Notes on Canonical Forms
2012-04-24 08:40

Jeffrey Diller

April 24, 2012

1 Introduction

Throughout these notes we take $T : V \to V$ to be a linear operator on a finite dimensional[1] vector space V over a field $\mathbf{F}$. Our goal is to relate the structure of T to that of its characteristic polynomial $p_{char}(\lambda) := \det(\lambda\,\mathrm{id} - T)$. The model for this project, and the best case scenario, occurs when T is diagonalizable: we have a basis $\mathcal{B} = \{\mathbf{v}_1, \dots \mathbf{v}_n\} \subset V$, in which each $\mathbf{v}_j$ is an eigenvector for some eigenvalue $\lambda_j \in \mathbf{F}$ of T. The matrix for T relative to $\mathcal{B}$ is then diagonal, given by

$$[T]_{\mathcal{B}\mathcal{B}} = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}$$

In particular the characteristic polynomial of T factors completely: $p_{char}(\lambda) = (\lambda - \lambda_1) \dots (\lambda - \lambda_n)$.

We have seen, however, that not all linear operators are diagonalizable. There are essentially two obstacles to diagonalization, which we illustrate by example. Consider the linear operators on $S, T : \mathbf{R}^2 \to \mathbf{R}^2$ with standard matrices

$$A_S = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \qquad A_T = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix},$$

respectively. The characteristic polynomial of S is $\lambda^2 + 1$, which has no real roots and therefore no (real) eigenvectors. This leaves us two independent eigenvectors short of a basis for $\mathbf{R}^2$. Note that the problem here is more with the field $\mathbf{R}$ than with the operator. If we change our field to $\mathbf{C}$, then the characteristic polynomial has two distinct complex roots $\pm i$ and therefore two independent eigenvectors, which suffice to diagonalize S as a complex linear operator.

The characteristic polynomial of T, on the other hand, is $(\lambda - 1)^2$. So T has two real (but equal) roots $\lambda = 1, 1$. The problem is that the eigenspace of $\lambda = 1$ is only one dimensional spanned by $(1, 0)$. So this time we are one eigenvector short of a basis. This problem is much less common than the one we had with T, but it is more serious: we cannot make T

[1] The assumption $\dim V < \infty$ isn't always needed, but ultimately it's the case we care about here.

diagonalizable by simply changing the field to $\mathbf{C}$. Indeed, *any* field $\mathbf{F}$ contains elements 0 and 1, so in fact the matrix A_T defines a linear operator $T : \mathbf{F}^2 \to \mathbf{F}^2$ over any field $\mathbf{F}$, and the missing eigenvector argument shows that T is not diagonalizable over $\mathbf{F}$ either.

All of this discussion is meant to suggest that the behavior of a linear operator T is closely tied to the way in which its characteristic polynomial can or cannot be factored. To make this idea more precise requires us to replace the notion of eigenvector with something more general and flexible. Namely,

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

Exercise 1.1 *Prove each of the following.*

- *A one dimensional subspace H is invariant if and only if it is generated by an eigenvector.*
- *The eigenspace of an eigenvalue is an invariant subspace (regardless of its dimension).*
- *If $H_1, H_2 \subset V$ are invariant, then so are $H_1 \cap H_2$ and $H_1 + H_2$.*
- *If $\mathbf{F} = \mathbf{C}$, then any invariant subspace $H \subset V$ contains a one dimensional invariant subspace.*

Exercise 1.2 *Find all invariant subspaces of the operators S and T above. Do the same for the operators with standard matrices*

$$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

These illustrate all the main possibilities for linear operators on two dimensional vector spaces. Can you prove this? What further possibilities are there for linear operators on three dimensional vector spaces?

Note that if H is a T-invariant subspace, then we obtain another a linear operator $T : H \to H$ simply by limiting the domain of T to consist only of vectors in H. We call this operator the *restriction of T to H* and denote it by $T|_H$. Below, we will discuss two important ways to locate and describe invariant subspaces of V. Before any of this, however, we digress a bit and consider polynomials as objects of interest in their own right.

2 Polynomials

In this section, we present and discuss some useful facts concerning the set

$$\mathbf{F}[x] := \{c_k x^k + \cdots + c_1 x + c_0 : c_j \in \mathbf{F}\}$$

of all polynomials with coefficients in $\mathbf{F}$. We have already seen one element of $\mathbf{F}[x]$, the characteristic polynomial, associated with the linear operator T. In this section we will find another polynomial, the minimal polynomial, that is canonically associated to T.

