Canonical forms of Linear Transformations

David M. Rocke
Department of Applied Science
UC Davis

April 16, 2003

1. Overview

A *canonical form* of a linear transformation is a matrix representation in a basis chosen to make that representation simple in form. The most common canonical form is a diagonal matrix. In this section of the course, we explore canonical forms with three main types of results:

1. For real symmetric or complex Hermitian matrices, we show that these are always diagonalizable; i.e., there exists a basis wrt which the matrix representation is diagonal.

2. For general linear transformations, we show that it is diagonalizable if and only if its minimal polynomial is a product of distinct linear factors.

3. For any linear transformation for which the characteristic polynomial factors completely (this is all linear transformations if the field is $\mathbb{C}$), there is a matrix representation in Jordan canonical form.

Most of this will be shown directly in class, assuming the standard facts about real and complex numbers and solution and factoring of polynomials. The Cayley-Hamilton theorem is given without proof, as this would be too extensive.

2. Eigenvalues and Eigenvectors

Let $T \in L(V)$. If $Tx = \lambda x$ for some $x \in V$, we say that λ is an *eigenvalue* with associated eigenvector x. If λ is an eigenvalue of T, then the *eigenspace* associated with λ is $\{x \in V | Tx = \lambda x\}$.

The following is an immediate consequence of the definitions.

Theorem 1 *Suppose $T \in L(V)$, with V finite dimensional. Then the following are equivalent:*

1. λ is an eigenvalue of T.

2. $T - \lambda I$ *is singular and therefore not invertible.*

3. $|T - \lambda I| = 0$.

If u is a scalar variable, then $|T - uI|$ is a polynomial of degree n in u called the *characteristic polynomial.* It's roots are exactly the eigenvalues according to the theorem above. Next, we show that the characteristic polynomial is independent of the basis representation.

Lemma 1 *Let $T \in L(V)$, V finite dimensional, and let A be a matrix representation of T in a particular basis. If P is a basis transformation, then a point x in the old basis is Px in the new basis, and T has matrix representation $B = P^{-1}AP$ in the new basis. Then*

$$\begin{aligned} |B - uI| &= |P^{-1}AP - uI| \\ &= |P^{-1}AP - uP^{-1}P| \\ &= |P^{-1}(A - uI)P| \\ &= |P^{-1}||(A - uI)||P| \\ &= |(A - uI) \end{aligned}$$

since $|P^{-1}| = |P|^{-1}$.

A linear transformation T is *diagonalizable* if there exists a basis wrt which the matrix representation of T is diagonal, or, equivalently, if there is a basis of the whole vector space consisting of eigenvectors of T. If T has a diagonal representation, and $\lambda_1, \lambda_2, \ldots, \lambda_k$ are the distinct eigenvalues, then T's matrix representation consists of k diagonal blocks, of which the ith block is the diagonal matrix $\lambda_i I_{d_i \times d_i}$ and the characteristic polynomial of T is $(u - \lambda_1)^{d_1}(u - \lambda_2)^{d_2} \cdots (u - \lambda_k)^{d_k}$. For diagonalizable transformations, the dimension of the eigenspace of λ_i is d_i.

Lemma 2 *If $T \in L(V)$, V finite dimensional of dimension n, and if T has n distinct eigenvalues, then T is diagonalizable.*

Proof. Let $(\lambda_1, x_1), (\lambda_2, x_2), \ldots, (\lambda_n, x_n)$ be the distinct eigenvalues and an associated eigenvector for each. We show that $\{x_1, x_2, \ldots, x_n\}$ is a linearly independent set. Then this must be a basis of V since the space spanned by them is of dimension n and must therefore be the whole space.

