_{kk} is

$$y = \begin{pmatrix} (\lambda_k I - C_{k-1,k-1} - c_{k-1,k}q_{k,k-1}^{\mathrm{H}})^{-1}c_{k-1,k} \\ 1 + q_{k,k-1}^{\mathrm{H}}(\lambda_k I - C_{k-1,k-1} - c_{k-1,k}q_{k,k-1}^{\mathrm{H}})^{-1}c_{k-1,k} \end{pmatrix}.$$

Now it follows from (3.6) that as $\gamma_k \to 0$, we have $C_{kk} \to B_{kk}D_k$. Hence for γ_k small we have $\tilde{\gamma}_{k-1} \le \alpha\gamma_{k-1}$ for some constant α near one. Hence as $\gamma_{k-1} \to 0$, we have

$$q_k^{\mathrm{H}} = c_{k,k-1}^{\mathrm{H}}C_{k-1,k-1}^{-1} + O(\gamma_{k-1})$$

$$\lambda_k = c_{kk} - c_{k,k-1}^{\mathrm{H}}C_{k-1,k-1}^{-1}c_{k-1,k} + \delta_k O(\gamma_{k-1}).$$

and

$$y = \begin{pmatrix} -C_{k-1,k-1}^{-1} c_{k-1,k} + \delta_k D_{k-1}^{-1} O(\gamma_{k-1}) \\ 1 + O(\gamma_{k-1}) \end{pmatrix}.$$

From (3.2), the kth eigenvector of A is given by

$$x_k = \begin{pmatrix} -C_{k-1,k-1}^{-1} c_{k-1,k} + \delta_k D_{k-1}^{-1} O(\gamma_{k-1}) \\ 1 + O(\gamma_{k-1}) \\ -P_{n,k-1} C_{k-1,k-1}^{-1} c_{k-1,k} + p_{n,k} + O(\gamma_{k-1}) \end{pmatrix}$$

We now observe that as $\gamma_k \to 0$, $C_{k-1,k-1} D_{k-1}^{-1} \to B_{k-1,k-1}$, and similarly for the other components in the partition of C_{kk}. Hence we have the following theorem.

Theorem 3.4. *As γ_{k-1} and γ_k approach zero, we have*

$$\lambda_k = \delta_k (b_{kk} - b_{k,k-1}^{\mathrm{H}} B_{k-1,k-1}^{-1} b_{k-1,k}) + \delta_k O(\max\{\gamma_{k-1}, \gamma_k\}) \tag{3.7}$$

and

$$x_k = \begin{pmatrix} -\delta_k D_{k-1}^{-1} B_{k-1,k-1}^{-1} b_{k-1,k} \\ 1 \\ B_{nk} B_{kk}^{-1} \mathbf{e}_k \end{pmatrix} + \begin{pmatrix} \delta_k D_{k-1}^{-1} O(\max\{\gamma_{k-1}, \gamma_k\}) \\ 0 \\ O(\max\{\gamma_{k-1}, \gamma_k\}). \end{pmatrix}, \tag{3.8}$$

where $\mathbf{e}_k$ the last column of the $k \times k$ identity matrix.

The expressions in Theorem 3.4 represent a nice division of labor. The matrix B determines the unscaled structure of the eigenvalue and eigenvector; the matrix D determines their scale. We will exploit this division of labor in the next section, where we determine condition numbers of eigenvalues and eigenvectors.

The expressions confirm the observations made in §2.1 and §2.2. Their validity depends only on the sizes of the local grading coefficients γ_{k-1} and γ_k. The approximation (3.7) to λ_k is δ_k times the Schur complement of $B_{k-1,k-1}$ in B_{kk} — precisely the kth diagonal element of the U-factor in the LU decomposition of A. The approximation (3.8) to the eigenvector has the scaling shown in (2.2). In fact the approximate eigenvector can be quite good. For example, here is x_3 from the matrix (2.1) compared with its approximation.

x_3	approximation
1.2086e−08	1.2087e−08
−7.0094e−05	−7.0097e−05
1.0000e+00	1.0000e+00
−2.4855e−02	−2.4947e−02
1.1495e+00	1.1496e+00

With the exception of the unusually small fourth component, the vectors agree to four figures, which is consonant with the grading ratios $\gamma_2 = 1.1{\cdot}10^{-3}$ and $\gamma_3 = 1.3{\cdot}10^{-3}$ for this example.

