NOTES ON CANONICAL FORMS

2009-04-14 14:20

1. Introduction: linear operators, invariant subspaces, and polynomials

Throughout these notes V will denote a vector space with finite dimension n over a field $\mathbf{F}$, and $T : V \to V$ will be a linear operator. In many ways, T is easiest to understand if it is 'diagonalizable'—that is, if there exists a basis for V relative to which the matrix for T is diagonal. Diagonalization is not always possible, however. Our main concerns in these notes will be to find ways of presenting operators as nicely as possible and to understand thoroughly why they cannot be presented better.

Recall that there are essentially two obstacles to diagonalization. The first is that the characteristic polynomial for a given operator might not have sufficiently many roots (or at least roots belonging to the scalar field). For instance, if V is a vector space over $\mathbf{R}$ and $T : V \to V$ has characteristic polynomial $\lambda^2 + 1$, then T has no real eigenvalues. The second and more subtle obstacle to diagonalization occurs when the characteristic polynomial has a root λ that occurs with multiplicity $k > 1$, but the eigenspace associated to λ has dimension smaller than k. The standard example of this latter problem is the linear transformation $T : \mathbf{R}^2 \to \mathbf{R}^2$ whose matrix relative to the standard basis is

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}.$$

The only eigenvalue for T is $\lambda = 1$, which occurs twice as a root of the characteristic polynomial. However the eigenspace for this eigenvalue is one dimensional, generated by the vector $(1, 0)$.

1.1. Invariant subspaces.

Confronted with either of these obstacles, we need something to replace the missing eigenvectors. It turns out that it is better to think in terms of subspaces than eigenvectors.

Definition 1.1. *A subspace $H \subset V$ is* T-invariant *if $T(H) \subset H$.*

Invariant subspaces have the following basic properties, whose verification we leave as an exercise.

Proposition 1.2. *If $H_1, H_2 \subset V$ are T-invariant subspaces, then so are $H_1 + H_2$ and $H_1 \cap H_2$. If $S : V \to V$ is another linear operator and $H \subset V$ is invariant with respect to both S and T, then H is also invariant with respect to $S \circ T$ and with respect to any linear combination of S and T.*

If $\mathbf{v} \in V$ is an eigenvector for T, then one has that $H = \operatorname{span}\{\mathbf{v}\}$ is a one dimensional T-invariant subspace. In fact, a one dimensional subspace of V is T-invariant *only if* it is spanned by an eigenvector of T. However, one can create invariant subspaces from *any* vector, eigen or not.

Definition 1.3. *Given $\mathbf{v} \in V$, we call the set*

$$\mathrm{orb}_T(\mathbf{v}) = \{\mathbf{v}, T\mathbf{v}, T^2\mathbf{v}, \dots\}$$

*the T-*orbit *of $\mathbf{v}$. We call*

$$H_{T,\mathbf{v}} := \operatorname{span} \mathrm{orb}_T(\mathbf{v})$$

*the T-*cyclic subspace generated by $\mathbf{v}$.

Proposition 1.4. *For any $\mathbf{v} \in V$, the subspace $H_{T,\mathbf{v}}$ is T-invariant. If, moreover, $H \subset V$ is any other T-invariant subspace containing $\mathbf{v}$, then $H_{T,\mathbf{v}} \subset H$.*

Proof. If $\mathbf{w} \in H_{T,\mathbf{v}}$, then we can write

$$\mathbf{w} = c_0\mathbf{v} + \cdots + c_kT^k\mathbf{v}$$

as a linear combination of vectors in $\mathrm{orb}_{T,\mathbf{v}}$. Thus

$$T(\mathbf{w}) = c_0T(\mathbf{v}) + \cdots + c_kT^{k+1}(\mathbf{v}) \in H_{T,\mathbf{v}},$$

too. Hence $H_{T,\mathbf{v}}$ is T-invariant. Now if $H \subset V$ is a T-invariant subspace containing $\mathbf{v}$, it follows inductively that $T\mathbf{v}, T^2\mathbf{v}, \cdots \in H$, too. Since H is closed with respect to linear combinations, it follows that $H_{T,\mathbf{v}} \subset H$. □

Some examples might help here. Both V and $\{\mathbf{0}\}$ are always T-invariant. A one dimensional subspace $H \subset V$ is T-invariant if and only if $H = H_{T,\mathbf{v}}$ for some eigenvector $\mathbf{v} \in V$.

We saw in class last semester that if $\mathbf{F} = \mathbf{R}$ and $\lambda = a + bi \in \mathbf{C}$ is a non-real root of the characteristic polynomial, then there are vectors $\mathbf{u}, \mathbf{v} \in V$ such $T(\mathbf{u}) = a\mathbf{u} - b\mathbf{v}$ and $T(\mathbf{v}) = b\mathbf{u} + a\mathbf{v}$. In this case $H_{T,\mathbf{u}} = H_{T,\mathbf{v}} = \operatorname{span}(\mathbf{v}, \mathbf{u})$ is a two dimensional T-invariant subspace. In this last case, one can show that in fact $H_{T,\mathbf{v}}$ is 'irreducible' in the sense that it contains no smaller non-trivial T-invariant subspaces.

Finally, if $T : \mathbf{R}^3 \to \mathbf{R}^3$ is the diagonal transformation $T(x_1, x_2, x_3) = (x_1, x_2, 2x_3)$, then the vector $\mathbf{v} = (0, 1, 1)$ generates a two dimensional cyclic subspace $H_{T,\mathbf{v}} = \{(0, x_2, x_3) \in \mathbf{R}^3\}$. It is, however, 'reducible' in the sense that it contains the smaller T-cyclic subspace spanned by the eigenvector $(0, 0, 1)$.

If H is a T-invariant subspace, then it turns out that T 'induces' a linear transformation on the quotient space V/H.

Theorem 1.5. *Given a T-invariant subspace $H \subset V$, let $\tilde{V} = V/H$ and for any $\mathbf{v} \in V$, let $\tilde{\mathbf{v}} \in \tilde{V}$ denote the equivalence class of $\mathbf{v}$. Then there is a well-defined linear operator $\tilde{T} : \tilde{V} \to \tilde{V}$ given by*

$$\tilde{T}(\tilde{\mathbf{v}}) = \widetilde{(T\mathbf{v})}.$$

If, moreover, $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ is a basis for V obtained by extending a basis $\mathcal{B}\{\mathbf{b}_1, \dots, \mathbf{b}_k\}$ for H, then the matrix for T relative to $\mathcal{B}$ has block upper triangular form

$$A = \begin{bmatrix} A_H & * \\ 0 & \tilde{A} \end{bmatrix},$$

where $A_H \in M_{k\times k}(\mathbf{F})$ is the matrix of the restricted operator $T|_H : H \to H$ relative to $\mathcal{B}_H$, and $\tilde{A} \in M_{(n-k)\times(n-k)}(\mathbf{F})$ is the matrix of $\tilde{T} : \tilde{V} \to \tilde{V}$ relative to the basis $\tilde{\mathcal{B}} = \{\tilde{\mathbf{b}}_{k+1}, \dots, \tilde{\mathbf{b}}_n\}$. In particular, the characteristic polynomial of $T : V \to V$ is the product of those of $T|_H$ and $\tilde{T}$.

Note that in what follows, we will sometimes write T/H instead of $\tilde{T}$ in order to emphasize the role of H.

Proof. To check that $\tilde{T}$ is well-defined, we suppose that $\mathbf{v}_1, \mathbf{v}_2 \in V$ satisfy $\tilde{\mathbf{v}}_1 = \tilde{\mathbf{v}}_2$. Then $\mathbf{v}_1 - \mathbf{v}_2 \in H$. Hence $T(\mathbf{v}_1) - T(\mathbf{v}_2) = T(\mathbf{v}_1 - \mathbf{v}_2) \in H$, too, because H is T-invariant. Hence $\widetilde{T\mathbf{v}_1} = \widetilde{T\mathbf{v}_2}$, so that our definition of $\tilde{T}(\tilde{\mathbf{v}})$ does not depend on which vector $\mathbf{v}$ is chosen to represent the equivalence class $\tilde{\mathbf{v}}$.

That $\tilde{T}$ is linear follows more or less immediately from linearity of T; e.g.

$$\tilde{T}(\tilde{\mathbf{v}}_1 + \tilde{\mathbf{v}}_2) = T(\widetilde{\mathbf{v}_1 + \mathbf{v}_2}) = T\widetilde{\mathbf{v}_1 + T\mathbf{v}_2} = \widetilde{T\mathbf{v}_1} + \widetilde{T\mathbf{v}_2} = \tilde{T}(\tilde{\mathbf{v}}_1) + \tilde{T}(\tilde{\mathbf{v}}_2),$$

so $\tilde{T}$ respects vector addition. A similar computation shows that $\tilde{T}$ respects scalar multiplication.

Now we turn to the assertion relating the matrices for T, $T|_H$ and $\tilde{T}$. Letting A denote the matrix for $T : V \to V$ relative to $\mathcal{B}$, we recall that the jth column of M is equal to $[T\mathbf{b}_j]_{\mathcal{B}}$, the coordinates of $T\mathbf{b}_j$ relative to the basis $\mathcal{B}$. Now if $1 \leq j \leq k$, then $\mathbf{b}_j \in \mathcal{B}_H$ is a vector in H. So by invariance of H, we have $T\mathbf{b}_j \in H$ too. Thus the last $n-k$ coordinates of $[T\mathbf{b}_j]_{\mathcal{B}}$ vanish and the first k coordinates are equal to $[T\mathbf{b}_j]_{\mathcal{B}_H}$. In short, each of the first k columns of A have the form

$$[T\mathbf{b}_j]_{\mathcal{B}} = \begin{pmatrix} [T\mathbf{b}_j]_{\mathcal{B}_H} \\ \mathbf{0} \end{pmatrix},$$

so that taken altogether, the first k columns of A comprise a matrix of block form $\begin{pmatrix} A_H \\ 0 \end{pmatrix}$.