Suppose, on the contrary, that

$$\sum_{i=1}^{n} a_i x_i = 0 \tag{2.1}$$

for some set of constants a_i not all zero. Without loss of generality, let $a_1 \neq 0$ Applying T to (2.1), we get

$$\sum_{i=1}^{n} a_i \lambda_i x_i = 0$$

so that

$$\begin{aligned}
0 &= \lambda_n \sum_{i=1}^{n} a_i x_i - \sum_{i=1}^{n} a_i \lambda_i x_i \\
&= \sum_{i=1}^{n} a_i(\lambda_n - \lambda_i) x_i \\
&= \sum_{i=1}^{n-1} a_i(\lambda_n - \lambda_i) x_i
\end{aligned}$$

We can then apply T again, multiplying by λ_{n-1} and subtracting, and continue the process to obtain

$$a_1(\lambda_n - \lambda_1)(\lambda_{n-1} - \lambda_1)\cdots(\lambda_2 - \lambda_1) = 0$$

Since all the eigenvalues are distinct by hypothesis, this means that $a_1 = 0$, a contradiction. Thus, the x_i are linearly independent. $\square$

3. Symmetric and Hermitian matrices

A *symmetric* matrix A is one in which $A^\top = A$. A *Hermitian* matrix A has $A^* = A$, where A^* is the transpose complex conjugate of A. An *orthogonal* matrix A is one in which $A^\top A = I$ and a *unitary* matrix A is one in which $A^* A = I$.

Lemma 3 *All eigenvalues of a symmetric or Hermitian matrix are real.*

PROOF. Let λ be an eigenvalue of a Hermitian matrix A and let x be the associated eigenvector. Then $Ax = \lambda x$, so multiplying by x^* on the left we obtain $x^*Ax = \lambda x^*x$. Taking the adjoint of both sides of the equation defining the eigenvalues, we obtain $x^*A^* = \bar{\lambda}x^*$, yielding $x^*A^*x = \bar{\lambda}x^*x$ by multiplying by x on the right. Since $A^* = A$, this implies that $\lambda x^*x = \bar{\lambda}x^*x$. Since x^*x cannot be zero for nonzero x, this implies $\bar{\lambda} = \lambda$, so that λ is real. The proof for symmetric matrices is the same $\square$

Lemma 4 *Eigenvectors corresponding to distinct eigenvalues of a real symmetric or Hermitian matrix are orthogonal.*

PROOF. Suppose that (λ_1, x_1) and (λ_2, x_2) are eigenvalue/eigenvector pairs for a Hermitian matrix A, and suppose that $\lambda_1 \neq \lambda_2$. Then $Ax_1 = \lambda_1 x_1$ and so $x_2^* A x_1 = \lambda_1 x_2^* x_1$ (multiplying on the left by x_2^*). Also, $Ax_2 = \lambda_2 x_2$, so $x_2^* A^* = x_2^* A = \lambda_2 x_2^*$ (using the definition of Hermitian and the previous lemma), and thus $x_2^* A x_1 = \lambda_2 x_2^* x_1$ (multiplying on the right by x_1). This in turn implies that $\lambda_1 x_2^* x_1 = \lambda_2 x_2^* x_1$. Since $\lambda_1 \neq \lambda_2$, $x_2^* x_1 = 0$; that is, the two vectors are orthogonal. The proof for symmetric matrices is the same. $\square$

Theorem 2 *Let A be a Hermitian (real symmetric) matrix. Then A is diagonalizable, and the basis transformation matrix P can be chosen to be unitary (orthogonal).*

Proof. Let A be Hermitian and let (λ_1, x_1) be an eigenvalue/eigenvector pair. Construct an orthonormal basis of V consisting of $y_1 = x_1/||x_1||$, y_2, $\ldots$, y_n. The basis transformation into this new coordinate system has a matrix P satisfying $P^*P = I$ because the new basis is orthonormal. Let B be the matrix in the new coordinate system, so that $B = P^*AP$. Note that B is still Hermitian because $B^* = (P^*AP)^* = P^*A^*P = P^*AP = B$. Also, B has the same eigenvalues as A since it has the same characteristic polynomial, and y_1 is an eigenvector because y_1 in the new coordinate system is $P^*x_1/||x_1||$ and $By_1 = P^*APP^*x_1/||x_1|| = P^*Ax_1/||x_1|| = \lambda_1 P^*x_1/||x_1|| = \lambda_1 y_1$. The fact that $Ay_1 = \lambda_1 y_1$, implies that the matrix representation of the linear transformation in the new basis has a first column consisting of $(\lambda_1, 0, \ldots, 0)^\top$. This is so because we know that the vector with coordinates $(1, 0, \ldots, 0)$ is mapped into $(\lambda_1, 0, \ldots, 0)$. Now consider a basis vector y_i with $i \neq 1$. $By_i \in V$, so $By_i = cy_1 + \tilde{y}$, where $\tilde{y}$ is a linear combination of y_2, y_3,$\ldots$, y_n and is therefore orthogonal to y_1. Now

$$\begin{aligned} y_1^* B y_i &= c y_1^* y_1 + y_1^* \tilde{y} \\ &= c. \end{aligned}$$