The *degree* of a polynomial $p(x) = c_k x^k + \cdots + c_0$ is the largest power k of x for which the coefficient c_k is non-zero. We adopt the convention that $\deg 0 = -\infty$. We call a non-zero polynomial *monic* if its leading coefficient $c_k = 1$.

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 $\mathbf{F}[x]$. 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 $\deg b \leq \deg a$.

Even when b does not divide a, we may still perform 'division with remainder'. This is arguably the most important basic fact concerning polynomials with field coefficients.

Theorem 2.1 (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

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

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 2.2 *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*.*

[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*.

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 2.3 *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 2.4 *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 2.5 *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 2.6 *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 2.3. 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 2.7 *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 $ap + bq = 1$.

Now let us return to the linear operator $T : V \to V$ which is the main object considered in these notes. If $p(x) = c_n x^n + \cdots + c_0 \in \mathbf{F}[x]$ is any polynomial, then we can define a new linear operator $p(T) : V \to V$ by substituting T for the unknown x:

$$p(T) := c_n T^n + \cdots + c_0 \,\mathrm{id}.$$

Note here that T^j means the j-fold composition $T \circ T \circ \cdots \circ T$.

We note several convenient features of this construction as follows, leaving the reader the exercise of verifying them.

Proposition 2.8 *If $p, q \in \mathbf{F}[x]$ are polynomials, then*

- *every T-invariant subspace of V is also $p(T)$ invariant;*
- *$p(T) \circ q(T) = q(T) \circ p(T) = (pq)(T)$;*
- *$\ker p(T)$ is a T-invariant subspace.*
- *If $p|q$, then $\ker p(T) \subset \ker q(T)$.*

The second assertion in Proposition 2.8 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 2.8 affords us one of two basic means for finding invariant subspaces of V. Usually, $\ker p(T) = \{\mathbf{0}\}$ is trivial, but we will see below that when p is e.g. a factor of the characteristic polynomial of T, the subspace $\ker p(T)$ is more interesting. For instance, when $p(x) = x - \lambda$ for some root λ of p_{char}, then $\ker p(T)$ is just the eigenspace for λ.

A particularly important case occurs when $\ker p(T) = V$.

Proposition 2.9 *The set*

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

of polynomials that 'annihilate' T is a non-trivial ideal.

Proof. We show only that T is non-trivial, leaving the reader to verify that $\mathcal{I}_T$ is an ideal. Note that the vector space $\mathcal{L}(V)$ of linear operators on V is a finite dimensional vector space. Hence when $N = \dim \mathcal{L}(V) = (\dim V)^2$, it follows that the $N+1$ operators $\mathrm{id}, T, \ldots, T^N$ are dependent: there exists a non-trivial combination

$$c_N T^N + \cdots + c_1 T + c_0 \mathrm{id} = 0$$

that vanishes. Hence $c_N x^N + \ldots c_1 x + c_0 \in \mathcal{I}(x)$. □

Definition 2.10 *The generator of $\mathcal{I}_T$ is called the* minimal polynomial *$p_{min}(x)$ of T.*

The proof of Proposition 2.9 shows that $\deg p_{min} \leq (\dim V)^2$. We will see later that in fact $\deg p_{min} \leq \dim V$. For now, we content ourselves with the following exercise, which suggests that the minimal and characteristic polynomials of T are closely related.

Exercise 2.1 *Suppose that T is diagonalizable.*

1. *Show that if the roots of p_{char} are all distinct, then $p_{min} = p_{char}$.*
2. *Show that if the roots of p_{char} are all equal to λ (i.e. $T = \lambda \mathrm{id}$), then $p_{min}(x) = x - \lambda$.*
3. *Show most generally that if $p_{char}(x) = (x - \lambda_1)^{m_1} \dots (x - \lambda_\ell)^{m_\ell}$, where $\lambda_j \in \mathbf{F}$ are all distinct, then*
$$p_{min}(x) = (x - \lambda_1) \dots (x - \lambda_\ell).$$

Finally, give an example of an operator T (evidently not diagonalizable) for which $p_{min}(x) = (x-1)^2$.

3 Quotient Spaces

Let $H \subset V$ be any subspace of our vector space V. In this section, we show how to define a new vector space V/H whose role is similar to that of a complementary subspace. Recall that since V is finite dimensional, we always have subspaces $H' \subset V$ complementing H. However, these are generally far from being unique. The *quotient space* V/H, on the other hand, is uniquely defined, and will therefore serve as a canonical replacement for the choice of a complement.