Turning to the remaining columns $k+1 \leq j \leq n$ of A, we note that if

$$T\mathbf{b}_j = c_1\mathbf{b}_1 + \ldots c_n\mathbf{b}_n,$$

then because each vector in H is equivalent to $\mathbf{0}$ modulo H,

$$\tilde{T}\tilde{\mathbf{b}}_j = c_{k+1}\mathbf{b}_{k+1} + \cdots + c_n\mathbf{b}_n.$$

So the last $n-k$ coordinates of $T\mathbf{b}_j$ relative to $\mathcal{B}$ are equal to the coordinates of $\tilde{T}\tilde{\mathbf{b}}_j$ relative to $\tilde{\mathcal{B}}$. That is, $[T\mathbf{b}_j]_{\mathcal{B}} = \begin{pmatrix} * \\ \left[\tilde{T}\tilde{\mathbf{b}}_j\right]_{\tilde{\mathcal{B}}} \end{pmatrix}$. This means that taken together, the last $n-k$ columns of A comprise a matrix of the form $\begin{pmatrix} * \\ \tilde{A} \end{pmatrix}$. So A has the block form asserted in the theorem.

The assertion concerning characteristic polynomials now follows immediately from the fact that

$$\lambda I - M = \begin{bmatrix} \lambda I - A & -B \\ 0 & -C \end{bmatrix}$$

also has block upper triangular form, and therefore $\det(\lambda I - M) = \det(\lambda I - A)\det(\lambda I - C)$. $\square$

1.2. **Polynomials.** Having seen that the notion of eigenvector of T can be generalized to that of an invariant subspace, we now indicate how that of an eigenvalue can by similarly generalized by considering certain special (i.e T-singular) polynomials associated to T.

We will use $\mathbf{F}[x]$ to denote the set of all polynomials $p(x) = c_k x^k + \cdots + c_1 x + c_0$, with coefficients $c_k, \ldots, c_0 \in \mathbf{F}$. The *degree* of p is the largest index k for which the coefficient c_k is non-zero. If $p = 0$, then we adopt the convention that $\deg p = -\infty$. We call a non-zero $p \in \mathbf{F}[x]$ *monic* if it's leading coefficient $c_k = 1$. Finally, if $a, b, q \in \mathbf{F}[x]$ are polynomials such that $a = bq$, then we will say that a is *divisible by* b or, more commonly, that b *divides* a, signifying the relationship by writing $b|a$. We will often take advantage of the fact that $b|a$ implies $\deg b \leq \deg a$. Hence, for instance, the final assertion in Theorem 1.5 may be restated in slightly weaker form by saying that the characteristic polynomials of $T|_H$ and T/H both divide that of T.

Given a polynomial $p(x) \in \mathbf{F}[x]$ as above, one can replace the independent variable x by things other than just elements of $\mathbf{F}$. In our case, we will substitute the linear operator T, defining:

$$p(T) = c_k T^k + \cdots + c_1 T + c_0 \mathrm{id}$$

Here the power T^j should be understand as the j-fold composition $T \circ T \circ \cdots \circ T$ of T with itself. One can readily verify the basic features of this kind of substitution, which we now summarize.

Proposition 1.6. *If $p, q \in \mathbf{F}[x]$ are polynomials, and $H \subset V$ is a T-invariant subspace, then*

- *H is a $p(T)$ invariant subspace;*
- *$p(T) \circ q(T) = q(T) \circ p(T) = (pq)(T)$;*
- *$\ker p(T)$ is a T-invariant subspace.*

The second assertion in Proposition 1.6 says among other things that for any $p, q \in \mathbf{F}[x]$, the operators $p(T)$ and $q(T)$ commute. Typically in what follows, we will write $p(T)q(T)$ instead of $p(T) \circ q(T)$. Besides emphasizing the connection between composition of operators and multiplication of polynomials, this abbreviation accords well with our tendency to write $T\mathbf{v}$ instead of $T(\mathbf{v})$ when the parentheses start to pile up.

The third assertion in Proposition 1.6 gives the connection between eigenvalues and polynomials: in the case of a *linear* polynomial $p(x) = x - \lambda$, we have that $\ker p(T) = \ker(\lambda \mathrm{id} - T)$ is the eigenspace for λ, and in particular λ is an eigenvalue if and only if $\ker p(T)$ is non-trivial. Let us more generally call a polynomial (of any degree) *T-singular* if $\ker p(T)$ is non-trivial.

There is a close, albeit not perfectly complementary, connection between T-singular polynomials and T-cyclic subspaces. Fixing $\mathbf{v} \in V$, we observe first that any element $\mathbf{w} \in H_{T,\mathbf{v}}$ can be written

$$\mathbf{w} = c_k T^k \mathbf{v} + c_{k-1} T^{k-1} \mathbf{v} + \cdots + c_0 \mathbf{v} = p(T)\mathbf{v}$$

for some polynomial $p(x) = c_k x^k + \cdots + c_0 \in \mathbf{F}[x]$. Hence the T-cyclic subspace generated by $\mathbf{v}$ may be alternatively presented

$$H_{T,\mathbf{v}} = \{p(T)\mathbf{v} \in V : p \in \mathbf{F}[x]\}.$$

Moreover, since $\dim H_{T,\mathbf{v}} \leq \dim V$ is finite, there is a smallest non-negative integer $k \in \mathbf{N}$ such that $\{\mathbf{v}, T\mathbf{v}, \ldots, T^k\mathbf{v}\}$ is dependent; i.e.

$$\mathbf{0} = p_{T,\mathbf{v}}(T)\mathbf{v}$$

for some polynomial $p_{T,\mathbf{v}} \in \mathbf{F}[x]$ of minimal degree k. Dividing $p_{T,\mathbf{v}}$ by its leading (non-zero) coefficient c_k, we may assume that $p_{T,\mathbf{v}}$ is monic. We summarize this discussion in a definition.

Definition 1.7. *The* T-minimal polynomial[1] *of* $\mathbf{v} \in V$ *is the monic polynomial* $p_{T,\mathbf{v}}$ *of smallest non-negative degree such that* $p_{T,\mathbf{v}}(T)\mathbf{v} = \mathbf{0}$.

There might a priori be many linear combinations of $\mathbf{v}, \dots, T^k\mathbf{v}$ that vanish, so we need to know first that this definition is not ambiguous.

Proposition 1.8. *The polynomial* $p_{T,\mathbf{v}}$ *is well-defined, and* $H_{T,\mathbf{v}} \subset \ker p_{T,\mathbf{v}}(T)$. *Moreover,* $\dim H_{T,\mathbf{v}} = \deg p_{T,\mathbf{v}}$ *and more precisely,*

$$H_{T,\mathbf{v}} = \operatorname{span}\{\mathbf{v}, T\mathbf{v}, \dots, T^{k-1}\mathbf{v}\}.$$

Later we will use a general fact about polynomials to deduce the stronger statement that if $\mathbf{v} \in \ker q(T)$, then q is actually a multiple of p. Note that it is not true in general that $H_{T,\mathbf{v}} = \ker p_{T,\mathbf{v}}(T)$. If for instance, $\mathbf{v} \in V$ is an eigenvector for the eigenvalue λ, then $H_{T,\mathbf{v}}$ is always one dimensional, generated by $\mathbf{v}$, whereas $p_{T,\mathbf{v}} = x - \lambda$, and therefore the eigenspace $\ker p_{T,\mathbf{v}}(T)$ is the full (and possibly higher dimensional) eigenspace for λ.

Proof. To see that $p_{T,\mathbf{v}}$ is well-defined, suppose that $p \in \mathbf{F}[x]$ is a second monic polynomial of degree k such that $p(T)\mathbf{v} = \mathbf{0}$, then $\deg(p - p_{T,\mathbf{v}}) < k$ since leading terms cancel, and

$$(p - p_{T,\mathbf{v}})(T)\mathbf{v} = \mathbf{0}.$$

Because $p_{T,\mathbf{v}}$ was chosen to have minimal non-negative degree, it follows that $p - p_{T,\mathbf{v}} = 0$.

Since $\ker p_{T,\mathbf{v}}(T)$ is a T-invariant subspace containing $\mathbf{v}$, Proposition 1.4 tells us that $H_{T,\mathbf{v}} \subset \ker p_{T,\mathbf{v}}(T)$.

Finally, we have by definition of $k = \deg p_{T,\mathbf{v}}$ that $\{\mathbf{v}, \dots, T^{k-1}\mathbf{v}\}$ is an independent set. Hence $\dim H_{T,\mathbf{v}} \geq k$. Suppose (in order to get a contradiction) that the inequality is actually strict. That is, there is an element $\mathbf{w} = p(T)\mathbf{v} \in H$ that is not a linear combination of the vectors $\mathbf{v}, \dots, T^{k-1}\mathbf{v}$. We can assume that $\ell := \deg p \geq k$ is as small as possible, and after dividing through by the leading coefficient, that p is monic. Then

$$\mathbf{w} = \mathbf{w} - \mathbf{0} = \mathbf{w} - p_{T,\mathbf{v}}(T)T^{\ell-k}\mathbf{v} = (p - x^{\ell-k}p_{T,\mathbf{v}})(T)\mathbf{v}.$$

However, $\deg(p - x^{\ell-k}p_{T,\mathbf{v}}) < \deg p$, contradicting the fact that p was supposed to have minimal degree. Hence $\{\mathbf{v}, \dots, T^{k-1}\mathbf{v}\}$ is actually a basis for $H_{T,\mathbf{v}}$. □

1.3. **The Cayley-Hamilton Theorem.** As a first application of the ideas we have been discussing, we now prove a very important fact about linear operators. Recall that the *characteristic polynomial* of T is

$$p_{char}(x) := \det(x\,\mathrm{id} - T) \in \mathbf{F}[x].$$

In particular $p_{char}(x)$ is a monic polynomial with degree equal to the dimension of V.

As it turns out, $\ker p_{char}(T)$ is *all* of V. To prove this we need a preliminary result.

Lemma 1.9. *If* $\mathbf{v} \in V$ *is a non-zero vector, then the* T*-minimal polynomial* $p_{T,\mathbf{v}}$ *of* $\mathbf{v}$ *is also the characteristic polynomial of the restricted operator* $T : H_{T,\mathbf{v}} \to H_{T,\mathbf{v}}$.