However, $y_1^* B = \lambda_1 y_1^*$ so $y_1^* B y_i = \lambda_1 y_1^* y_i = 0$. Thus, $c = 0$. This means that B has a first row in which all the off-diagonal elements are zero and is of the form

$$B = \begin{pmatrix} \lambda_1 & 0 \\ 0 & C \end{pmatrix}$$

Here C is a Hermitian (real symmetric) matrix of degree $n-1$ since

$$B^* = \begin{pmatrix} \lambda_1 & 0 \\ 0 & C \end{pmatrix}^* = \begin{pmatrix} \lambda_1 & 0 \\ 0 & C^* \end{pmatrix} = B = \begin{pmatrix} \lambda_1 & 0 \\ 0 & C \end{pmatrix}$$

so that $C^* = C$.

We now finish the proof by induction on the dimension n. For $n = 1$, the result is obvious since 1×1 matrices are trivially diagonal. If the result is true for Hermitian matrices of dimension $n-1$, then there is a unitary matrix Q such that Q^*CQ is diagonal. This implies that B is diagonalized by the unitary matrix

$$R = \begin{pmatrix} 1 & 0 \\ 0 & Q \end{pmatrix}$$

□

4. A Criterion for General Diagonalizability

In this section, we derive a criterion for diagonalizability based on the minimal polynomial, which is related to the characteristic polynomial. Consider a diagonal matrix A with distinct eigenvalues λ_1, λ_2,..., λ_k. Then as previously observed, A is of the form

$$A = \begin{pmatrix} \lambda_1 I_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 I_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & 0 \\ 0 & 0 & \cdots & \lambda_k I_k \end{pmatrix}$$

where I_i is a $d_i \times d_i$ identity matrix. This means that the characteristic polynomial of A is $p(u) = (u-\lambda_1)^{d_1}(u-\lambda_2)^{d_2}\cdots(u-\lambda_k)^{d_k}$. Now $A - \lambda_1 I$ has the first d_1 diagonal entries equal to zero, $(A-\lambda_2 I)$ has the next d_2 entries equal to zero, so that the matrix $(A-\lambda_1 I)(A-\lambda_2 I)\cdots(A-\lambda_k I) = 0$, the matrix of all zeroes. Possible a lower degree polynomial may serve as well, because if one of the eigenvalues, say λ_k, is zero, then one may omit that factor and the resulting polynomial in A is still the zero polynomial. Thus, for any diagonal matrix A, there is a polynomial $p(u)$ of degree $k \leq n$ which is a product of distinct linear factors and which when A is substituted in for the variable produces a zero linear transformation. Note also, that this polynomial has the same factors as the characteristic polynomial (or possibly one fewer), but in any case divides the characteristic polynomial. The rest of the section is spent showing that this condition is sufficient as well as necessary. To do this, we will need to take a short detour in the study of polynomial ideals.

4.1. Polynomial Ideals

The set of polynomials $\mathcal{P}$ is a vector space, and in fact is a commutative algebra with ordinary multiplication of polynomials as the algebra's multiplication. A *polynomial ideal* is a subspace $\mathcal{F} \subset \mathcal{P}$ which is closed under multiplication by polynomials; that is, whenever $f \in \mathcal{F}$, and $p \in \mathcal{P}$, then $fp \in \mathcal{F}$. If $f \in \mathcal{F}$, the *principal ideal* generated by f, denoted $\langle f \rangle$, is the smallest polynomial ideal containing f. It consists of all polynomials with f as a factor. As it turns out, the principal ideals are all polynomial ideals:

Lemma 5 *Let $\mathcal{F} \subset \mathcal{P}$ be a polynomial ideal. Then there exists a monic (leading coefficient 1) polynomial f with $\mathcal{F} = \langle f \rangle$.*

