Similarity Transformations

Definition Two matrices A and B are said to be *similar* if there exists a nonsingular matrix S such that

$$B = S^{-1} A\, S$$

Theorem A matrix A is similar to a diagonal matrix D if and only if A has n linearly independent eigenvectors. In this case $D = S^{-1} A\, S$ where the columns of S are the eigenvectors of A and the diagonal entries of D are the corresponding eigenvectors of A.

An important observations about similarity transformations is that they can be composed. For example, if A is similar to B and B is similar to C, then A is similar to C:

$$B = S_1^{-1} A\, S_1$$

$$C = S_2^{-1} B\, S_2 = S_2^{-1} S_1^{-1} A\, S_1\, S_2 = \left(S_1\, S_2\right)^{-1} A \left(S_1\, S_2\right)$$

This suggests a systematic strategy for finding all of the eigenvalues and eigenvectors of a matrix A: we seek a sequence of similarity transformations that systematically transform the original A into a diagonal matrix D. According the theorem above the diagonal entries of D will end up being the eigenvalues of A and the matrix that results from multiplying out the product of all of the similarity matrices will be the matrix whose columns are the eigenvectors of A.

The Householder transformation

One way to carry out this strategy starts by shooting for an intermediate stage in which we make the original matrix A similar to a tridiagonal matrix. Once we get to the tridiagonal matrix we get the rest of the way to diagonal by using an algorithm called the QR decomposition.

To get from the original A to a tridiagonal B we apply a series of applications of a *Householder transformation*. The following theorem sets up this transformation:

Definition If $\mathbf{w}$ is a vector whose l_2 norm is 1 the matrix

$$P = I - 2\, \mathbf{w}\, \mathbf{w}^t$$

is called a Householder transformation.

Theorem A Householder matrix P is symmetric and orthogonal, so $P^{-1} = P^t = P$.

One the first round of the process we seek a vector $\mathbf{w}^{(1)}$ and associated transformation matrix $P^{(1)}$ that has the effect of turning A into a matrix $A^{(1)}$ that has 0s in its first column from its third row to the last row. If the original A is symmetric the transformation will also produce 0s in the first row of the new matrix from the third column to the last column.

The textbook shows how to compute the correct $\mathbf{w}^{(1)}$: the result is summarized by

$$\alpha = -sgn(a_{21}) \left(\sum_{j=2}^{n} a_{j1}{}^{2} \right)^{1/2}$$

$$r = \left(\frac{1}{2}\, \alpha^2 - \frac{1}{2}\, a_{21}\, \alpha \right)^{1/2}$$

$$\mathbf{w}^{(1)}{}_1 = 0$$

$$\mathbf{w}^{(1)}{}_2 = \frac{a_{21} - \alpha}{2\, r}$$

$$\mathbf{w}^{(1)}{}_j = \frac{a_{j1}}{2\, r} \text{ for } j = 3, \ldots, n$$

This transforms A into $A^{(2)}$. The general step that transforms $A^{(k)}$ into $A^{(k+1)}$ uses a vector $\mathbf{w}^{(k)}$ whose components are computed by

$$\alpha = -sgn(a^{(k)}{}_{k+1,k}) \left(\sum_{j=k+1}^{n} a^{(k)}{}_{j,k}{}^{2} \right)^{1/2}$$

$$r = \left(\frac{1}{2}\, \alpha^2 - \frac{1}{2}\, a^{(k)}{}_{k+1,k}\, \alpha \right)^{1/2}$$

$$\mathbf{w}^{(k)}{}_1 = \cdots\ \mathbf{w}^{(k)}{}_k = 0$$

$$\mathbf{w}^{(k)}{}_{k+1} = \frac{a^{(k)}{}_{k+1,k} - \alpha}{2\, r}$$

$$\mathbf{w}^{(k)}{}_j = \frac{a^{(k)}{}_{j,k}}{2\, r} \text{ for } j = k+2, \ldots, n$$

After a series of Householder transformations we will have converted our original A into a tridiagonal $A^{(n)}$.

The QR method

The QR method is a similarity transformation that seeks to bring a symmetric, tridiagonal matrix closer to being a pure diagonal matrix.

The QR method starts with the concept of a rotation matrix. A rotation matrix P is a matrix that looks mostly like an identity matrix, with the exception of the entries $P_{i,i}$ $P_{i,i+1}$, $P_{i+1,i}$, and $P_{i+1,i+1}$, which take the form

$$\begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix}$$

The value of the parameter θ is selected to cause the entry in row $i+1$ and column i of the product $P\,A$ to vanish. By multiplying an original matrix A by a series of rotation matrices P_i we can produce a new matrix R whose below diagonal entries are all 0.

What makes all of this interesting to us is that we can turn this process into a similarity transformation:

$$R = P_{n-1}\,P_{n-2} \cdots P_2\,P_1\,A = Q^{-1}\,A$$

$$R\,Q = Q^{-1}\,A\,Q$$

Since rotation matrices are orthogonal matrices we can easily compute the matrix Q:

$$Q = \left(P_{n-1}\,P_{n-2} \cdots P_2\,P_1\right)^{-1} = P_1^{\,t}\,P_2^{\,t} \cdots P_{n-2}^{\,t}\,P_{n-1}^{\,t}$$

After a typical round of the QR method we will end up transforming our original matrix to a new matrix that is closer to being diagonal. In the limit of a large number of iterations of the method we get a diagonal matrix whose diagonal entries are the eigenvalues of the original matrix A.

Accelerating convergence

One final problem with the QR method is that although the off-diagonal entries of the matrix converge to 0, they may do so slowly. Furthermore, the rate at which the off-diagonal entry $A_{j+1,j}$ converges to 0 affects the rate at which the diagonal entry $A_{j,j}$ converges to the eigenvalue λ_j. It turns out that the rate at which $A_{j+1,j}$ converges to 0 is proportional to $|\lambda_j/\lambda_{j-1}|$. (Here we are assuming that the eigenvalues have been ordered in order of decreasing absolute value.)

If, for example, the ratio is $|\lambda_n/\lambda_{n-1}|$ is not small enough, we can employ a shifting trick to modify the eigenvalues of A to get a better ratio. In a shifting trick we introduce a factor σ and replace A with the matrix

$$A - \sigma\,I$$

The new matrix will have eigenvalues equal to $\lambda_j - \sigma$, so we will want to select σ to make the ratio

$$\frac{\lambda_n - \sigma}{\lambda_{n-1} - \sigma}$$

small. The book says that a commonly used strategy is to compute the eigenvalues of the small matrix

$$\begin{bmatrix} A_{n-1,n-1} & A_{n-1,n} \\ A_{n,n-1} & A_{n,n} \end{bmatrix}$$

and then select the eigenvalue of this matrix that is closest to $A_{n,n}$ in absolute value and then set σ equal to that eigenvalue.

After computing Q and R for the shifted matrix we then make the new matrix be

$$R\,Q + \sigma\,I$$

This final step undoes the shift and restores the eigenvalues of the matrix back to their original values.