If A is real, then all the quantities in (3.7) are real, and consequently λ_k is real. Moreover, if we allow nonpositive scaling factors, we can change the sign of λ_k, or even make it complex with any argument we wish. In particular, a general complex matrix whose leading principal submatrices are nonsingular can be graded so that all its eigenvalues are real (Fisher and Fuller [6]).

The grading coefficients for A^{H} are the same as for A. Consequently, the left eigenvectors of A are as well (or ill) behaved as the right eigenvectors. However, the approximation (3.8) is no longer valid, since A^{H} is graded by rows. However, the correct approximation can be recovered by transforming (3.8) into an approximation suitable for a matrix graded by rows; i.e., by multiplying it by D.

3.4. Gentle grading

In the foregoing we have assumed that the local grading coefficients were sufficiently small to allow the matrix to split as in Theorem 3.3. But even when the grading is gentle, the structure we have described persists, albeit in a rough way. Figure 3.1 contains a mesh plot of the common logarithms of the absolute values of the components of the eigenvectors of a matrix obtained by column grading a random matrix of order 100 with grading ratios of 0.69 (i.e., grading factors running from 1 to 10^{-16}). As above the first k components tend to be constant and then the components show a decrease at a rate determined by the grading ratio. The behavior is not uniform: note the ridges formed groups of the eigenvectors. But the plot never deviates far from the normative behavior. Why this should be so is an open research question.

4. Condition Numbers

In this section we will derive approximate perturbation bounds for the eigenvalues and eigenvectors of a column graded matrix. The bounds themselves are certainly overestimates. But they give us reason to believe that, baring untoward circumstances, the small eigenvalues and the small components of their eigenvectors are determined to high relative accuracy.

An interesting fact that will emerge from our analysis is that the condition of the eigenvalues and eigenvectors of a graded matrix depends on the grading impediments, not the grading coefficients. Of course, the grading coefficients must be small enough for our approximations to be valid. But once they are, further reducing the grading coefficients by reducing the grading ratios has little effect on the condition of the eigenvalues and eigenvectors.