Definition 3.1 *Two vectors $\mathbf{v}_1, \mathbf{v}_2 \in V$ are* equivalent modulo H *if their difference $\mathbf{v}_1 - \mathbf{v}_2$ belongs to H. In this case we write $\mathbf{v}_1 \sim_H \mathbf{v}_2$. To each $\mathbf{v} \in V$ we associate the subset*

$$\tilde{v} := \{\mathbf{w} \in V : \mathbf{w} \sim \mathbf{v}\},$$

which we call the equivalence class *of $\mathbf{v}$; conversely, we call any vector $\mathbf{w} \in \tilde{\mathbf{v}}$ a* representative *of $\tilde{\mathbf{v}}$. We let*

$$V/H := \{\tilde{\mathbf{v}} : \mathbf{v} \in V\}$$

denote the set of all possible equivalence classes of vectors in V. We call V/H the quotient space *of V modulo H.*

Exercise 3.1 *Do each of the following.*

1. *Show that $\sim$ is an* equivalence relation*. More precisely, for any vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$, show that $\sim$ is*
   - *reflexive: $\mathbf{v} \sim \mathbf{v}$;*
   - *symmetric: $\mathbf{v} \sim \mathbf{w}$ implies $\mathbf{w} \sim \mathbf{v}$;*
   - *transitive: $\mathbf{u} \sim \mathbf{v}$ and $\mathbf{v} \sim \mathbf{w}$ implies that $\mathbf{u} \sim \mathbf{w}$.*

2. *Show that, consequently, V/H is a* partition *of V. That is,*

   - *If two equivalence classes $\tilde{\mathbf{v}}$, $\tilde{\mathbf{w}}$ intersect, then $\tilde{\mathbf{v}} = \tilde{\mathbf{w}}$;*
   - *V is the union of all equivalence classes $\tilde{\mathbf{v}} \in V/H$.*

One might visualize V/H as a deck of cards, where $H = \tilde{\mathbf{0}}$ is the card through the origin, and any other 'card' $\tilde{\mathbf{v}}$ in the deck is obtained by translating H away from $\mathbf{0}$ by the vector $\mathbf{v}$. Note that V/H is not a subset of V but rather a set of subsets of V. The great thing is that we can add these sets to one another and multiply them by scalars in a well-defined way. The idea is deceptively simple. For any vectors $\mathbf{v}, \mathbf{w} \in V$ and any scalar $\lambda \in \mathbf{F}$, we declare

- $\tilde{\mathbf{v}} + \tilde{\mathbf{w}} = \widetilde{\mathbf{v} + \mathbf{w}}$;
- $\lambda\tilde{\mathbf{v}} = \widetilde{\lambda\mathbf{v}}$;

That is, in order to e.g. add equivalence classes, we first choose vectors representing each class, add these representatives and then take the equivalence class of the sum. The problem is one of 'well-definedness'; we need to know that the final result does not depend on which vectors we choose to represent our equivalence classes.

Theorem 3.2 *The operations $+$ and $\cdot$ are well-defined on V/H, and with these operations V/H becomes a vector space over $\mathbf{F}$. In particular,*

- *the additive identity in V/H is $\tilde{\mathbf{0}}$.*
- *the additive inverse of any vector $\tilde{\mathbf{v}} \in V/H$ is given by $-\tilde{\mathbf{v}} = \widetilde{-\mathbf{v}}$ for any*
- *the function $\pi(\mathbf{v}) = \tilde{\mathbf{v}}$ is a surjective linear transformation $\pi : V \to V/H$.*

Finally, $\dim V/H = \dim V - \dim H$.

Proof. To see that addition on V/H is well-defined, suppose that $\mathbf{v}_1, \mathbf{v}_2, \mathbf{w}_1, \mathbf{w}_2 \in V$ satisfy $\mathbf{v}_1 \sim \mathbf{v}_2$ and $\mathbf{w}_1 \sim \mathbf{w}_2$; i.e. suppose that $\tilde{\mathbf{v}}_1 = \tilde{\mathbf{v}}_2$ and $\tilde{\mathbf{w}}_1 = \tilde{\mathbf{w}}_2$. Then by definition $\mathbf{v}_1 - \mathbf{v}_2, \mathbf{w}_1 - \mathbf{w}_2 \in H$. Hence

$$(\mathbf{v}_1 + \mathbf{w}_1) - (\mathbf{v}_2 + \mathbf{w}_2) = (\mathbf{v}_1 - \mathbf{v}_2) + (\mathbf{w}_1 - \mathbf{w}_2) \in H,$$

since H is closed with respect to addition. Thus $\mathbf{v}_1 + \mathbf{w}_1 \sim \mathbf{v}_2 + \mathbf{w}_2$, and addition is well-defined. We leave it to the reader to verify that scalar multiplication is also well-defined.