[1] In Hoffman and Kunze's book $p_{T,\mathbf{v}}$ is called the T-*annihilating polynomial* of $\mathbf{v}$

Proof. The degree k of $p_{T,\mathbf{v}}$ is chosen so that $\mathcal{B} := \{\mathbf{v}, \ldots, T^{k-1}\mathbf{v}\}$ is a basis for $H_{T,\mathbf{v}}$, and the coefficients of $p_{T,\mathbf{v}}(x) = x^k + c_{k-1}x^{k-1} + \cdots + c_0$ are chosen so that $T^k\mathbf{v} = -c_{k-1}T^{k-1}\mathbf{v} - \cdots - c_0\mathbf{v}$. Thus relative to $\mathcal{B}$, the restricted operator $T : H_{T,\mathbf{v}} \to H_{T,\mathbf{v}}$ has matrix

$$M := \begin{bmatrix} 0 & 0 & 0 & \ldots & 0 & -c_0 \\ 1 & 0 & 0 & \ldots & 0 & -c_1 \\ 0 & 1 & 0 & \ldots & 0 & -c_2 \\ 0 & 0 & 1 & \ldots & 0 & -c_3 \\ & & & \vdots & & \\ 0 & 0 & 0 & \ldots & 1 & -c_{k-1} \end{bmatrix}$$

Using cofactor expansion about the last column and working from the bottom entry up, we obtain the characteristic polynomial for $T|_{H_{T,\mathbf{v}}}$:

$$\det(\lambda I - M) = (\lambda + c_{k-1} \det M_{kk} - c_{k-2} \det M_{k-1,k} + \cdots + (-1)^{k+1} c_0 \det M_{1k}$$

where M_{jk} are the jk minors of M. On closer inspection, one finds that these have block diagonal form

$$M_{jk} = \begin{bmatrix} A_j & 0 \\ 0 & B_j \end{bmatrix}$$

where A_j is a $(j-1) \times (j-1)$ lower triangular matrix with all entries on the main diagonal equal to λ and B_j is a $(k-j) \times (k-j)$ upper triangular matrix with all entries on the main diagonal equal to -1. Hence we may finish computing the characteristic polynomial for $T|_{H_{T,\mathbf{v}}}$:

$$\begin{aligned} \det(\lambda I - M) &= (\lambda + c_{k-1}) \det A_k \det B_k - c_{k-2} \det A_{k-1} \det B_k + \cdots + (-1)^{k+1} c_0 \det A_1 \det B_1 \\ &= \lambda^{k-1}(\lambda + c_{k-1}) + c_{k-2}\lambda^{k-2} + \cdots + c_0 \\ &= p_{T,\mathbf{v}}. \end{aligned}$$

□

Theorem 1.10 (Cayley-Hamilton)**.** *If V is a finite dimensional vector space, and $T : V \to V$ is a linear operator, then $p_{char}(T) = 0$.*

Proof. Given $\mathbf{v} \in V$, we may apply Theorem 1.5 to $H_{T,\mathbf{v}}$ and then Lemma 1.9 to deduce that $p_{char} = p_{T,\mathbf{v}}q$ for some $q \in \mathbf{F}[x]$. Thus

$$p_{char}(T)\mathbf{v} = q(T)p_{T,v}(T)\mathbf{v} = q(T)\mathbf{0} = \mathbf{0}.$$

This proves for any non-zero $\mathbf{v} \in V$ that $p_{char}(T)\mathbf{v} = \mathbf{0}$. We conclude that $p_{char}(T)$ is the zero operator. □

2. Direct sum decompositions and linear transformations

Here we return to a topic that we touched on at the beginning of last semester. As above, we let V be a finite dimensional vector space over a field $\mathbf{F}$. Recall that the *sum* of subspaces $H_1, \ldots, H_k \subset V$ is the subspace

$$H_1 + \cdots + H_k := \{\mathbf{v}_1 + \cdots + \mathbf{v}_k \in V : \mathbf{v}_j \in H_j\}.$$

Definition 2.1. *We say that $H_1, \ldots, H_k \subset V$ are* independent *if the only vectors $\mathbf{v}_1 \in H_1, \ldots, \mathbf{v}_k \in H_k$ satisfying*

$$\mathbf{v}_1 + \cdots + \mathbf{v}_k = \mathbf{0}$$

are $\mathbf{v}_1 = \cdots = \mathbf{v}_k = \mathbf{0}$.

When $H_1, \ldots, H_k$ are independent, we say that their sum is *direct* and denote it by $H_1 \oplus \cdots \oplus H_k$. It will be useful to us below to be able to verify independence of a collection of subspaces inductively. To this end we prove

Proposition 2.2. *Suppose that $H_1, \ldots, H_{k-1} \subset V$ are independent subspaces and that $H_k \subset V$ is another subspace that intersects $H_1 \oplus \cdots \oplus H_{k-1}$ trivially. Then $H_1, \ldots, H_k$ are independent subspaces.*

Proof. Suppose that $\mathbf{v}_j \in H_j$, $1 \leq j \leq k$ are given and

$$\mathbf{v}_1 + \cdots + \mathbf{v}_k = \mathbf{0}.$$

Then

$$\mathbf{v}_1 + \cdots + \mathbf{v}_{k-1} = -\mathbf{v}_k$$

is a vector in $(H_1 \oplus \cdots \oplus H_{k-1}) \cap H_k$. But this intersection is trivial by hypothesis, so both sides of the last equation must be zero. In particular, independence of $H_1, \ldots, H_{k-1}$ and the vanishing of the left side imply that $\mathbf{v}_1 = \cdots = \mathbf{v}_{k-1} = \mathbf{0}$. Hence $H_1, \ldots, H_k$ are independent subspaces. □

As the next proposition indicates, a collection of subspaces whose direct sum is V is analogous to a basis for V.

Proposition 2.3. *Suppose that $V = H_1 \oplus \cdots \oplus H_k$. Then for any $\mathbf{v} \in V$, there are unique vectors $\mathbf{v}_1 \in H_1, \ldots \mathbf{v}_k \in H_k$ such that*

$$\mathbf{v} = \mathbf{v}_1 + \cdots + \mathbf{v}_k.$$

The vectors $\mathbf{v}_j$ in the statement of this proposition are in a sense the 'coordinates' of $\mathbf{v}$ relative to the decomposition $V = H_1 \oplus \cdots \oplus H_r$.

Proof. Given $\mathbf{v} \in V$, we have (since V is the sum of $H_1, \ldots, H_k$) $\mathbf{v}_1 \in H_1, \ldots, \mathbf{v}_k \in H_k$ such that $\mathbf{v} = \mathbf{v}_1 + \cdots + \mathbf{v}_k$.

In order to establish uniqueness, suppose that we have some other vectors $\mathbf{v}_1' \in H_1, \ldots, \mathbf{v}_k' \in H_k$ satisfying $\mathbf{v} = \mathbf{v}_1' + \ldots \mathbf{v}_k'$. Then

$$\mathbf{0} = \mathbf{v} - \mathbf{v} = (\mathbf{v}_1 - \mathbf{v}_1') + \cdots + (\mathbf{v}_k - \mathbf{v}_k').$$

Since $\mathbf{v}_j - \mathbf{v}_j' \in H_j$ it follows from independence of $H_1, \ldots, H_k$ that $\mathbf{v}_j = \mathbf{v}_j'$ for each j. That is, the vectors $\mathbf{v}_j$ are unique. □

To amplify the analogy between direct sums and bases, we offer

Proposition 2.4. *Suppose that $V = H_1 \oplus \cdots \oplus H_k$ and that for each $1 \leq j \leq k$, we are given a basis $\mathcal{B}_j$ for H_j. Then $\mathcal{B} := \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_k$ is a basis for V.*

Let us say that a basis $\mathcal{B} := \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_k$ as in the statement of this proposition is *compatible* with the decomposition $V = H_1 \oplus \cdots \oplus H_k$.

Proof. Clearly $\mathcal{B}$ spans each of the subspaces H_j, and since any vector in V can be written as a sum of vectors in $H_1, \ldots, H_k$, it follows that $\mathcal{B}$ spans V.

To see that $\mathcal{B}$ is independent, suppose that some linear combination of vectors in $\mathcal{B}$ vanishes. Taking advantage of the fact that $\mathcal{B} = \bigcup \mathcal{B}_j$, we may express this assumption as follows:

$$\sum_{j=1}^{k} \sum_{\mathbf{b} \in \mathcal{B}_j} c_{\mathbf{b}} \mathbf{b} = \mathbf{0}.$$

Since $\sum_{\mathbf{b} \in \mathcal{B}_j} c_{\mathbf{b}} \mathbf{b} \in V_j$, it follows from independence of the subspaces V_j that $\sum_{\mathbf{b} \in \mathcal{B}_j} c_{\mathbf{b}} \mathbf{b} = \mathbf{0}$ for each j separately. Since $\mathcal{B}_j$ is independent, it then follows further that $c_{\mathbf{b}} = 0$ for each $\mathbf{b} \in \mathcal{B}_j$. Thus all coefficients in the linear combination vanish, and we conclude that $\mathcal{B}$ is independent. □

Corollary 2.5. *If $V = H_1 \oplus \cdots \oplus H_k$, then $\dim V = \dim H_1 + \cdots + \dim H_k$.*

Proof. For each $j \in \{1, \ldots, k\}$, let $\mathcal{B}_j$ be a basis for H_j. Because the H_j are independent, we have $\mathcal{B}_j \cap \mathcal{B}_i = \emptyset$ for $i \neq j$. Thus by the previous proposition, we have

$$\dim V = \# \cup \mathcal{B}_j = \sum \# \mathcal{B}_j = \sum \dim H_j.$$

□

Fixing a decomposition $V = H_1 \oplus \cdots \oplus H_k$ and a compatible basis $\mathcal{B} = \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_k$, we return to the linear transformation $T : V \to V$ introduced at the beginning of these notes. We will say that the decomposition $V = H_1 \oplus \cdots \oplus H_k$ is *T-invariant* if each of the subspaces H_j involved is T-invariant.

Theorem 2.6 (Block diagonalization)**.** *Let $V = H_1 \oplus \cdots \oplus H_k$ be a T-invariant decomposition of V, and let $\mathcal{B} = \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_k$ be a compatible basis for V. For each j, let A_{jj} be the matrix relative to $\mathcal{B}_j$ of the restricted transformation $T : H_j \to H_j$, and let p_j be its characteristic polynomial. Then the matrix of T relative to $\mathcal{B}$ has block diagonal form*

$$\begin{bmatrix} A_{11} & 0 & \ldots & 0 \\ 0 & A_{22} & \ldots & 0 \\ & & \vdots & \\ 0 & 0 & \ldots & A_{kk} \end{bmatrix}.$$

In particular the characteristic polynomial of $T : V \to V$ is $p_1 \ldots p_k$.