Proof. Let f be a nonzero polynomial in $\mathcal{F}$ of smallest degree, and wlog let f be monic. We show that $\mathcal{F} = \langle f \rangle$. Let $g \in \mathcal{F}$. It is sufficient to show that f divides g. By polynomial division with remainders, we can write $g = fq + r$, where $\deg(r) < \deg(f)$. But by hypothesis, $g \in \mathcal{F}$, and $fq \in \mathcal{F}$ because $\mathcal{F}$ is a polynomial ideal, so $r \in \mathcal{F}$ since $\mathcal{F}$ is a vector subspace of $\mathcal{P}$. But f is by choice a polynomial in $\mathcal{F}$ of lowest degree possible for nonzero polynomials, so $r = 0$, and g is therefore a multiple of f. $\square$

4.2. Cyclic Subspaces and Annihilating Polynomials

Consider $T \in L(V)$, V of finite dimension n, and let $x \in V$. Consider the sequence of images of x consisting of x, Tx, T^2x, ...and consider the sequence of subspaces $V_0 = \langle x \rangle$,

$V_1 = \langle x, Tx \rangle$, $V_2 = \langle x, Tx, T^2x \rangle$, ... At some point, the sequence of subspaces must stop increasing in size, since the dimension increases by 1 at each step if the size increases, and the dimension is limited to n. Thus, at some point $T^kx \in V_{k-1}$, so

$$T^kx = -c_0x - c_1Tx - c_2T^2x - \cdots - c_{k-1}T^{k-1}x$$

so $p(T)x = 0$ for the polynomial $p(u) = c_0 + c_1u + \cdots + c_{k-1}u^{k-1} + u^k$.

Now this is for a specific vector x, but we can also look at polynomials in T that are zero for every x. First, there must be at least one of degree at most $n \times n$, since $I, T, T^2, \ldots, T^{n^2}$ cannot be linearly independent, since $\mathbb{R}^{n\times n}$ is of dimension n^2 and the set of powers of T has n^2+1 members. Thus there is some linear combination $c_0I + c_1T + c_2T^2 + \cdots + c_{n^2}T^{n^2} = 0$. Thus the polynomial $p(u) = c_0 + c_1u + c_2u^2 + \cdots + c_{n^2}u^{n^2}$ has $p(T) = 0$. Such a polynomial is said to *annihilate* T. Now consider the set $\mathcal{A}_T \subset \mathcal{P}$ of polynomials that annihilate a fixed polynomial T. The set is non-empty, and is a vector subspace since $0+0 = 0$ and the 0 linear transformation times any scalar is also zero. Furthermore, it is a polynomial ideal, since if $p(T) = 0$, then also $p(T)q(T) = 0$ for any polynomial q. Thus, it must be the principal ideal generated by a unique monic polynomial $m(u)$, called the *minimal polynomial.*

Lemma 6 *Let $m(u)$ be the minimal polynomial of a linear transformation T and let λ be an eigenvalue of T. Then $m(\lambda) = 0$. Conversely, if $m(c) = 0$, then c is an eigenvalue of T.*

Proof. First, suppose that λ is an eigenvalue of T. Then $Tx = \lambda x$, so $0 = m(T)x = m(\lambda)x$, so $m(\lambda) = 0$. Conversely, if $m(c) = 0$, then $m(u) = (u - c)q(u)$ and $q(T) \neq 0$, since q is of lower degree than the minimal polynomial u. Let x be such that $y = q(T)x \neq 0$. Then $0 = m(T)x = (T - cI)q(T)x = (T - cI)y$, so c is an eigenvalue of T. □

The last result means that the minimal polynomial and the characteristic polynomial have the same linear factors. Furthermore, if T is diagonalizable, then the minimal polynomial for T must be exactly the product of $(u - \lambda_i)$ for all non-zero distinct eigenvalues λ_i. We state the following without proof.

Theorem 3 (Cayley-Hamilton) *Let p be the characteristic polynomial for $T \in L(V)$, V finite dimensional. Then $p(T) = 0$.*

This implies that the minimal polynomial must divide the characteristic polynomial. We now show that T is diagonalizable if and only if the minimal polynomial for T is a product of distinct linear factors.