The vector space axioms for V/H and linearity of the quotient map $\pi : V \to V/H$ now follow from chasing definitions and applying the vector space axioms on V. Let us verify, for instance, that scalar multiplication is distributive over addition: by definition of the vector space operations on V/H, we have

$$\lambda(\tilde{\mathbf{v}} + \tilde{\mathbf{u}}) := \lambda(\widetilde{\mathbf{v} + \mathbf{u}}) := \lambda\widetilde{(\mathbf{v} + \mathbf{u})} = \widetilde{\lambda\mathbf{v} + \lambda\mathbf{u}} := \widetilde{\lambda\mathbf{v}} + \widetilde{\lambda\mathbf{u}} := \lambda\tilde{\mathbf{v}} + \lambda\tilde{\mathbf{u}}.$$

To give another example, let us also verify that π respects scalar multiplication:

$$\pi(\lambda\mathbf{v}) := \widetilde{\lambda\mathbf{v}} := \lambda\tilde{\mathbf{v}} := \lambda\pi(\mathbf{v}).$$

We leave the other verifications to the reader.

Concerning the formula for the dimension of V/H, observe that by definition of V/H, the quotient map π is surjective and has kernel equal to H. So applying the rank theorem to π gives us that

$$\dim V/H + \dim H = \dim V.$$

□

Theorem 3.3 *Let $H' \subset V$ be any subspace complementary to H. Then the restriction $\pi : H' \to V/H$ is an isomorphism. V/H is isomorphic to any subspace $H' \subset V$ complementing H. More precisely, if $\mathcal{B}' = \{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is a basis for H', then $\tilde{\mathcal{B}} = \{\tilde{\mathbf{v}}_1, \ldots, \tilde{\mathbf{v}}_k\}$ is a basis for V/H.*

Proof. Since $\ker \pi = H$ and $H \cap H' = \{\mathbf{0}\}$, we have that $\ker \pi|_{H'}$ is trivial; i.e. $\pi|_{H'}$ is injective. Since $\dim H' = \dim V/H$, we conclude that π is also surjective and therefore and isomorphism. The final assertion proceeds from the fact that isomorphisms carry bases to bases. □

Now let us bring our linear operator $T : V \to V$ back into the picture.

Theorem 3.4 *If $H \subset V$ is a T-invariant subspace, then T 'induces' a linear operator $\tilde{T} : V/H \to V/H$ given by*

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

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

$$[T]_{\mathcal{B}\mathcal{B}} = \begin{bmatrix} [T|_H]_{\mathcal{B}_H} & * \\ 0 & [\tilde{T}]_{\tilde{\mathcal{B}}} \end{bmatrix},$$

where $\tilde{\mathcal{B}} = \{\tilde{\mathbf{v}}_{k+1}, \ldots, \tilde{\mathbf{v}}_n\}$ is the induced basis for V/H.

Proof. Note that the proposed definition of $\tilde{T}(\tilde{\mathbf{v}})$ depends which $\mathbf{v}$ we choose to represent $\tilde{\mathbf{v}}$. However, if $\mathbf{u} \sim \mathbf{v}$ is another vector representing $\tilde{\mathbf{v}}$, then $\mathbf{u} - \mathbf{v} \in H$ and by invariance $T(\mathbf{u} - \mathbf{v}) \in H$, too. Therefore, linearity of the projection map and of T give us that

$$\widetilde{T(\mathbf{u})} - \widetilde{T(\mathbf{v})} = \widetilde{T(\mathbf{u} - \mathbf{v})} = \mathbf{0}.$$

Hence $\tilde{T}$ is well-defined. Linearity of $\tilde{T}$ is inherited from T; e.g. given $\tilde{\mathbf{v}} \in V/H$ and $\lambda \in \mathbf{F}$, we have

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

Now we turn to the assertion relating the matrices for T, $T|_H$ and $\tilde{T}$. For $1 \leq j \leq k$, we have $\mathbf{b}_j \in H$ and therefore $T\mathbf{b}_j \in H$. Hence

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

So we see that the first k columns of $[T]_{\mathcal{B}\mathcal{B}}$ are $\left(\begin{array}{c} [T|_H]_{\mathcal{B}_H} \\ 0 \end{array} \right)$

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

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

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

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

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, $[T\mathbf{b}_j]_{\mathcal{B}} = \left(\begin{array}{c} * \\ \left[\tilde{T}\tilde{\mathbf{b}}_j\right]_{\tilde{\mathcal{B}}} \end{array} \right)$. This means that taken together, the last $n-k$ columns of A comprise a matrix of the form $\left(\begin{array}{c} * \\ [\tilde{T}]_{\tilde{\mathcal{B}}} \end{array} \right)$. So A has the block form asserted in the theorem.