Proof. This is most easily done by induction on the number k of subspaces in the decomposition. If $k = 1$ there is nothing to prove. We treat the case $k = 2$ separately because the induction step relies on it implicitly. In this case we have $\mathcal{B} = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ where $\mathcal{B}_1 = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ and $\mathcal{B}_2 = \{\mathbf{v}_{k+1}, \ldots, \mathbf{v}_n\}$. Since $H_1 = \operatorname{span} \mathcal{B}_1$ and $H_2 = \operatorname{span} \mathcal{B}_2$ are T-invariant, it follows as in the proof of Theorem 1.5 that the jth column of $[T]_{\mathcal{B}}$ is given by

$$[T\mathbf{v}_j]_{\mathcal{B}} = \begin{pmatrix} [T\mathbf{v}_j]_{\mathcal{B}_1} \\ \mathbf{0} \end{pmatrix} \text{ if } 1 \leq j \leq k, \text{ and } [T\mathbf{v}_j]_{\mathcal{B}} = \begin{pmatrix} \mathbf{0} \\ [T\mathbf{v}_j]_{\mathcal{B}_2} \end{pmatrix} \text{ if } k+1 \leq j \leq n.$$

Putting all the columns together gives

$$[T]_{\mathcal{B}} = \begin{bmatrix} A_1 & 0 \\ 0 & A_2 \end{bmatrix},$$

so the assertion is proved when $k = 2$.

Supposing now that the assertion is proved when $k = K - 1$, I consider the case $k = K$. I have $V = H_1 \oplus H'$ where $H' = H_2 \oplus \cdots \oplus H_K$ is also an invariant subspace. So by the case $k = 2$,

$$[T]_{\mathcal{B}} = \begin{bmatrix} A_1 & 0 \\ 0 & A' \end{bmatrix},$$

where A' is the matrix for $T : H' \to H'$ relative to the basis $\mathcal{B}' = \mathcal{B}_2 \cup \cdots \cup \mathcal{B}_K$. And by the induction hypothesis we further have

$$A' = \begin{bmatrix} A_{22} & 0 & \dots & 0 \\ 0 & A_{33} & \dots & 0 \\ & & \vdots & \\ 0 & 0 & \dots & A_{kk} \end{bmatrix},$$

so the assertion for $k = K$ follows immediately. □

With this theorem we can now better state the goal of these notes: find a T-invariant direct sum decomposition $V = H_1 \oplus \cdots \oplus H_k$ in which the dimensions of the subspaces H_j are as small as possible. As the theorem indicates, this will allow us to find a matrix representing T that is as 'diagonal' as possible. Our goal requires us to take a closer look at the characteristic polynomial of T, and indeed at polynomials generally.

3. Background concerning polynomials

In this section, we present and discuss some useful facts concerning the set $\mathbf{F}[x]$ of polynomials with coefficients in the field $\mathbf{F}$. Polynomials in $\mathbf{F}[x]$ can be added, subtracted and multiplied in the usual way, and all the relevant axioms for arithmetic hold. In contrast with $\mathbf{F}$ itself, however, there is no operation of division[2] on $\mathbf{F}[x]$. There is, however, 'division with remainder' and this is arguably the most fundamental fact concerning polynomials with field coefficients.

Theorem 3.1 (Division algorithm)**.** *For any polynomials $a(x), b(x) \in \mathbf{F}[x]$, there are unique $q(x), r(x) \in \mathbf{F}[x]$ such that $\deg r < \deg b$ and*

$$a = bq + r.$$

Proof. Let $S \subset \mathbf{F}[x]$ be the set of all polynomials of the form $a - bp$ for some $p \in \mathbf{F}[x]$. Let $r = a - bq \in S$ be a polynomial (possibly zero) of minimal degree. Suppose $\deg r(x) = k$ with leading coefficient $c_k \neq 0$ and that $\deg b(x) = \ell$ with leading coefficient c'_ℓ. If $k \geq \ell$, then

$$r(x) - \frac{c_k}{c'_\ell} x^{k-\ell} b = a - (q + \frac{c_k}{c_\ell} x^{k-\ell}) b \in S$$

is a polynomial with degree strictly smaller than k, because the leading terms in the difference on the left cancel each other. This contradicts the minimality of $\deg r$, so it must be instead that $k < \ell$. We conclude that $a = bq + r$ where $\deg r < \deg b$, as the theorem asserts.

To prove that $r, q \in \mathbf{F}[x]$ are unique, suppose that $\tilde{r}, \tilde{q} \in \mathbf{F}[x]$ also satisfy the conclusion of the theorem. Then $bq + r = b\tilde{q} + \tilde{r}$. Rearranging, we find that

$$b(q - \tilde{q}) = r - \tilde{r}.$$

[2] in mathematical parlance, this state of affairs is summarized by saying that $\mathbf{F}[x]$ is not a field, but rather a *commutative ring*.

Comparing degrees then gives

$$\deg b + \deg(q - \tilde{q}) = \deg(r - \tilde{r}) < \deg b,$$

which implies that $\deg(q - \tilde{q}) < 0$; i.e. $q = \tilde{q}$, and therefore $r = \tilde{r}$. So the polynomials $q, r \in \mathbf{F}[x]$ are unique. □

All the other results in this section, whether we prove them or not, depend ultimately on the division algorithm. The reader might note in all this that there is a very compelling analogy between polynomials and integers, with the notion of 'degree' for polynomials playing the role of 'absolute value' for integers. In particular, the notion of 'prime number' is replaced by that of 'irreducible polynomial'.

Definition 3.2. *A non-constant polynomial $p \in \mathbf{F}[x]$ is called* irreducible *if the only the polynomials in $\mathbf{F}[x]$ that divide p are constants and constant multiples of p.*

Any polynomial of degree one is irreducible. The fundamental theorem of algebra ('every complex polynomial of degree at least one has a complex root') implies that when $\mathbf{F} = \mathbf{C}$, the converse statement holds: any irreducible polynomial in $\mathbf{C}[x]$ has degree one.

For arbitrary fields, it is a tricky thing to determine whether a given polynomial of degree two or higher is irreducible. For instance $x^2 + 1$ is irreducible as a polynomial in $\mathbf{R}[x]$ but not as a polynomial in $\mathbf{C}[x]$. Likewise $x^2 - 2$ is irreducible as a polynomial in $\mathbf{Q}[x]$ but not as a polynomial in $\mathbf{R}[x]$. Keeping this in mind might make the next two theorems seem a little less 'obvious'. The hard part of each theorem is the uniqueness.

Theorem 3.3. *Given any two polynomials $a, b \in \mathbf{F}[x]$, not both equal to zero, there is a unique monic $d \in \mathbf{F}[x]$ such that $d|a$, $d|b$ and $\deg d \geq \deg \tilde{d}$ for every other $\tilde{d} \in \mathbf{F}[x]$ that divides both a and b. In fact if $\tilde{d} \in \mathbf{F}[x]$ divides both a and b, then $\tilde{d}|d$, too.*

The polynomial d is called the *greatest common divisor* of a and b and denoted $\gcd(a, b)$. If $\gcd(a, b) = 1$, then a and b are said to be *relatively prime*. It turns out, for reasons we discuss below, that $\gcd(a, b)$ is *not* very sensitive to the underlying field. For instance

$$\gcd(x^4 - 1, 3x^3 + 3x) = x^2 + 1$$

regardless of whether $x^4 - 1$ and $x^3 + x$ are though of as polynomials in $\mathbf{Q}[x]$, in $\mathbf{R}[x]$, or in $\mathbf{C}[x]$. This makes the concept of 'relatively prime polynomials' more straightforward in many cases than that of 'irreducible polynomial'.

Theorem 3.4. *Every non-constant polynomial $p(x) \in \mathbf{F}[x]$ can be factored*

$$p = q_1 \dots q_k$$

into irreducible polynomials $q_j \in \mathbf{F}[x]$. The factorization is unique except for the order and leading coefficients of the polynomials q_j.

The decomposition of p into irreducible polynomials is called the *prime factorization of p*. Often the ambiguity concerning leading coefficients in prime factorizations is addressed by requiring p and all the factors q_j to be monic. Moreover, it is common to acknowledge repeated factors explicitly in prime factorizations by writing the factorization in the alternative form

$$p = q_1^{m_1} \dots q_\ell^{m_\ell},$$

and implicitly assuming that all the q_j are distinct (i.e. $i \neq j$ implies $q_i \neq cq_j$ for any $c \in \mathbf{F}$).

Here is the concept that links the division algorithm to the previous two results.

Definition 3.5. *A non-empty set of polynomials $S \subset \mathbf{F}[x]$ is called an* ideal *if for any $a, b \in S$ and $p \in \mathbf{F}[x]$, we have that $a + b \in S$ and $ap \in S$.*

The resemblance between the notion of an 'ideal' of $\mathbf{F}[x]$ and that of a 'subspace' of a vector space is not a coincidence. The main fact concerning ideals of $\mathbf{F}[x]$ is that they are all 'one dimensional.'

Theorem 3.6. *Suppose that $S \subset \mathbf{F}[x]$ is an ideal containing at least one non-zero polynomial. Then S contains a unique (up to constant multiple) non-zero polynomial of smallest possible degree, and in fact*

$$S = p\mathbf{F}[x] := \{pq : q \in \mathbf{F}[x]\}$$

is the set of all polynomial multiples of p.

The polynomial p in the statement of this theorem is called the *generator* of S. We can (and usually do) assume with no loss of generality that p is monic.

Proof. Given p as in the theorem, we have by definition of ideal that S contains every polynomial multiple pq, $q \in \mathbf{F}[x]$ of p; i.e. that $p\mathbf{F}[x] \subset S$. Suppose now (to get a contradiction) that S contains something that is *not* a multiple of p. That is, suppose there exists $\tilde{p} \in S$ such that p does not divide $\tilde{p}$. Then by the division algorithm, we have $r, q \in \mathbf{F}[x]$ such that $\deg r < \deg p$ and $\tilde{p} = pq + r$. Since p does not divide $\tilde{p}$, it follows that $r \neq 0$. Moreover, since $r = pq - \tilde{p}$ we have from the definition of ideal that $r \in S$. That is, there is a non-zero element of S whose degree is smaller than that of p—a contradiction. We conclude that $\tilde{p}$ does not exist and that S is precisely equal to $p\mathbf{F}[x]$.