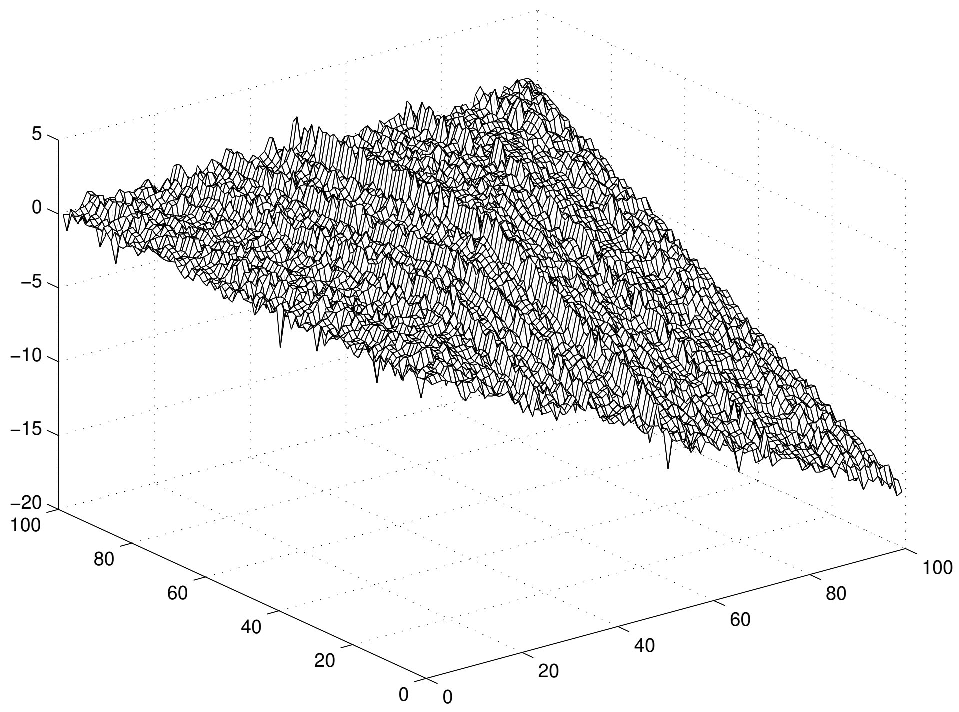

Figure 3.1: Eigenvectors of a gently graded matrix

4.1. Generalities

In order to derive condition numbers for graded matrices, we must first decide what it means to perturb a graded matrix. It seems natural that such a perturbation should itself be graded, we will adopt the following definition.

Definition 4.1. *Let $A = BD$ be a graded matrix and let $B + E$ be a perturbation of B. Then $BD + ED$ is a graded perturbation of A.*

Just as a graded matrix is generated by grading a base matrix, a graded perturbation is generated by grading a base perturbation of the base matrix. Our analysis will be entirely in terms of this base perturbation. Note that a graded perturbation need not represent a small componentwise perturbation in the elements of A since we do not exclude large relative perturbations in small elements of B.

As we shall see, it is easy to plug E into the expressions in Theorem 3.4 to get first order perturbation expansions. Taking norms in these expansions gives first order bounds from which we can derive putative condition numbers. Unfortunately, we are not working with eigenvalues and eigenvectors of A but with approximations to them, and these approximations may be in greater error than the first order error bounds we derive. It is therefore uncertain what such bounds actually mean.

To see what is going on, let us suppose we have a function $\varphi(\gamma, \epsilon)$, where $\varphi(0,0)$ represents our eigenvalue or eigenvector and $\varphi(\gamma, 0)$ represents its approximation. The quantities $\varphi(0, \epsilon)$ and $\varphi(\gamma, \epsilon)$ represents the perturbations of the original quantity and its approximation. Now a first order perturbation expansion gives

$$\varphi(0, \epsilon) \cong \varphi(0,0) + \varphi_\epsilon(0,0)\epsilon,$$

where φ_ϵ is the derivative of φ with respect to ϵ. What we actually compute is

$$\varphi(\gamma, \epsilon) \cong \varphi(\gamma, 0) + \varphi_\epsilon(\gamma, 0)\epsilon.$$

Depending on the size of ϵ the perturbations $\varphi_\epsilon(0,0)\epsilon$ and $\varphi_\epsilon(\gamma, 0)\epsilon$ may be far smaller than $\varphi(\gamma, 0) - \varphi(0,0)$. Nonetheless, if φ_ϵ is differentiable with respect to γ, $\varphi_\epsilon(\gamma, 0) - \varphi_\epsilon(0,0) = O(\gamma)$. Hence

$$\varphi_\epsilon(\gamma, 0)\epsilon - \varphi_\epsilon(0,0)\epsilon = O(\gamma\epsilon).$$

It follows that if γ is small enough, whether we base our perturbation theory on $\varphi(0,0)$ or $\varphi(\gamma, 0)$: they give essentially the same correction.

4.2. Eigenvalues

Since the approximation (3.7) for the eigenvalue is derived by multiplying the Schur complement

$$\mu_k = b_{kk} - b_{k,k-1}^{\mathrm{H}} B_{k-1,k-1}^{-1} b_{k-1,k}$$

by δ_k, it is sufficient to derive a perturbation expansion for the Schur complement. If E is partitioned conformally with B, the perturbed Schur complement becomes

$$\tilde{\mu}_k = b_{kk} + e_{kk} - (b_{k,k-1}^{\mathrm{H}} + e_{k,k-1}^{\mathrm{H}})(B_{k-1,k-1} + E_{k-1,k-1})^{-1}(b_{k-1,k} + e_{k-1,k}).$$