Lemma 7 *Let $T \in L(V)$, V of finite dimension n, and let $m(u)$ be the minimal polynomial of T, Suppose that $m(u) = (u - c_1)^{r_1}(u - c_2)^{r_2} \cdots (u - c_k)^{r_k}$. Let $W \subset V$, $W \neq V$, and suppose that W is* invariant *under T, meaning that if $w \in W$, then $Tw \in W$. Then there exists $y \in V$ with $y \notin W$ and some eigenvalue λ of T such that $(T - \lambda I)y \in W$.*

Proof. Fix $x \in V \setminus W$ and consider the set $\mathcal{F}$ of polynomials $p(u)$ such that $p(T)x \in W$. The minimal polynomial $m(u)$ is in $\mathcal{F}$, so the set is non-empty. Also, we can show that $\mathcal{F}$

is a polynomial ideal. First, it is clearly a vector space. Next, if $p \in \mathcal{F}$ and $q \in \mathcal{P}$ then $p(T)x \in W$ because $p \in \mathcal{F}$, and $q(T)[p(T)x] \in W$ because W is invariant under T, and thus $pq \in \mathcal{F}$. Since $\mathcal{F}$ is a polynomial ideal, it must be the principal ideal generated by some monic polynomial g. Since $m \in \mathcal{F}$, $g|m$, so $g(u) = (u-c_1)^{e_1}(u-c_2)^{e_2}\cdots(u-c_k)^{e_k}$, with $e_i \leq r_i$. For at least one i, we must have $e_i > 0$ since $g \neq 1$, say e_j. Then $g(u) = (u-c_j)h(u)$ and $h(u) \notin \mathcal{F}$ because g is a polynomial of minimal degree in $\mathcal{F}$. Now let $y = h(T)x \notin W$. We have $(T - c_jI)y = (T - c_j)h(T)x = g(T)x \in W$, as required. $\square$

Theorem 4 *Let $T \in L(V)$ with V of finite dimension n. Then T is diagonalizable if and only if the minimal polynomial for T is a product of distinct linear factors $m(u) = (u-\lambda_1)(u-\lambda_2)\cdots(u-\lambda_k)$, where the λ_i are the distinct eigenvalues of T*

We have already shown that if T is diagonalizable, then the minimal polynomial of T is of the required form. Now we show the converse. Suppose that the minimal polynomial of T is of the required form. If V is spanned by eigenvectors of T, then T has a basis of eigenvectors, and is thus diagonal in that basis. So suppose that the subspace spanned by the eigenvectors is W, a proper subspace of V. W is invariant under T since any eigenvector of T is taken onto another eigenvector with the same eigenvalue. Then, by the previous lemma, there exists a vector $x \in V \setminus W$ and an eigenvalue λ_j of T such that $y = (T - \lambda_j I)x \in W$. By hypothesis, $m(u) = (u-\lambda_j)q(u)$ with $q(\lambda_j) \neq 0$. The polynomial $q(u) - q(\lambda_j)$ has a zero at λ_j, so factors as $q(u) - q(\lambda_j) = (u-\lambda_j)h(u)$. Now $q(T)x - q(\lambda_j)x = h(T)(T-\lambda_j I)x = h(T)y \in W$. But $0 = m(T)x = (T-\lambda_j I)q(T)x$, so $q(T)x$ is an eigenvector and thus lies in W. Thus $q(\lambda_j)x$ must also lie in W, which is a contradiction since $q(\lambda_j) \neq 0$. $\square$

5. Jordan Canonical Form

We give the following result without proof.

Theorem 5 *Let $T \in L(V)$, with V of finite dimension n. Suppose that the characteristic polynomial of T factors completely as $(u-\lambda_1)^{d_1}(u-\lambda_2)^{d_2}\cdots(u-\lambda_k)^{d_k}$. Then there exists a basis wrt which the matrix representation A of T has the following form: A is block diagonal with blocks A_1, A_2,…,A_k. Each block corresponds to one eigenvalue and is in turn block diagonal with blocks B_{i1}, B_{i2},…, B_{in_i}. Each of the B_{ij} has a diagonal consisting of the eigenvalue λ_i, has 1 on each entry directly below the diagonal, and has zeroes elsewhere. There is exactly one eigenvector corresponding to each B_{ij}.*