To see that p is unique, suppose that $\tilde{p} \in S$ is another non-zero polynomial of smallest degree. Then, as we have just shown, $\tilde{p} = pq$ for some $q \in \mathbf{F}[x]$. Since $\deg p = \deg \tilde{p} = \deg p + \deg q$, it follows that $\deg q = 0$. That is, $q = c_0 \in \mathbf{F}$ is a constant. □

We illustrate the power of the 'ideal' concept as follows.

Proof of Theorem 3.3. Given $a, b \in \mathbf{F}[x]$ as in the theorem, we let

$$S = \{ap + bq : p, q \in \mathbf{F}[x]\}$$

be the set of all polynomial combinations of a and b. The reader will (on pain of lightening strike for failing to comply) verify that S is an ideal of $\mathbf{F}[x]$ and that S contains a non-zero element. Hence $S = d\mathbf{F}[x]$, where $d \in S$ is the unique non-zero and monic element of smallest degree.

Then on the one hand, we have $d|a$ and $d|b$, since $a, b \in S$. And on the other hand d belongs to S, so we have by definition of S that

$$d = ap + bq$$

for some $p, q \in \mathbf{F}[x]$. From this, one may (i.e. you will now pull out pencil and paper in order to) deduce that any other common factor $\tilde{d}$ of a and b also divides d. In particular, if $\deg \tilde{d} \leq \deg d$, and if $\deg \tilde{d} = \deg d$, then $\tilde{d}$ and d are just constant multiples of one another. Hence $d = \gcd(a, b)$ is unique. □

Incidentally, the same idea leads to a very efficient method for actually *computing* greatest common divisors called the *Euclidean algorithm*. I'll be happy to provide further details in person. Beyond showing the usefulness of ideals, our discussion contains some facts that we will need later. These I summarize as follows.

Theorem 3.7. *For any non-zero polynomials $a, b \in \mathbf{F}[x]$, there are $p, q \in \mathbf{F}[x]$ such that*

$$ap + bq = \gcd(a, b).$$

In particular, if a and b are relatively prime, then there are $p, q \in \mathbf{F}[x]$ such that $ap + bq = 1$.

Returning to larger context of these notes, we define two more relevant ideals. Recall that V is a finite dimensional vector space over $\mathbf{F}$ and $T : V \to V$ is a linear operator. Then one may consider the set

$$A_T = \{p \in \mathbf{F}[x] : p(T) = 0\}$$

of polynomials that 'annihilate' T, and given $\mathbf{v} \in V$ the set

$$A_{T,\mathbf{v}} = \{p \in \mathbf{F}[x] : p(T)\mathbf{v} = 0\}$$

of polynomials p such that $p(T)$ annihilates $\mathbf{v}$. Clearly, $A_T \subset A_{T,\mathbf{v}}$. The reader will verify (or suffer greatly premature and total hair loss) that A_T and $A_{T,\mathbf{v}}$ are both ideals of $\mathbf{F}[x]$.

In terms of this new notation, the generator p_{min} of A_T is called the *minimal polynomial* of T, and in light of Theorem 3.7 the Cayley-Hamilton Theorem may be restated by saying that 'the minimal polynomial of T divides the characteristic polynomial of T'.

The minimal degree, non-zero element of $A_{T,\mathbf{v}}$ is (by definition) $p_{T,\mathbf{v}}$. Hence Theorem 3.6 implies that $p_{T,\mathbf{v}}$ divides $p_{min} \in A_T \subset A_{T,\mathbf{v}}$.

4. Primary decomposition of linear operators

A first application of our results concerning polynomials will be a kind of 'course' decomposition of the vector space V into T-invariant subspaces based on the *primary* factorization of the characteristic polynomial p_{min} for T.

Lemma 4.1. *Suppose that $p, q \in \mathbf{F}[x]$ are relatively prime polynomials. Then*

$$\ker p(T)q(T) = \ker p(T) \oplus \ker q(T).$$

Proof. The hypothesis implies that there exist $a, b \in \mathbf{F}[x]$ such that $ap + bq = 1$. Hence if $\mathbf{v} \in \ker p(T) \cap \ker q(T)$, then

$$\mathbf{v} = \mathrm{id}(\mathbf{v}) = (a(T)p(T) + b(T)q(T))\mathbf{v} = a(T)p(T)\mathbf{v} + b(T)q(T)\mathbf{v} = a(T)\mathbf{0} + b(T)\mathbf{0} = \mathbf{0}.$$

This proves that $\ker p(T) \cap \ker q(T) = \mathbf{0}$, i.e. that $\ker p(T)$ and $\ker q(T)$ are independent subspaces.

If, moreover, $\mathbf{v} = \mathbf{u} + \mathbf{w}$ where $\mathbf{u} \in \ker p(T)$ and $\mathbf{w} \in \ker q(T)$, then

$$p(T)q(T)\mathbf{v} = q(T)p(T)\mathbf{v} + p(T)q(T)\mathbf{u} = q(T)\mathbf{0} + p(T)\mathbf{0} = \mathbf{0}.$$

Hence $\ker p(T) \oplus \ker q(T) \subset \ker p(T)q(T)$.

Finally, suppose $\mathbf{v} \in \ker p(T)q(T)$. From the first paragraph, we have

$$\mathbf{v} = \mathrm{id}(\mathbf{v}) = (a(T)p(T) + b(T)q(T))\mathbf{v} = \mathbf{w} + \mathbf{u}$$

where $\mathbf{w} = a(T)p(T)\mathbf{v}$ and $\mathbf{u} = b(T)q(T)\mathbf{v}$. Thus

$$q(T)\mathbf{w} = q(T)a(T)p(T)\mathbf{v} = a(T)p(T)q(T)\mathbf{v} = a(T)\mathbf{0} = \mathbf{0},$$

so $\mathbf{w} \in \ker q(T)$. Similarly, $\mathbf{u} \in \ker p(T)$. Hence $\mathbf{v} = \mathbf{w} + \mathbf{u} \in \ker p(T) \oplus \ker q(T)$. We conclude that $\ker p(T)q(T) = \ker p(T) \oplus \ker q(T)$, as desired. □

If $q \in \mathbf{F}[x]$ is any irreducible polynomial, then we will call the set

$$\{\mathbf{v} \in V : q(T)^m \mathbf{v} = \mathbf{0} \text{ for some } m \in \mathbf{N}\}$$

the *primary subspace associated to q (and T).*

Proposition 4.2. *For any irreducible polynomial $q \in \mathbf{F}[x]$, the associated primary subspace is a T-invariant subspace of V. It is non-trivial only if q divides the minimal polynomial p_{min} of T. More specifically, if the multiplicity of q as a factor of p_{min} is m_q, then the primary subspace associated to q is just $\ker q(T)^{m_q}$.*

The last assertion amounts to noting that in the definition of primary subspace for q, one need not consider $\ker q(T)^m$ for arbitrary m but instead only $m = m_q$.

Proof. We leave this as a homework exercise, pointing out only that it is similar to the proof of Lemma 4.1 above. □

Theorem 4.3 (Primary decomposition theorem–version I)**.** *Suppose that T is a non-zero operator whose minimal polynomial p_{min} has distinct prime factors $q_1, \ldots, q_\ell \in \mathbf{F}[x]$ and associated primary subspaces $H_1, \ldots, H_\ell \subset V$. Then*

$$V = H_1 \oplus \cdots \oplus H_\ell.$$

In addition, for each factor q_j, the minimal polynomial of the restriciton $T : H_j \to H_j$ is $q_j^{m_j}$ where m_j is multiplicity of q_j as a factor of p_{min}.

Proof. We claim that for each $j = 1, \ldots, \ell$ that

$$\ker q_1(T)^{m_1} \ldots q_j(T)^{m_j} = H_1 \oplus \cdots \oplus H_j.$$

When $j = 1$, there is nothing to show. Assuming inductively that the claim is valid for all $j < J$, we consider the case $j = J$. Since $q_J^{m_J}$ is relatively prime to $q_1^{m_1} \ldots q_{J-1}^{m_{J-1}}$, we may apply Lemma 4.1 to get

$$\begin{aligned} \ker q_1(T)^{m_1} \ldots q_J(T)^{m_J} &= \ker(q_1(T)^{m_1} \ldots q_{J-1}(T)^{m_{J-1}}) \oplus \ker q_J(T)^{m_J} \\ &= (\ker q_1(T)^{m_1} \oplus \cdots \oplus \ker q_{J-1}(T)^{m_{J-1}}) \oplus \ker q_J(T)^{m_J} \\ &= \ker q_1(T)^{m_1} \oplus \cdots \oplus \ker q_{J-1}(T)^{m_{J-1}} \oplus \ker q_J(T)^{m_J} \end{aligned}$$

where the second inequality follows from the induction hypothesis. This verifies our claim. Applying it when $j = \ell$, we find that

$$V = \ker p_{min}(T) = H_1 \oplus \cdots \oplus H_\ell.$$

Turning to the final assertion in the theorem, we let p_j be the minimal polynomial for the restriction $T : H_j \to H_j$. Then $p_j | q_j^{m_j}$ by Proposition 4.2, so $p_j = q_j^{k_j}$ for some $k_j \leq m_j$. On the other hand, any vector $\mathbf{v} \in V$ can be decomposed $\mathbf{v} = \mathbf{v}_1 + \cdots + \mathbf{v}_\ell$ where $\mathbf{v}_j \in H_j$. In particular

$$p_1(T) \ldots p_j(T) \ldots p_\ell(T)\mathbf{v} = \mathbf{0}$$

simply because $p_j(T)\mathbf{v} = \mathbf{0}$ and all the factors on the left commute with each other. Hence $\mathbf{v} \in \ker(p_1 \ldots p_\ell)(T)$ for any $\mathbf{v} \in V$, and it follows that $p_{min} = q_1^{m_1} \ldots q_\ell^{m_\ell}$ must divide $p_1 \ldots p_\ell = q_1^{k_1} \ldots q_\ell^{k_\ell}$. Since the q_j are distinct irreducible polynomials and $k_j \leq m_j$, we conclude that $m_j = k_j$ for each j. That is, $p_j = q_j^{m_j}$ as asserted. □

We will next sharpen the conclusions of the Primary Decomposition Theorem by computing the dimensions of the subspaces H_j. We do this by relating H_j to the primary factors of the *characteristic* polynomial p_{char} of T. As noted at the end of the last section, Cayley-Hamilton may be restated by saying that $p_{min}|p_{char}$. In particular, every irreducible factor of p_{min} is also a factor of p_{char}. Though p_{min} need not equal p_{char}, this last statement remains true with p_{min} and p_{char} reversed.