Replacing $(B_{k-1,k-1} + E_{k-1,k-1})^{-1}$ by the first order expansion

$$(B_{k-1,k-1} + E_{k-1,k-1})^{-1} = B_{k-1,k-1}^{-1} - B_{k-1,k-1}^{-1} E_{k-1,k-1} B_{k-1,k-1}^{-1}$$

and dropping second order terms, we get

$$\begin{aligned}\tilde{\mu}_k \cong \mu_k + e_{kk} + b_{k,k-1}^{\mathrm{H}} B_{k-1,k-1}^{-1} e_{k-1,k} + e_{k,k-1}^{\mathrm{H}} B_{k-1,k-1} \\ - b_{k,k-1}^{\mathrm{H}} B_{k-1,k-1}^{-1} E_{k-1,k-1} B_{k-1,k-1}^{-1} b_{k-1,k}.\end{aligned}$$

Taking norms we get and bounding terms like $B_{k-1,k-1}^{-1} b_{k-1,k}$ by κ_{k-1}, we get

$$|\tilde{\mu}_k - \mu_k| \lesssim (1 + \kappa_{k-1})^2 \|E\|.$$

Hence

$$\frac{|\tilde{\mu}_k - \mu_k|}{|\mu_k|} \lesssim (1 + \kappa_{k-1})^2 \frac{\|B\|}{|\mu_k|} \frac{\|E\|}{\|B\|}.$$

Multiplying the numerator and denominator of the left-hand side of this relation by δ_k, we get

$$\frac{|\tilde{\lambda}_k - \lambda_k|}{|\lambda_k|} \lesssim (1 + \kappa_{k-1})^2 \frac{\|B\|}{|\mu_k|} \frac{\|E\|}{\|B\|}.$$

This bound shows that the relative condition of λ_k is governed by two factors. The first is essentially the square of the grading impediment κ_{k-1}. The second is the ratio of $\|B\|$ to the kth diagonal element of the U-factor of B. Since μ_k^{-1} is the (k,k)-element of B_{kk}^{-1}, this ratio is bounded by κ_k.

The second factor has an interesting interpretation. In an ordinary ungraded eigenvalue problem, a small eigenvalue, even if it is well condition in an absolute sense, will be ill-conditioned in a relative sense. An analogous phenomenon holds for eigenvalues of graded matrices, but it is not the size of the eigenvalue that determines the ill-conditioning but the size of the Schur complement μ_k with respect to the base matrix.

4.3. Eigenvectors

Bounds for eigenvectors are complicated by the fact that the expression (3.8) has two distinct formulas. We begin by writing

$$x_k = \begin{pmatrix} -\delta_k D_{k-1}^{-1} y \\ 1 \\ z \end{pmatrix},$$

where $y = B_{k-1,k-1}^{-1} b_{k-1,k}$ and $z = B_{nk} B_{kk}^{-1} \mathbf{e}_k$.

The perturbation of y is the same as the perturbation of the system $B_{k-1,k-1} y = b_{k-1,k}$. We can therefore use standard perturbation theory [13, §3.3.1]to get

$$\frac{\|\tilde{y} - y\|}{\|y\|} \lesssim \kappa_{k-1} \left(\frac{\|E_{k-1,k-1}\|}{\|B_{k-1,k-1}\|} + \frac{\|e_{k-1,k}\|}{\|b_{k-1,k}\|} \right).$$

Since the ith component of x_i is approximated by $x_i^{(k)} \cong -\delta_k \delta_i^{-1} y_i$, we have

$$\frac{|\tilde{x}_i^{(k)} - x_i^{(k)}|}{\delta_k^{-1} \delta_i \|y\|} \lesssim \kappa_{k-1} \left(\frac{\|E_{k-1,k-1}\|}{\|B_{k-1,k-1}\|} + \frac{\|e_{k-1,k}\|}{\|b_{k-1,k}\|} \right).$$

This bound has the following interpretation. The number κ_k is the relative condition number for all components of the upper half of the eigenvector for which y_i is not much smaller as $\|y\|$. However, as y_i becomes smaller, its relative accuracy deteriorates.

