METHODS FOR CHAPTER 7

Sections 1–2

Suppose that we have a linear operator T on an n-dimensional vector space V. In order to understand T, we seek a diagonalization, if one exists.

• 1. Calculate the representing matrix $A = [\![T]\!]_{\mathcal{S}}^{\mathcal{S}}$ for T, relative to some convenient basis $\mathcal{S}$, such as the standard basis, or the best basis we know of.

A is an $n \times n$ matrix. If T is L_A, matrix multiplication by A, then of course the representing matrix (relative to the standard basis for $\mathbb{R}^n$) is just A.

• 2. Find the characteristic polynomial $p(\lambda)$ of T, which is $\det(\lambda I - A)$, and its roots $\lambda_1, \ldots, \lambda_k$. Let m_i be the multiplicity of the root λ_i. (For large n, this step might need to be altered.)

This polynomial does not depend on our choice of $\mathcal{S}$. It is a polynomial of degree n in the variable λ. The roots $\lambda_1, \ldots, \lambda_k$ are the eigenvalues of T and of A. (A scalar λ is an eigenvalue if and only if $\dim E_\lambda > 0$, which happens if and only if the matrix $\lambda I - A$ is singular.)

Over $\mathbb{R}$, we have $p(t) = (t-\lambda_1)^{m_1} \cdots (t-\lambda_k)^{m_k} \cdot$(irreducible quadratics). And $m_1 + \cdots + m_k \le n$.

Over $\mathbb{C}$, we have simply $p(t) = (t-\lambda_1)^{m_1} \cdots (t-\lambda_k)^{m_k}$ and $m_1 + \cdots + m_k = n$.

• 3. For each eigenvalue λ_i, find a basis for the nullspace of $\lambda_i I - A$. By decoordinatizing, find a basis for E_{λ_i}.

E_{λ_i} is isomorphic to the nullspace of $\lambda_i I - A$ under to coordinate map $[\![\ ]\!]_{\mathcal{S}}$. We know that

$$1 \le \dim E_{\lambda_i} \le m_i$$

that is, the "geometric multiplicity" $\dim E_{\lambda_i}$ does not exceed the "algebraic multiplicity" m_i.

Combine the bases for $E_{\lambda_1}, \ldots, E_{\lambda_k}$. This gives a maximal linearly independent set of eigenvectors for T. (It is linearly independent by the theorem on independence of eigenvectors for different eigenvalues.) And the number of independent eigenvectors is

$$(\star) \qquad \dim E_{\lambda_1} + \cdots + \dim E_{\lambda_k} \le m_1 + \cdots + m_k \le n.$$

• 4. Now there are two cases.

4A. The good case: $\sum_i \dim E_{\lambda_i} = n$ (i.e., equality holds in $(\star)$).

This happens if and only if both (a) the characteristic polynomial splits completely into linear factors (no irreducible quadratics), and (b) for each i for which $m_i > 1$, the geometric multiplicity $\dim E_{\lambda_i}$ equals the algebraic multiplicity m_i. (For example, (a) always happens if the field is $\mathbb{C}$. And (b) always happens if $m_i = 1$ for each i, i.e., if every root of the characteristic polynomial is a simple root.)

Then we have a basis $\mathcal{E}$ of n eigenvectors. The matrix $D = [\![T]\!]_{\mathcal{E}}^{\mathcal{E}}$ is an $n \times n$ diagonal matrix, and the n entries on the diagonal are exactly the eigenvalues, each repeated according to its multiplicity (i.e., λ_i is repeated m_i times). The diagonal matrix D will equal $Q^{-1}AQ$, where $Q = [\![I_V]\!]_{\mathcal{S}}^{\mathcal{E}}$, so that the columns of Q are the coordinate vectors of the eigenvectors in $\mathcal{E}$ (i.e., the column vectors from step 3).

4B. The bad case: $\sum_i \dim E_{\lambda_i} < n$.

This happens if and only if either (a) the characteristic polynomial does not split completely into linear factors, or (b) for some i, the eigenspace E_{λ_i} has dimension less than m_i. In this case, T is not diagonalizable. (The maximal linearly independent set from step 3 fails to span V.) Even in the bad case, we can form a basis with as many eigenvectors as possible. The resulting matrix may not be diagonal, but it might be better than the one we started with. Beyond that, there is something called *Jordan normal form*; see §7.5.