Theorem 4.4. *p_{min} and p_{char} have the same irreducible factors (albeit with different multiplicities).*

Proof. Let $p_{char} = p_1 \dots p_k$ be the prime factorization of the characteristic polynomial. Our goal is to show that each factor p_j is also a prime factor of p_{min}. We will prove this by induction on the total number k of irreducible factors counted with multiplicity.

If $k = 0$, then $p_{char} = 1$ has no prime factors and the assertion is true vacuously.

For the induction step suppose that the assertion remains true for all $k < K$. If $k = K$, we choose any non-zero vector $\mathbf{v} \in V$ and let $H = H_{T,\mathbf{v}}$ be the cyclic subspace associated to $\mathbf{v}$. The characteristic and minimal polynomials of $T|_H$ are both equal to $p_{T,\mathbf{v}}$. Let $\tilde{p}_{char}$ and $\tilde{p}_{min}$ denote the characteristic and minimal polynomials of the induced operator $\tilde{T} : V/H \to V/H$. Then from Theorem 1.5 I have $p_{char} = p_{T,\mathbf{v}}\tilde{p}_{char}$.

There is no such simple formula relating the corresponding minimal polynomials $p_{min}, p_{T,\mathbf{v}}$, and $\tilde{p}_{min}$. However, I claim that the last two of these each divide the first. To see this, note that since $p_{min}(T) = 0$, I have in particular that $p_{min}(T)\mathbf{v} = \mathbf{0}$. Hence $p_{T,\mathbf{v}}|p_{min}$. Moreover $p_{min}(\tilde{T}) = \widetilde{p_{min}(T)} = \tilde{0} = 0$, so $\tilde{p}_{min}|p_{min}$, too. My claim is proved.

Now let p_j be any prime factor of $p_{char} = p_{T,\mathbf{v}}p_{char}$. Then either $p_j|p_{T,\mathbf{v}}$ or $p_j|\tilde{p}_{char}$. If $p_j|p_{T,\mathbf{v}}$, then $p_j|p_{min}$ by the previous paragraph. Suppose on the other hand that $p_j|\tilde{p}_{char}$. Since $\deg p_{T,\mathbf{v}} = \dim H_{T,\mathbf{v}} \geq 1$, I infer that $p_{T,\mathbf{v}}$ has at least one prime factor. Hence $\tilde{p}_{char}$ has less than K prime factors, and I may apply the induction hypothesis to see that $p_j|\tilde{p}_{min}$. Hence again $p_j|p_{min}$ by the previous paragraph. This completes the induction step and the proof. □

Theorem 4.5 (Primary decomposition–version II)**.** *Let $p_{char} = q_1^{r_1} \dots q_\ell^{r_\ell}$ be the primary factorization of the characteristic polynomial of T, and let $H_j = \ker q_j^{m_j}(T)$ as in Theorem 4.3. Then $q_1^{r_1}$ is the characteristic polynomial of the restriction $T : H_j \to H_j$, and in particular, the dimension of H_j is $n_j \deg q_j$.*

Proof. Let p_j be the characteristic polynomial of $T : H_j \to H_j$. Then $p_j = q_j^{R_j}$ for some $R_j \geq m_j$ by Cayley-Hamilton and Theorem 4.4. Moreover, by Theorem 2.6, we have $p_{char} = p_1 \dots p_\ell$. Since the q_j are distinct irreducible polynomials, it follows that $R_j = r_j$ for each j. □

It is instructive to consider the implications of Theorems 4.3 and 4.5 in the case where the underlying field $\mathbf{F}$ is $\mathbf{C}$. Then the primary decomposition of the characteristic polynomial is given by

$$p_{char}(x) = (x - \lambda_1)^{m_1} \dots (x - \lambda_\ell)^{m_\ell}.$$

Hence

$$V = \ker(T - \lambda_1 \mathrm{id})^{m_1} \oplus \dots \oplus \ker(T - \lambda_\ell \mathrm{id})^{m_\ell}.$$

Now suppose that $\mathcal{B} = \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_\ell$ is a basis for V obtained by concatenating bases for each the T-invariant subspaces $\ker(T - \lambda_j \mathrm{id})^{m_j}$. Then by Corollary 2.6, we see that the matrix for T relative to $\mathcal{B}$ has block diagonal form

$$A = \begin{bmatrix} A_1 & 0 & \dots & 0 \\ 0 & A_2 & \dots & 0 \\ & & \vdots & \\ 0 & 0 & \dots & A_k \end{bmatrix}$$

where A_j is the matrix for $T|_{\ker(T-\lambda_j \mathrm{id})^{m_j}}$ relative to $\mathcal{B}_j$. In particular $(A_j - \lambda_j I)^{m_j} = 0$. That is, $A_j = \lambda_j I + N_j$, where N_j is nilpotent (of order m_j). Reassembling we see that

$$A = \begin{bmatrix} \lambda_1 I & 0 & \dots & 0 \\ 0 & \lambda_2 I & \dots & 0 \\ & & \vdots & \\ 0 & 0 & \dots & \lambda_\ell I \end{bmatrix} + \begin{bmatrix} N_1 & 0 & \dots & 0 \\ 0 & N_2 & \dots & 0 \\ & & \vdots & \\ 0 & 0 & \dots & N_k \end{bmatrix},$$

where corresponding blocks in the two matrices each have the same sizes. Hence $A = S + N$ where S is diagonal, N is nilpotent and S and N commute. If we also use S and N to denote the linear operators on V given by these matrices, we arrive at

Theorem 4.6 (SN Decomposition)**.** *If $T : V \to V$ is a linear operator on a finite dimensional complex vector space, the $T = S + N$, where S is diagonalizable, N is nilpotent, and S and N commute.*

This theorem is very useful for computing e^A where A is a matrix with complex entries. The theorem tells us that $e^A = e^S P(N)$ where e^S is easily computed for diagonal S and P is the Taylor polynomial for e^x with degree one less than the order of the nilpotent matrix N.

5. Cyclic decomposition and Jordan Canonical Form

If the operator T is diagonalizable, then the primary subspaces of V are just the various eigenspaces for T. Of course in this case, one can further decompose an eigenspace into smaller T-invariant subspaces simply by choosing a basis. Since each element $\mathbf{v}$ of the basis is an eigenvector, it spans a one dimensional invariant subspace $H_{T,\mathbf{v}}$, and the (direct) sum of the subspaces $H_{T,\mathbf{v}}$ is the entire eigenspace.

The next theorem, whose statement and proof are the main goals of this section, says that the general situation is somewhat analogous to the diagonalizable one.

Theorem 5.1 (Cyclic decomposition theorem)**.** *Let $T : V \to V$ be a linear operator on a finite dimensional vector space over a field $\mathbf{F}$. Then there are non-zero vectors $\mathbf{v}_1, \dots, \mathbf{v}_k \in V$, positive integers $m_1, \dots, m_k$ and (not necessarily distint) monic irreducible polynomials $q_1, \dots, q_k \in \mathbf{F}[x]$ such that*

- *$p_{T,\mathbf{v}_j} = q^{m_j}$ for each $1 \le j \le k$.*
- *$V = H_{T,\mathbf{v}_1} \oplus \cdots \oplus \dots H_{T,\mathbf{v}_k}$.*

We will call any decomposition $V = H_{T,\mathbf{v}_1} \oplus \cdots \oplus \dots H_{T,\mathbf{v}_k}$ satisfying the conclusions of this theorem a *cyclic decomposition of V relative to T*. Note that each subspace $H_{T,\mathbf{v}_j}$ in the decomposition is contained in the primary subspace associated to the polynomial q. Since, as we learned in the previous section, the primary subspaces already give an invariant decomposition of T, it suffices to prove the theorem in the case where the minimal polynomial

for T has the form $p_{min} = q^m$ for some $r \geq 1$—i.e under the assumption that V is equal to the primary subspace associated to q. So until the proof of Theorem 5.1 is complete, we will operate under this assumption. The proof of this theorem is somewhat long and requires some preliminary lemmas.

Lemma 5.2. *Let $\mathbf{w} \in \ker q(T)$ be any non-zero vector. Then $p_{T,\mathbf{w}} = q$, and in particular, $\dim H_{T,\mathbf{w}} = \deg q$. Moreover, $H_{T,\mathbf{w}} = H_{T,\mathbf{w}'}$ for any non-zero vector $\mathbf{w}' \in H_{T,\mathbf{w}}$.*

Proof. Since $\mathbf{w} \neq \mathbf{0}$ and $q(T)\mathbf{w} = \mathbf{0}$, we have that $p_{T,\mathbf{w}}$ is a monic polynomial with degree at least 1 that divides q. Since q is irreducible, $p_{T,\mathbf{w}} = q$. If $\mathbf{w}' \in H_{T,\mathbf{w}}$ is another non-zero vector, then $H_{T,\mathbf{w}'} \subset H_{T,\mathbf{w}}$. But $\dim H_{T,\mathbf{w}'} = \deg q$, too, so $H_{T,\mathbf{w}'} = H_{T,\mathbf{w}}$. □

Lemma 5.3. *Let $\mathbf{v} \in V$ be any non-zero vector and let $W = H_{T,\mathbf{v}} \cap \ker q(T)$. Then $W = H_{T,\mathbf{w}}$ for any non-zero $\mathbf{w} \in W$. If $\mathbf{v} \notin W$, then for any vector $\mathbf{w} \in W$ there is a vector $\mathbf{u} \in H_{T,\mathbf{v}}$ such that $q(T)\mathbf{u} = \mathbf{w}$.*

Proof. Since $q(T)^m\mathbf{v} = \mathbf{0}$, it follows that $p_{T,\mathbf{v}}$ divides q^m and is therefore equal to q^r for some $r \leq m$. Since $\mathbf{v} \neq \mathbf{0}$, we have $r \geq 1$. Thus $\mathbf{w} = q(T)^{r-1}\mathbf{v}$ is a non-zero vector in $H_{T,\mathbf{v}}$ such that $q(T)\mathbf{w} = \mathbf{0}$, so W is non-trivial.