The perturbation expansion for z does not simplify as nicely as that of y. A straightforward analysis yields the following bound:

$$\frac{\|\tilde{z} - z\|}{\|z\|} \lesssim (1 + \kappa_k) \frac{\|b_{kk}^{(-1)}\| \|B\|}{\|z\|} \frac{\|E\|}{\|B\|},$$

where $b_{kk}^{(-1)}$ denotes the kth column of B_{kk}^{-1}. There is no need to break this bound into components, since z is not graded. The first factor in this bound is essentially the kth grading impediment. The second factor, which is bounded by $\kappa_k / \|z\|$, has the following interpretation. Since $z = B_{nk} b_{kk}^{(-1)}$, we have $\|z\| \leq \|B_{nk}\| \|b_{kk}^{(-1)}\|$, so that the factor is always greater than one. It is much greater than one when z is atypically small; i.e., when it does not reflect the size of $b_{kk}^{(-1)}$. As with the bound for y, only the larger components of z are determined with high relative accuracy.

Most bounds for eigenvectors, whether normwise or componentwise, invoke a gap hypothesis that says that the eigenvalue in question is sufficiently separated from its neighbors. No explicit gap appears in our expressions. The reason is that we have assumed that γ_{k-1} and γ_k are small. This forces the eigenvalue λ_k to be well enough separated from its neighbors for our bounds to hold.

5. Positive definite matrices and the singular value decomposition

In this section we will show how our theory applies to positive definite matrices. We will then use these results to describe the behavior of the singular value decomposition of a graded matrix.

5.1. Positive definite matrices

In treating positive definite matrices, it is natural to pass on the symmetry (and positive-definiteness) of B to A by grading B by diagonals so that $A = D^{\frac{1}{2}}BD^{\frac{1}{2}}$. The expression (3.8) for the eigenvectors must then be multiplied by $D^{\frac{1}{2}}$. When the kth component of x_k is normalized to one, the result is

$$x_k = \begin{pmatrix} -\delta_k^{\frac{1}{2}} D_{k-1}^{-\frac{1}{2}} B_{k-1,k-1}^{-1} b_{k-1,k} \\ 1 \\ \delta_k^{-\frac{1}{2}} D_n^{\frac{1}{2}} B_{nk} B_{kk}^{-1} \mathbf{e}_k \end{pmatrix} + \begin{pmatrix} \delta_k^{\frac{1}{2}} D_{k-1}^{-\frac{1}{2}} O(\max\{\gamma_{k-1}, \gamma_k\}) \\ 0 \\ \delta_k^{-\frac{1}{2}} D_n^{\frac{1}{2}} O(\max\{\gamma_{k-1}, \gamma_k\}) \end{pmatrix}. \tag{5.1}$$

Thus, when the grading ratios are constant, each eigenvector x_k exhibits constant grading downward above and below its kth component.

The grading impediments of a graded positive definite matrix are better behaved that those of a general graded matrix. Because of the interlacing properties of the eigenvalues of symmetric matrices, we have

$$\|B_{11}^{-1}\| \leq \|B_{22}^{-1}\| \leq \cdots \leq \|B_{nn}^{-1}\|.$$

Hence the grading impediments κ_k are nondecreasing and are bounded by κ_n. In particular, graded positive definite matrices cannot exhibit the intermediate ill behavior found in §2.2.

There is also a computational difference between graded positive definite matrices and graded general matrices. We have seen that for a general graded matrix A the eigenvalues corresponding to sufficiently small grading coefficients are approximated by the diagonal elements of the U-factor of A. Unfortunately, this U-factor must be computed by Gaussian elimination without pivoting, which will be unstable if any of the grading impediments are large. With positive definite matrices pivoting is unnecessary for a stable reduction.

