Section 7.1: Diagonalization of Symmetric Matrices

April 23, 2008

Abstract

As we begin chapter seven (and finish up the semester!), we should keep track of our specific objectives (other than relaxing after finals). We've got two goals:

(a) to analyze the structure of general matrices of information (like LandSat images, say, as described in the opening pages of the chapter, p. 447, or like statistical data sets) – we'll do this via the Singular Value Decomposition, wavelets, etc.; and

(b) to examine the behavior of symmetric matrices (those that satisfy $A^T = A$) as linear transformations (it turns out that they're fundamental to goal (a)).

Great things happen when you work with symmetric matrices: their special structure leads to some seemingly magical properties. Symmetric matrices are an important special case, as we found in working with the least-squares problems (where the left-hand side was $A^T A$, a symmetric matrix!).

Theorem 1: If A is symmetric, then any two eigenvectors from different eigenspaces are orthogonal.

Comment: In the past, when a matrix had two distinct eigenvalues λ_1 and λ_2, we could conclude that the corresponding eigenvectors were independent – but we couldn't conclude that the eigenvectors were orthogonal.

Example: #13, p. 454

Definition: *orthogonally diagonalizable* A matrix is orthogonally diagonalizable if there is an orthogonal matrix P and diagonal matrix D such that

$$A = PDP^T$$

Example: #22, p. 454

Theorem 2: $A_{n \text{ x } n}$ is orthogonally diagonalizable if and only if A is a symmetric matrix.

Theorem 3 (The Spectral Theorem): Symmetric $A_{n \text{ x } n}$ has the following properties:

(a) A has n real eigenvalues, counting multiplicities (no complex eigenvalues!).

(b) The dimension of the eigenspace for each eigenvalue λ equals the multiplicity of λ as a root of the characteristic equation (no "missing" dimensions).

(c) The eigenspaces are mutually orthogonal: eigenvectors corresponding to different eigenvalues are orthogonal (no shadows cast on each other).

(d) A is orthogonally diagonalizable.

Let's look at the geometry of this: if we think about transforming the n-dimensional unit ball into an ellipsoid, then it turns out that the eigenvectors are the major/minor axes of the ellipsoid, and the eigenvalues are the stretch factors.

Example: #31, p. 455

$$A = \underbrace{\begin{bmatrix} \mathbf{u}_1 \cdots \mathbf{u}_n \end{bmatrix}}_{P} \underbrace{\begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix}}_{D} \underbrace{\begin{bmatrix} \mathbf{u}_1^T \\ \vdots \\ \mathbf{u}_n^T \end{bmatrix}}_{P^T}$$

Since $A = PDP^T$, where p is an orthogonal matrix, we can write

$$A = \lambda_1 \mathbf{u}_1 \mathbf{u}_1^T + \lambda_2 \mathbf{u}_2 \mathbf{u}_2^T + \ldots + \lambda_n \mathbf{u}_n \mathbf{u}_n^T,$$

the **spectral decomposition** of A. Each matrix $\mathbf{u}_j \mathbf{u}_j^T$ is a **projection matrix**: the projection of vector $\mathbf{x}$ onto the subspace spanned by $\mathbf{u}_j$ is given by

$$\text{proj}_{\mathbf{u}_j} \mathbf{x} = \mathbf{u}_j \mathbf{u}_j^T \mathbf{x} = (\mathbf{x} \cdot \mathbf{u}_j)\mathbf{u}_j$$

(the last part of the equation is one way of thinking of the projection that I've emphasized).

Example: #34, p. 455

$$A = \begin{bmatrix} 3 & -2 & 4 \\ -2 & 6 & 2 \\ 4 & 2 & 3 \end{bmatrix} = 7 \begin{bmatrix} 1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & 0 & 1/\sqrt{2} \end{bmatrix} + 7 \begin{bmatrix} -1/\sqrt{18} \\ 4/\sqrt{18} \\ 1/\sqrt{18} \end{bmatrix} \begin{bmatrix} -1/\sqrt{18} & 4/\sqrt{18} & 1/\sqrt{18} \end{bmatrix}$$

$$+ -2 \begin{bmatrix} -2/3 \\ -1/3 \\ 2/3 \end{bmatrix} \begin{bmatrix} -2/3 & -1/3 & 2/3 \end{bmatrix}$$

The action of A as a linear transformation is well understood, therefore:

$$A\mathbf{x} = \lambda_1 \mathbf{u}_1 \mathbf{u}_1^T \mathbf{x} + \lambda_2 \mathbf{u}_2 \mathbf{u}_2^T \mathbf{x} + \ldots + \lambda_n \mathbf{u}_n \mathbf{u}_n^T \mathbf{x},$$

or

$$A\mathbf{x} = (\lambda_1 \mathbf{u}_1^T \mathbf{x})\mathbf{u}_1 + (\lambda_2 \mathbf{u}_2^T \mathbf{x})\mathbf{u}_2 + \ldots + (\lambda_n \mathbf{u}_n^T \mathbf{x})\mathbf{u}_n.$$

That is, we project $\mathbf{x}$ onto each basis vector, and then multiply each of these projections by the corresponding eigenvalue. Alternatively, if

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix}_P$$

where P represents the basis composed of its columns, then

$$A\mathbf{x} = \begin{bmatrix} \lambda_1 x_1 \\ \lambda_2 x_2 \\ \vdots \\ \lambda_n x_n \end{bmatrix}_P$$

Neat!