Clearly, $H_{T,\mathbf{w}} \subset W$. To see that the reverse inclusion holds, let $\mathbf{w}' \in W$ be any other vector. Then $\mathbf{w}' = p(T)\mathbf{v}$ for some polynomial $p \in \mathbf{F}[x]$. Moreover, $q(T)p(T)\mathbf{v} = q(T)\mathbf{w}' = \mathbf{0}$, so that q^r divides pq. Hence $p = q^{r-1}a$ for some $a \in \mathbf{F}[x]$, and

$$\mathbf{w}' = p(T)\mathbf{v} = a(T)q(T)^{r-1}\mathbf{v} = a(T)\mathbf{w} \in H_{T,\mathbf{w}}.$$

We conclude that $W \subset H_{T,\mathbf{w}}$, as desired.

Lemma 5.2 now tells us that $W = H_{T,\mathbf{w}}$ for *any* non-zero $\mathbf{w} \in W$. And if $\mathbf{v} \notin W$, then $r > 1$ and we further obtain that $\mathbf{w} = q(T)\mathbf{u}$ where $\mathbf{u} = a(T)q(T)^{r-2}\mathbf{v}$. □

Lemma 5.4. *Let $W \subset \ker q(T)$ be a T-invariant subspace. Then there are vectors $\mathbf{w}_1, \ldots, \mathbf{w}_s$ such that*

$$\ker q(T) = W \oplus H_{T,\mathbf{w}_1} \oplus \cdots \oplus H_{T,\mathbf{w}_s}.$$

The number of vectors $\mathbf{w}_j$ is uniquely determined by the formula $s = \frac{\dim \ker q(T) - \dim W}{\deg q}$.

Proof. Suppose $W' \subset \ker q(T)$ is any invariant subspace and $\mathbf{w} \in \ker q(T)$ a non-zero vector. If there exists a non-zero vector $\mathbf{w}' \in H_{T,\mathbf{w}} \cap W'$, then by invariance of W', we have $H_{T,\mathbf{w}'} \subset W$ and by Lemma 5.2 that $\mathbf{w} \in W'$. Hence we arrive at a dichotomy: either $\mathbf{w} \in W'$, or $W' \cap H_{T,\mathbf{w}}$ is trivial. Based on this, we now construct our list of vectors $\mathbf{w}_j$.

Suppose in fact that $\mathbf{w}_1, \ldots, \mathbf{w}_s \in \ker q(T)$ is *any* list of non-zero vectors such that the subspaces $W, H_{T,\mathbf{w}_1}, \ldots, H_{T,\mathbf{w}_s}$ are independent. Then using Lemma 5.2 we compare dimensions and find that

$$\dim W + s \deg q \leq \dim \ker q(T)$$

with equality if and only if $\ker q(T) = W \oplus H_{T,\mathbf{w}_1} \oplus \cdots \oplus H_{T,\mathbf{w}_s}$. In particular, the number s of vectors in the list is bounded above, and we may choose the list so that s is as large as possible. If $\ker q(T) \neq W' := W \oplus H_{T,\mathbf{w}_1} \oplus \cdots \oplus H_{T,\mathbf{w}_s}$, then we can choose a non-zero $\mathbf{w}_{s+1} \in \ker q(T) - W'$. But according to the dichotomy established in the first paragraph, $H_{T,\mathbf{w}_{s+1}} \cap W'$ is trivial. And therefore Proposition 2.2 tells us that $W, H_{T,\mathbf{w}_1}, \ldots, H_{T,\mathbf{w}_{s+1}}$ are independent subspaces, contradicting the assumption that our list of subspaces was as long as possible. We conclude that $\ker q(T) = W'$ after all. □

Having laid some groundwork, we can now proceed to the proof of our main result.

Proof of Theorem 5.1. Recall that we are assuming with no loss of generality that $p_{min} = q^m$, where $q \in \mathbf{F}[x]$ is irreducible. We will proceed by induction on m. If $m = 1$, then $V = \ker q(T)$, and the theorem follows from applying Lemma 5.4 with $W = \{\mathbf{0}\}$.

Supposing that the theorem is true for all $m < M$, we consider the case $m = M$. Let $\tilde{V} = V/\ker q(T)$. Since $\ker q(T)$ is invariant, we have an induced operator $\tilde{T} : \tilde{V} \to \tilde{V}$. Given any $\mathbf{v} \in V$, we have $q(T)^m \mathbf{v} = \mathbf{0}$, so $\mathbf{v} \in \ker q(T)^{M-1}$. Hence $q(\tilde{T})^{M-1}\tilde{\mathbf{v}} = \widetilde{q(T)^{M-1}\mathbf{v}} = \tilde{\mathbf{0}}$. Thus the minimal polynomial for $\tilde{T} : \tilde{V} \to \tilde{V}$ is q^r for some $r < M$. It follows from our induction hypothesis that there are non-zero vectors $\tilde{\mathbf{v}}_1, \ldots, \tilde{\mathbf{v}}_p \in \tilde{V}$ such that

$$\tilde{V} = H_{\tilde{T},\tilde{\mathbf{v}}_1} \oplus \cdots \oplus H_{\tilde{T},\tilde{\mathbf{v}}_p}.$$

Working back in V, we let $W = (H_{T,\mathbf{v}_1} + \cdots + H_{T,\mathbf{v}_p}) \cap \ker q(T)$, which is invariant. Lemma 5.4 gives us vectors $\mathbf{v}_{p+1}, \ldots, \mathbf{v}_k$ such that $\ker q(T) = W \oplus H_{T,\mathbf{v}_{p+1}} \oplus \cdots \oplus H_{T,\mathbf{v}_k}$. We claim that

$$V = H_{T,\mathbf{v}_1} \oplus \cdots \oplus H_{T,\mathbf{v}_k},$$

which completes the induction step.

To prove the claim, let us first show that $H_{T,\mathbf{v}_1}, \ldots, H_{T,\mathbf{v}_k}$ are independent. That is, suppose we have $\mathbf{w}_j \in H_{T,\mathbf{v}_j}$ such that

$$\mathbf{w}_1 + \cdots + \mathbf{w}_k = \mathbf{0}.$$

Then since $\mathbf{w}_{p+1}, \ldots, \mathbf{w}_k \in \ker q(T)$, we may pass to the quotient $\tilde{V}$ and obtain

$$\tilde{\mathbf{w}}_1 + \cdots + \tilde{\mathbf{w}}_p = \tilde{\mathbf{0}}.$$

Since $H_{\tilde{T},\tilde{\mathbf{v}}_1}, \ldots, H_{\tilde{T},\tilde{\mathbf{v}}_p}$ are independent, it follows that $\tilde{\mathbf{w}}_1, \ldots, \tilde{\mathbf{w}}_p = \tilde{\mathbf{0}}$. That is, $\mathbf{w}_j \in \ker q(T) \cap H_{T,\mathbf{v}_j} \subset W$ for $1 \leq j \leq p$. From independence of $W, H_{T,\mathbf{v}_{p+1}}, \ldots H_{T,\mathbf{v}_k}$, we then obtain that $\mathbf{0} = \mathbf{w}_1 + \cdots + \mathbf{w}_p = \mathbf{w}_{p+1} = \cdots = \mathbf{w}_k$.

By Lemma 5.3 moreover, there are vectors $\mathbf{u}_j \in H_{T,\mathbf{v}_j}$ such that $q(T)\mathbf{u}_j = \mathbf{w}_j$. Hence

$$q(T)(\mathbf{u}_1 + \cdots + \mathbf{u}_p) = \mathbf{w}_1 + \cdots + \mathbf{w}_p = \mathbf{0}.$$

In other words, moving to the quotient again, we get

$$\tilde{\mathbf{u}}_1 + \cdots + \tilde{\mathbf{u}}_p = \tilde{\mathbf{0}}$$

Using independence of the subspaces $H_{\tilde{T},\tilde{\mathbf{v}}_j}$ again, we infer that $\mathbf{u}_j \in \ker q(T)$ and therefore $\mathbf{w}_j = q(T)\mathbf{u}_j = \mathbf{0}$, for $1 \leq j \leq p$. This proves that $H_{T,\mathbf{v}_1}, \ldots, H_{T,\mathbf{v}_k}$ are independent subspaces.

To complete the proof of the claim, we must show $V = H_{T,\mathbf{v}_1} + \cdots + H_{T,\mathbf{v}_k}$. Given $\mathbf{v} \in V$, we have $\tilde{\mathbf{v}} = \tilde{\mathbf{u}}_1 + \ldots \tilde{\mathbf{u}}_p$ for some $\mathbf{u}_j \in H_{T,\mathbf{v}_j}$, $1 \leq j \leq p$. Hence

$$\mathbf{v} = \mathbf{u}_1 + \cdots + \mathbf{u}_p + \mathbf{w}$$

for some $\mathbf{w} \in \ker q(T)$. Additionally, $\mathbf{w} = \mathbf{w}_1 + \ldots \mathbf{w}_k$ for some $\mathbf{w}_j \in H_{T,\mathbf{v}_j} \cap \ker q(T)$. Hence $\mathbf{v} \in H_{T,\mathbf{v}_1} + \cdots + H_{T,\mathbf{v}_k}$, as desired, and the claim is proved. □

Unlike the subspaces in the Primary Decomposition Theorem, those in the Cyclic Decomposition Theorem are not uniquely determined. They depend on the choices of vectors $\mathbf{v}_j$ made along the way. However, some aspects of the subspaces do not depend on these choices.

Theorem 5.5. *Let*

$$V = H_{T,\mathbf{v}_1} \oplus \cdots \oplus H_{T,\mathbf{v}_k} = H_{T,\mathbf{w}_1} \oplus \cdots \oplus H_{T,\mathbf{w}_\ell}$$

be two cyclic decompositions of V relative to T. Then in fact $k = \ell$ and by reindexing the vectors in one of the decompositions, one may arrange that for each $1 \leq j \leq k$, the minimal polynomial for the restriction $T : H_{T,\mathbf{v}_j} \to H_{T,\mathbf{v}_j}$ is the same as that of $T : H_{T,\mathbf{w}_j} \to H_{T,\mathbf{w}_j}$. In particular, $\dim H_{T,\mathbf{v}_j} = \dim H_{T,\mathbf{w}_j}$.