5.2. The singular value decomposition

In this subsection we will derive the structure of the singular value decomposition of graded matrices. For definiteness we will consider an $m\times n$ ($m \geq n$) base matrix Y and a column graded matrix $X = YD^{\frac{1}{2}}$. (Results for row graded matrices can be obtained by considering the transpose matrix.) We will write the singular value decomposition of X as

$$X = U\Sigma V^{\mathrm{H}},$$

where $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_n)$ $(\sigma_1 \geq \cdots \geq \sigma_n \geq 0)$ and U and V are orthonormal. The columns u_i of U are called the right singular vectors of X and the columns v_i of V are called the left singular vectors of X.

We will be chiefly concerned with a qualitative description of the structure of the graded singular value decomposition; however, formulas and bounds can easily be obtained from our previous results. The key observation is that the squares of the singular values of X are the eigenvalues of $A = X^{\mathrm{H}}X$ and the right singular vectors of X are the eigenvectors of A. Moreover if v_k is a singular vector of X corresponding to the singular value σ_k, then $u_k = \sigma_k^{-1} X v_k$.

Now let $B = Y^{\mathrm{H}}Y$, so that $A = D^{\frac{1}{2}} B D^{\frac{1}{2}}$, and let $B = S^{\mathrm{H}}S$ be the Cholesky factorization of B. Assuming that the grading coefficients γ_{k-1} and γ_k of A are sufficiently small, we have the following results.

1. The square of the kth singular value of X is approximately

$$\delta_k s_{kk}^2 + \delta_k O(\max\{\gamma_{k-1}, \gamma_k\}).$$

   It follows that

$$\sigma_k = \delta_k^{\frac{1}{2}}(s_{kk} + O(\max\{\gamma_{k-1}, \gamma_k\})).$$

2. The left singular vector v_k has the structure given in (5.1).

To determine the behavior of the right singular vectors, we will exploit a connection between the singular value decomposition of a graded matrix and its QR decomposition. Specifically, suppose that $\gamma_k \to 0$. Then the columns of $(I\ P_{nk}^{\mathrm{H}})^{\mathrm{H}}$ [see (3.3)] span the eigenspace of A corresponding to the first k right singular vectors. It follows that the columns of

$$X_{nk} + X_{nn}P_{nk} = Y_{nk}D_k + Y_{nk}D_n B_{nk} B_{kk}^{-1} + O(\gamma_k)$$

span the space $\mathcal{U}_k$ spanned by the first k right singular vectors of X. Postmultiplying by D_k^{-1}, we find that the columns of

$$Y_{nk} + O(\gamma_k)$$

span the same subspace. Thus, in the limit $\mathcal{U}_k$ is the column space of Y_{nk} or X_{nk}.

Now suppose that γ_{k-1} also approaches zero. Then $\mathcal{U}_{k-1}$ is well approximated by the column space of $Y_{n,k-1}$. Since $\mathcal{U}_k$ obtained by appending u_k to $\mathcal{U}_{k-1}$, up to terms of order $\max\{\gamma_{k-1}, \gamma_k\}$ the vector u_k must be the result of orthogonalizing the kth column of Y against $Y_{n,k-1}$. This is just the kth vector in the orthogonal part of the QR factorization of Y or X.

Recalling that the R-factor R in the QR decomposition of X is the Cholesky factor of $X^{\mathrm{T}}X$ and that $R = SD^{\frac{1}{2}}$, we have the following theorem.

Theorem 5.1. *Let $X = QR$ be the QR factorization of X. If γ_{k-1} and γ_k are sufficiently small then*

$$\sigma_k = r_{kk} + \delta_k^{\frac{1}{2}} O(\max\{\gamma_{k-1}, \gamma_k\}),$$

and

$$u_k = q_k + O(\max\{\gamma_{k-1}, \gamma_k\}).$$

We have scaled X by $D^{\frac{1}{2}}$ to retain consistency with our earlier results. However, it is the γ_i, which are proportional to the δ_i, that control the convergence of our approximations in Theorem 5.1. Thus with respect to actual gradings, approximations for the singular value decomposition converge faster than approximations for the eigenvalue problem.

