Note 14: Symmetric Matrices and Spectral Theorem

1 Overview

In previous notes, we have explored the usefulness of diagonal and diagnolizable matrices. Another useful type of matrix is the **symmetric matrix**—matrices that are equal to their transpose. In this note, we will discuss the Spectral Theorem, which provides us with useful properties of symmetric matrices.

2 Spectral Theorem

Now it is time to discuss one of the most important and fundamental theorems in linear algebra.

Theorem 1 (Spectral Theorem for Real Symmetric Matrices)
Let $A \in \mathbb{R}^{n \times n}$ be real and symmetric. Then:

(i) The eigenvalues of A are real.

(ii) A is diagonalizable.

(iii) There is an orthonormal basis of $\mathbb{R}^n$ consisting of eigenvectors of A.

In short, A may be *orthonormally diagonalized*: $A = V \Lambda V^\top$ where $V \in \mathbb{R}^{n \times n}$ is an orthonormal matrix of eigenvectors of A, and $\Lambda \in \mathbb{R}^{n \times n}$ is a real diagonal matrix of eigenvalues.

Orthonormal diagonalization is a key ingredient in proofs and algorithms involving symmetric matrices.

Proof of Theorem 1 part (i).

Take an arbitrary eigenvector λ of A with corresponding eigenvector $\vec{v}$. Then

$$A\vec{v} = \lambda\vec{v}. \tag{1}$$

Taking the conjugate of both sides and using the fact that A is real so that $A = \overline{A}$, we get

$$A\overline{\vec{v}} = \overline{A}\,\overline{\vec{v}} = \overline{A\vec{v}} = \overline{\lambda\vec{v}} = \overline{\lambda}\,\overline{\vec{v}}. \tag{2}$$

Taking advantage of the fact that A is symmetric so that $A = A^\top$, we take the transpose of both sides to get

$$A\overline{\vec{v}} = \overline{\lambda}\,\overline{\vec{v}} \tag{3}$$

$$\overline{\vec{v}}^\top A^\top = \overline{\lambda}\,\overline{\vec{v}}^\top \tag{4}$$

$$\overline{\vec{v}}^\top A = \overline{\lambda}\,\overline{\vec{v}}^\top. \tag{5}$$

Then we multiply by $\vec{v}$ on both sides to get

$$\overline{\vec{v}}^\top A\vec{v} = \overline{\lambda}\,\overline{\vec{v}}^\top\vec{v} \tag{6}$$

$$\lambda \overline{\vec{v}}^\top \vec{v} = \overline{\lambda}\, \overline{\vec{v}}^\top \vec{v} \tag{7}$$

$$0 = (\lambda - \overline{\lambda}) \overline{\vec{v}}^\top \vec{v} \tag{8}$$

using the fact that $(\vec{v}, \lambda)$ is an eigenvector-eigenvalue pair of of A. Now

$$\overline{\vec{v}}^\top \vec{v} = \sum_{i=1}^{n} \overline{v_i} \cdot v_i = \sum_{i=1}^{n} |v_i|^2. \tag{9}$$

so $\overline{\vec{v}}^\top \vec{v}$ is nonzero if and only if $\vec{v}$ is nonzero. Since $\vec{v}$ is an eigenvector and thus nonzero, we know that $\overline{\vec{v}}^\top \vec{v}$ is nonzero, and thus that $\lambda - \overline{\lambda} = 0$. Thus $\lambda = \overline{\lambda}$ so λ is real. Since λ is an arbitrary eigenvalue, all eigenvalues of A are real.

□

2.0.1 Proof that Symmetric Matrices Have Orthogonal Eigenvectors

The interplay between symmetric matrices and orthogonal eigenvectors is a cornerstone in linear algebra. Firstly, a symmetric matrix $A \in \mathbb{R}^{n \times n}$ is guaranteed to have orthogonal eigenvectors. Let us try to prove this fact.

Consider a symmetric matrix $A \in \mathbb{R}^{n \times n}$, and let λ_1 and λ_2 be distinct eigenvalues of A with corresponding eigenvectors $\vec{v}_1$ and $\vec{v}_2$ respectively. We aim to show that $\vec{v}_1$ and $\vec{v}_2$ are orthogonal.

From the definition of eigenvectors and eigenvalues, we have:

$$A\vec{v}_1 = \lambda_1 \vec{v}_1,$$
$$A\vec{v}_2 = \lambda_2 \vec{v}_2.$$

Multiplying both sides of the first equation by $\vec{v}_2^T$ and both sides of the second by $\vec{v}_1^T$ gives:

$$\vec{v}_2^T A \vec{v}_1 = \lambda_1 \vec{v}_2^T \vec{v}_1,$$
$$\vec{v}_1^T A \vec{v}_2 = \lambda_2 \vec{v}_1^T \vec{v}_2.$$

Notice that this is a scalar so: $\vec{v}_1^T A \vec{v}_2 = (\vec{v}_1^T A \vec{v}_2)^T = \vec{v}_2^T A^T \vec{v}_1 = \vec{v}_2^T A \vec{v}_1$. The last step is true because A is symmetric, hence $A = A^T$. Therefore, we can equate the right-hand sides of the above equations:

$$\lambda_1 \vec{v}_2^T \vec{v}_1 = \lambda_2 \vec{v}_1^T \vec{v}_2.$$

Given $\lambda_1 \neq \lambda_2$, it follows that $\vec{v}_2^T \vec{v}_1 = \vec{v}_1^T \vec{v}_2 = 0$, demonstrating that $\vec{v}_1$ and $\vec{v}_2$ are orthogonal.