As observed after the statement of Theorem 5.1, any cyclic decomposition of V relative to T refines the primary decomposition of V. Hence we may again assume that $p_{min} = q^m$, where $q \in \mathbf{F}[x]$ is monic and irreducible. Under this assumption it suffices to show only that we can reindex the vectors in the two decompositions so that $\dim H_{T,\mathbf{v}_j} = \dim H_{T,\mathbf{w}_j}$ for all $1 \leq j \leq \min\{k, \ell\}$. Everything else follows from the facts that $\dim V = \sum_{j=1}^{k} \dim H_{T,\mathbf{v}_j} = \sum_{j=1}^{\ell} \dim H_{T,\mathbf{w}_j}$ and that the minimal polynomial of the restriction $T : H_{T,\mathbf{v}} \to H_{T,\mathbf{v}}$ to a cyclic subspace is q^r where $r \deg q = \dim H_{T,\mathbf{v}}$.

We will need some more auxiliary results. Let $\tilde{V} = V/\ker q(T)$ and $\tilde{T}$ be as in the proof of Theorem 5.1.

Lemma 5.6. *The minimal polynomial of $\tilde{T}$ is q^{m-1}.*

Proof. Given any $\tilde{\mathbf{v}} \in \tilde{V}$, we have $q(T)^m \mathbf{v} = \mathbf{0}$. Hence $q(T)^{m-1}\mathbf{v} \in \ker q(T)$. Hence $q(\tilde{T})\tilde{\mathbf{v}} = \tilde{\mathbf{0}}$. From this, we conclude that the minimal polynomial of T is q^r for some $r \leq m-1$.

On the other hand, the logic reverses: since $q(\tilde{T})^r \tilde{\mathbf{v}} = \mathbf{0}$ for all $\tilde{\mathbf{v}} \in \tilde{V}$, we have $q(T)^r \mathbf{v} \in \ker q(T)$. Thus $q(T)^{r+1}\mathbf{v} = \mathbf{0}$ for all $\mathbf{v} \in V$, and from this we see that Hence $p_{min} = q^m$ where $m \leq r+1$. Reconciling the two paragraphs, we obtain $r = m-1$ as asserted. □

Lemma 5.7. *Suppose that $H_1, \ldots, H_k \subset V$ are independent subspaces, each invariant by T. For each $1 \leq j \leq k$, set*

$$\tilde{H}_j = \{\tilde{\mathbf{v}} \in \tilde{V} : \mathbf{v} \in H_j\}.$$

Then $\tilde{H}_1, \ldots, \tilde{H}_K \subset \tilde{V}$ are independent subspaces.

Proof. Given $\tilde{\mathbf{w}}_j \in \tilde{H}_j$, we may suppose $\mathbf{w}_j \in H_{T,\mathbf{v}_j}$. If

$$\tilde{\mathbf{w}}_1 + \cdots + \tilde{\mathbf{w}}_k = \tilde{\mathbf{0}},$$

then $\mathbf{w}_1 + \cdots + \mathbf{w}_k \in \ker q(T)$. That is,

$$q(T)\mathbf{w}_1 + \cdots + q(T)\mathbf{w}_k = \mathbf{0}.$$

By invariance of H_j we have $q(T)\mathbf{w}_j \in H_j$, and hence by independence of the H_j, we have $q(T)\mathbf{w}_j = \mathbf{0}$ for each j. So $\tilde{\mathbf{w}}_j = \tilde{\mathbf{0}}$ for each j, which proves that the subspaces $H_{\tilde{T},\tilde{\mathbf{v}}_j}$ are independent. □

Proof of Theorem 5.5. Again we proceed by induction on m, the case $m = 1$ being an immediate consequence of Lemmas 5.2 and 5.4.

Assume the corollary has been established for all $m < M$ and consider the case $m = M$. Suppose the subspaces $H_{T,\mathbf{v}_1}, \ldots, H_{T,\mathbf{v}_k}$ are reordered, and $p \leq k$ chosen so that $H_{T,\mathbf{v}_j} \subset \ker q(T)$ if and only if $j > p$. Then $H_{\tilde{T},\tilde{\mathbf{v}}_j}$ is trivial if and only if $j > p$. So by Lemma 5.7, $H_{\tilde{T},\tilde{\mathbf{v}}_1}, \ldots, H_{\tilde{T},\tilde{\mathbf{v}}_p}$ furnishes a cyclic decomposition of $\tilde{V}$. Likewise, we may reorder the

subspaces $H_{T,\mathbf{w}_1}, \ldots, H_{T,\mathbf{w}_\ell}$ and choose $s \leq \ell$ so that $H_{T,\mathbf{w}_{s+1}}, \ldots, H_{T,\mathbf{w}_\ell} \subset \ker q(T)$ and $H_{\tilde{T},\tilde{\mathbf{w}}_1}, \ldots H_{\tilde{T},\tilde{\mathbf{w}}_s}$ furnish a cyclic decomposition of $\tilde{V}$.

Now from Lemma 5.6 and our induction hypothesis, we have that $p = s$ and after reordering subspaces again, $\dim H_{\tilde{T},\tilde{\mathbf{v}}_j} = \dim H_{\tilde{T},\tilde{\mathbf{w}}_j}$ for all $1 \leq j \leq p$. The remaining subspaces on each list all have dimension equal to $\deg q$, so the induction step is complete and the theorem is proved. □

As a final remark about general linear operators, we should point out that a cyclic decomposition of V relative to T is the finest possible decomposition of V into T-invariant subspaces.

Theorem 5.8. *Suppose $V = H_1 \oplus \ldots H_\ell$, where $H_j \subset V$ are T-invariant subspaces. Then there is a cyclic decomposition*

$$V = H_{T,\mathbf{v}_1} \oplus \ldots H_{T,\mathbf{v}_k}$$

of V relative to T such that for each j, we have $H_{T,\mathbf{v}_j} \subset H_i$ for some i (depending on j).

Proof. Since H_j is T-invariant, we can apply Theorem 5.1 to the restriction $T : H_j \to H_j$ and obtain a cyclic decomposition

$$H_j = H_{T,\mathbf{v}_{1,j}} \oplus \cdots \oplus H_{T,\mathbf{v}_{k_j,j}}.$$

Then the subspaces $\{H_{T,\mathbf{v}_{i,j}} : 1 \leq i \leq \ell \text{ and } 1 \leq j \leq k_i\}$ form a cyclic decomposition of V relative to T. □

As with the primary decomposition theorem, we would like to see what the cyclic decomposition theorem says in the particular case when the vector space V is complex. In this case, we have that $q(x) = (x - \lambda)$ for some root $\lambda \in \mathbf{C}$ of p_{char}. So if $\mathbf{v} = \mathbf{v}_j$ is one of the vectors in the conclusion of Theorem 5.1 and $p_{T,\mathbf{v}} = q^r$, then $\{\mathbf{v}, \ldots, T^{r-1}\mathbf{v}\}$ is a basis for $H_{T,\mathbf{v}}$.

Proposition 5.9. *An alternative basis for $H_{T,\mathbf{v}}$ is*

$$\mathcal{B}' = \{\mathbf{v}, (T - \lambda)\mathbf{v}, \ldots, (T - \lambda)^{r-1}\mathbf{v}\}.$$

The matrix for $T : H_{T,\mathbf{v}} \to H_{T,\mathbf{v}}$ relative to $\mathcal{B}'$ is

$$\begin{bmatrix} \lambda & 1 & 0 & \ldots & 0 \\ 0 & \lambda & 1 & \ldots & 0 \\ 0 & 0 & \lambda & \ldots & 0 \\ & & & \vdots & \\ 0 & 0 & 0 & \ldots & 1 \\ 0 & 0 & 0 & \ldots & \lambda \end{bmatrix}$$

Proof. To prove the first assertion, it suffice to show that $\mathcal{B}'$ is merely independent, because $\#\mathcal{B}' = \dim H_{T,\mathbf{v}}$. So suppose that

$$\mathbf{0} = c_0\mathbf{v} + \cdots + c_{r-1}(T - \lambda)^{r-1}\mathbf{v} = p(T)\mathbf{v},$$

where $p(x) = \sum_{j=0}^{r-1} c_j(x - \lambda)^j$ is a polynomial of degree at most $r - 1$. On the other hand, by definition of $p_{T,v} = (T - \lambda)^r$, we have that $(x - \lambda)^r$ must divide p. The only way this can happen is if $p = 0$. Since $\{1, x - \lambda, \ldots, (x - \lambda)^{r-1}\}$ are independent polynomials, we infer that all the coefficients c_j vanish. This proves that $\mathcal{B}'$ is independent and therefore a basis.

Turning to the matrix for $T : H_{T,\mathbf{v}} \to H_{T,\mathbf{v}}$ relative to $\mathcal{B}'$, we note that

$$T(T-\lambda)^j\mathbf{v} = (T-\lambda)^{j+1}\mathbf{v} + \lambda(T-\lambda)^j\mathbf{v}.$$

Hence the jth column in $[T]_{\mathcal{B}'}$ is (presented horizontally) $[T(T-\lambda)^j\mathbf{v}]_{\mathcal{B}'} = (0,\ldots,1,\lambda,\ldots,0)$, where the λ falls in the jth entry of the column. □

Applying Proposition 5.9 to each subspace in the cyclic decomposition of each non-trivial primary subspace associated to $T : V \to V$, we arrive at the famous

Corollary 5.10 (Jordan canonical form)**.** *Suppose that $T : V \to V$ is a linear operator on a finite dimensional complex vector space V. Then there is a basis $\mathcal{B}$ for V such that the matrix T relative to $\mathcal{B}$ has block diagonal form*

$$\begin{bmatrix} A_1 & 0 & \ldots & 0 \\ 0 & A_2 & \ldots & 0 \\ & & \vdots & \\ 0 & 0 & \ldots & A_k \end{bmatrix}$$

where each matrix A_j is a $k_j \times k_j$ matrix of the form

$$\begin{bmatrix} \lambda_j & 1 & 0 & \ldots & 0 \\ 0 & \lambda_j & 1 & \ldots & 0 \\ 0 & 0 & \lambda_j & \ldots & 0 \\ & & & \vdots & \\ 0 & 0 & 0 & \ldots & 1 \\ 0 & 0 & 0 & \ldots & \lambda_j \end{bmatrix}.$$

for some root λ_j of the characteristic polynomial of T.