To illustrate this phenomenon consider the following matlab code.

```
%Y = randn(6,2);
err = [];
for l = 1:5
   D = diag(logspace(0,-l,2));
   X = Y*D;
   [Q, R] = qr(X);
   [U,S,V] = svd(X);
   U(:,2) = U(:,2)/sign(U(1,2));
   Q(:,2) = Q(:,2)/sign(Q(1,2));
   err = [err;
          [norm(U(:,2)-Q(:,2)), abs(abs(R(2,2))-S(2,2))/S(2,2)]];
end
```

It generates a random base matrix, successively scales the second column by 10^{-1} through 10^{-5}, and computes the error in the QR approximation to the second right singular vector and the corresponding singular value. The array `err` is

5.9262e−03	2.9763e−03
5.9264e−05	2.9633e−05
5.9264e−07	2.9632e−07
5.9264e−09	2.9632e−09
5.9264e−11	2.9631e−11

We see that the approximations are converging as 100^{-i}. The fact that the ratios of these errors quickly become constant suggests that the second order terms in γ are also converging.

6. Bibliographical notes

Problems involving scaling matrices by diagonal elements have a long history in modern matrix computations. Early work was directed to the effects of scaling on the condition of the matrix in question and the effects of rounding error in Gaussian elimination. Although this work does not concern us directly, it is appropriate to mention the seminal papers by Bauer [3, 1966], van der Sluis [16, 1969], and Skeel [12, 1979].

In 1958 Fisher and Fuller [6, 1958] showed that if the leading principle submatrices of a matrix B are nonsingular, there is a diagonal matrix D such that the eigenvalues of DB are positive. Although they do not mention grading explicitly, their construction amounts to chosing grading ratios so large that the eigenvalues of the resulting matrix are real. By allowing D to have negative elements, the eigenvalues can be made positive. Later Ballantine [1, 1970] gave a simple proof of the theorem.[2]

The first reference I can find to graded matrices as such is by Martin, Reinsch, and Wilkinson [8, 1968], who warned that their version of Householder tridiagonalization would destroy the accuracy of the small eigenvalues of a downward graded matrix. However, the analysis of the eigensystem of graded matrices began with Dalquist [4, 1985], whose application was to stiff ordinary equations. He introduces, block grading in terms of a base matrix, and shows that under certain conditions a graded matrix A can be written in the form $A = LRL^{-1}$, where L and R are close to the (block) L- and U-factors of A. He establishes this fact using a block LR-algorithm ([17, Ch. 8]) to triangularize A. By accumulating the transformations he gets error bounds on his approximations.

The idea of grading a base matrix was rediscovered by Barlow and Demmel [2, 1990] and Stewart and Zhang [15, 1991]. The latter paper, which like Dalquist's dealt

[2] Fisher and Fuller had in mind the solution of the linear system $Bx = c$ by an iterative method of the form $x_{k+1} = (I - DB)x_k + Dc$. If w choose D so that the eigenvalues of the iteration matrix lie in $[0, 2)$, the iteration matrix has spectral radius less than one. However, if the grading of B is strong, $I - DB$ will have eigenvalues very near one, and the iteration will converge very slowly.

only with eigenvalues, established its result by a direct block triangularization of the kind described here. This paper also introduced (though not by name) the grading impediments κ_k. Mathias [9, 1996] gave eigenvalue and eigenvector bounds for positive definite matrices and at at the thirteenth Householder Symposium in Pontresina (1996) observed, without proof that similar results hold for general graded matrices.

When the base matrix has special structure — e.g., when it is symmetric diagonally dominant or positive definite — different kinds of bounds can be obtained. This line of investigation was initiated by Barlow and Demmel [2, 1990] and continued by Demmel and Veselić [5, 1992], Mathias and Stewart [11, 1993], and Mathias [9, 1996] [10, 1997]. A typical result for eigenvalues is the following [5, Theorem 2.3].