2.0.2 Proof that Orthogonal Eigenvectors Imply a Symmetric Matrix

Assume a matrix A has a set of orthogonal eigenvectors $\{\vec{v}_1, \vec{v}_2, \ldots, \vec{v}_n\}$ with corresponding eigenvalues $\{\lambda_1, \lambda_2, \ldots, \lambda_n\}$. Let Q be the matrix with columns as the normalized eigenvectors of A and Λ be the diagonal matrix with the eigenvalues of A.

Since the eigenvectors are orthogonal and normalized, Q is an orthogonal matrix, and hence $Q^T Q = QQ^T = I$, where I is the identity matrix.

We can write A as:

$$A = Q \Lambda Q^T.$$

To show A is symmetric, we need to prove $A^T = A$. Taking the transpose of the above equation:

$$A^T = (Q\Lambda Q^T)^T = Q\Lambda^T Q^T.$$

Since Λ is diagonal, $\Lambda^T = \Lambda$. Thus, we obtain:

$$A^T = Q\Lambda Q^T = A,$$

which shows that A is symmetric.

3 Dyadic Decomposition

Dyadic decomposition is a concept from linear algebra that involves representing a matrix as a sum of dyadic products. Dyadic products, also known as outer products, are matrices obtained from two vectors. Specifically, if $\mathbf{u}$ and $\vec{v}$ are vectors, their dyadic product is a matrix $\mathbf{u}\vec{v}^T$, where $\vec{v}^T$ denotes the transpose of $\vec{v}$.

3.1 Dyadic Decomposition for Symmetric Matrices

Any symmetric matrix $A \in \mathbb{R}^{n\times n}$ can be represented as a sum of outer products of vectors, which can be considered a form of dyadic decomposition. This representation is closely related to the spectral decomposition of A. For a symmetric matrix, the dyadic decomposition can be expressed in terms of its eigenvectors and eigenvalues as follows:

$$A = \sum_{i=1}^{r} \lambda_i u_i u_i^T \tag{10}$$

where u_i are the eigenvectors of A, λ_i are the corresponding eigenvalues, and r is the rank of A. Each term $u_i u_i^T$ represents a dyadic product, contributing to the overall structure of A.

4 Importance and Application

Symmetric Matrices Symmetric matrices play a pivotal role in mathematics and its applications due to their natural occurrence in a multitude of contexts. For instance, the product of a matrix A and its transpose A^T results in a symmetric matrix, regardless of whether A itself is symmetric. This can be seen in expressions such as AA^T and A^TA, which are inherently symmetric. Furthermore, the construction of a matrix through the multiplication of A, a diagonal matrix D, and the transpose of A (i.e., ADA^T) also yields a symmetric matrix. These forms are not just mathematical curiosities; they are fundamental in various fields, including statistics, where covariance matrices arise.

Dyadic Decomposition Dyadic decomposition is particularly useful in understanding the structure of matrices and in applications where matrices can naturally be expressed as sums of outer products. This includes areas such as signal processing, data analysis, and machine learning, where understanding the components of a matrix in terms of simpler, interpretable outer products can be highly valuable.

5 Example

To recap what we have said before, the Spectral Theorem provides a powerful way to analyze matrices by decomposing them into components related to their eigenvalues and eigenvectors. For a symmetric matrix A, the theorem guarantees that A can be decomposed in a form $A = Q\Lambda Q^T$, where Λ is a diagonal matrix of eigenvalues and Q is an orthogonal matrix whose columns are the normalized eigenvectors of A.

(i) Finding the Decomposition To decompose the given matrix $A = \begin{bmatrix} 1 & 3 \\ 3 & 1 \end{bmatrix}$, we first find its eigenvalues and eigenvectors. The characteristic equation of A is given by $\det(A - I) = 0$, leading to two eigenvalues $\lambda_1 = 4$ and $\lambda_2 = -2$. Corresponding normalized eigenvectors can be calculated as $\vec{v}_1 = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\vec{v}_2 = \frac{1}{\sqrt{2}}\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.

(ii) Showing Orthogonality The eigenvectors of a symmetric matrix like A are orthogonal. This can be verified by calculating the dot product of $\vec{v}_1$ and $\vec{v}_2$. Specifically, $\vec{v}_1^T\vec{v}_2 = 0$, confirming their orthogonality. This property is a direct consequence of the Spectral Theorem and is crucial for the orthogonal decomposition of A.

(iii) Dyadic Form The matrix A can also be represented in dyadic form using its eigenvalues and eigenvectors: $A = \lambda_1\vec{v}_1\vec{v}_1^T + \lambda_2\vec{v}_2\vec{v}_2^T$. Substituting the eigenvalues and eigenvectors found earlier, we obtain: $A = 4\left(\frac{1}{\sqrt{2}}\begin{bmatrix} 1 \\ 1 \end{bmatrix}\right)\left(\frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \end{bmatrix}\right) + (-2)\left(\frac{1}{\sqrt{2}}\begin{bmatrix} -1 \\ 1 \end{bmatrix}\right)\left(\frac{1}{\sqrt{2}}\begin{bmatrix} -1 & 1 \end{bmatrix}\right)$.

This dyadic form directly shows how the matrix A is constructed from its eigencomponents, illustrating the fundamental principle behind the Spectral Theorem.

Contributors:
- Druv Pai.
- Rahul Arya.
- Anant Sahai.
- Matteo Guarrera.