Theorem 6.1. *Let B be positive definite with unit diagonal elements, and let $A = D^{\frac{1}{2}}BD^{\frac{1}{2}}$. Let $\|F\| \leq \lambda_{\min}(B)$. Let λ_i be the ith eigenvalue of A (in descending order), and let $\tilde{\lambda}_i$ be the ith eigenvalue of $A+E$, where $E = D^{\frac{1}{2}}FD^{\frac{1}{2}}$. Then if $\|F\| < \lambda_{\min}(B)$,*

$$\frac{|\tilde{\lambda}_i - \lambda_i|}{\lambda_i} \leq \frac{\|E\|}{\lambda_{\min}(B)}. \tag{6.1}$$

In addition to perturbation bounds for eigenvalues, this sequence of papers develops componentwise perturbation theory for eigenvectors and for the singular value decomposition.

It is evident that Theorem 6.1 has a different flavor from the expansions and bounds derived in this paper. It is global and simple in the sense that one bound serves all eigenvalues, whereas our condition numbers vary with the eigenvalue. The price to be paid for this simplicity is that it can be quite pessimistic. Our analysis makes it clear that the sensitivity of an eigenvalue depends only on the local grading impediments; whereas the reciprocal of $\lambda_{\min}(B)$ in the bound (6.1) represents the largest grading impediment. Thus an open problem in the perturbation theory of graded positive definite matrices is to derive bounds in the spirit of (6.1) that are local in nature.

References

[1] C. S. Ballantine. Stabilization by a diagonal matrix. *Proceedings of the American Mathematical Society*, 25:728–734, 1970.

[2] J. Barlow and J. Demmel. Computing accurate eigensystems of scaled diagonally dominant matrices. *SIAM Journal on Numerical Analysis*, 27:762–791, 1988.

[3] F. L. Bauer. Genauigkeitsfragen bei der Lösung linear Gleichungssysteme. *Zeitschrift für angewandte Mathematik und Mechanik*, 46:409–421, 1966.

[4] G. Dahlquist. On transformations of graded matrices, with applications to stiff ODE's. *Numerische Mathematik*, 47:363–385, 1995.

[5] J. Demmel and K. Veselić. Jacobi's method is more accurate than QR. *SIAM Journal on Matrix Analysis and Applications*, 13:1204–1245, 1992.

[6] M. E. Fisher and A. T. Fuller. On the convergence of matrices and the convergence of linear iterative processes. *Proceedings of the Cambridge Philosophical Society*, 45:417–425, 1958.

[7] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, second edition, 1989.

[8] R. S. Martin, C. Reinsch, and J. H. Wilkinson. Householder tridiagonalization of a real symmetric matrix. *Numerische Mathematik*, 11:181–195, 1968. Also in [18, pp. 212–226].

[9] R. Mathias. Fast accurate eigenvalue methods for graded positive definite matrices. *Numerische Mathematik*, 74:85–104, 1996.

[10] R. Mathias. Spectral perturbation bounds for positive definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 18:959–980, 1997.

[11] R. Mathias and G. W. Stewart. A block QR algorithm and the singular value decomposition. *Linear Algebra and Its Applications*, 182:91–100, 1993.

[12] R. D. Skeel. Scaling for numerical stability in Gaussian elimination. *Journal of the ACM*, 26:494–526, 1979.

[13] G. W. Stewart. *Matrix Algorithms I: Basic Decompositions*. SIAM, Philadelphia, 1998.

[14] G. W. Stewart and J.-G. Sun. *Matrix Perturbation Theory*. Academic Press, New York, 1990.

[15] G. W. Stewart and G. Zhang. Eigenvalues of graded matrices and the condition of numbers of a multiple eigenvalue. *Numerische Mathematik*, 58:703–712, 1991.

[16] A. van der Sluis. Condition numbers and equilibration of matrices. *Numerische Mathematik*, 14:14–23, 1969.

[17] J. H. Wilkinson. *The Algebraic Eigenvalue Problem*. Clarendon Press, Oxford, 1965.

[18] J. H. Wilkinson and C. Reinsch. *Handbook for Automatic Computation. Vol. II Linear Algebra*. Springer-Verlag, New York, 1971.