□

Corollary 3.5 *If $H \subset V$ is a T-invariant subspace, then $p_{char}(\lambda) = p_H(\lambda)\tilde{p}(\lambda)$ where p_H is the characteristic polynomial of $T|_H$ and $\tilde{p}$ is the characteristic polynomial of $\tilde{T}$.*

In the case where $\mathbf{F} = \mathbf{C}$, we can use Theorem 3.4 to 'upper triangularize' the operator T. We don't really need the result in what follows, but it costs us very little effort to state and prove it now.

Corollary 3.6 (Schur Representation) *If V is a complex vector space, then there is a basis $\mathcal{B} \subset V$ such that $[T]_{\mathcal{B}\mathcal{B}}$ is upper triangular with diagonal entries equal to the roots of p_{char}.*

Proof. We work by induction on $\dim V$. The case $\dim V = 1$ is immediate. Supposing the assertions are true when $\dim V = n-1$, we consider the case $\dim V = n$. Since $\mathbf{F} = \mathbf{C}$, the fundamental theorem of algebra gives us a root $\lambda \in \mathbf{C}$ of p_{char}. Let $\mathbf{v}_1 \in V$ be an eigenvector with eigenvalue λ. Then $H := \operatorname{span}\{\mathbf{v}\}$ is a one-dimensional invariant subspace of V, so Theorem 3.4 gives us an induced operator $\tilde{T} : V/H \to V/H$. Since $\dim V/H = n-1$, our inductive hypothesis gives us a basis $\tilde{\mathcal{B}} := \{\tilde{\mathbf{v}}_2, \dots, \tilde{\mathbf{v}}_n\} \subset V/H$ such that $[\tilde{T}]_{\mathcal{B}\mathcal{B}}$ is upper triangular.

We claim that $\mathcal{B} := \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n\}$ is a basis for V. To see this, note that independence of $\tilde{\mathbf{v}}_2, \dots, \tilde{\mathbf{v}}_n$ means that no non-trivial combination of $\mathbf{v}_2, \dots, \mathbf{v}_n$ lies in H. In particular, no non-trivial combination can vanish or be equal to $\mathbf{v}_1$; i.e. $\mathbf{v}_2, \dots, \mathbf{v}_n$ are independent and $\mathbf{v}_1$ is not in their span. This implies that $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ are independent and, since $n = \dim V$, form a basis for V as claimed.

Finally, Theorem 3.4 gives us that

$$[T]_{\mathcal{B}\mathcal{B}} == \begin{bmatrix} \lambda & * \\ 0 & [\tilde{T}]_{\tilde{\mathcal{B}}\tilde{\mathcal{B}}} \end{bmatrix},$$

which, since the bottom right block is upper triangular, is itself upper triangular. □

4 The Cayley-Hamilton Theorem

We saw above that the kernel of $p(T)$ for any $p \in \mathbf{F}[x]$ is a T-invariant subspace. We begin this section by describing a second way to construct invariant subspaces. Let $\mathbf{v} \in V$ be any vector, and consider the *forward orbit* $\mathbf{v}, T\mathbf{v}, T^2\mathbf{v}, \dots$ of $\mathbf{v}$ under T. This is typically an infinite set of vectors, but since V is finite dimensional, the finite segment $\mathbf{v}, \dots, T^k\mathbf{v}$ will be linearly dependent for $k \in \mathbf{N}$ large enough. Taking k to be the first such integer, we have that $\{\mathbf{v}, \dots, T^{k-1}\mathbf{v}\}$ is independent and therefore a basis for the subspace $H_\mathbf{v}$ that it spans. We call $H_\mathbf{v}$ the *cyclic subspace generated by T and* $\mathbf{v}$.

Theorem 4.1 *The cyclic subspace $H_\mathbf{v}$ generated by $\mathbf{v}$ is T-invariant, and the minimal and characteristic polynomials of $T|_{H_\mathbf{v}}$ are the same.*

Proof. By definition of $H_\mathbf{v}$, any vector $\mathbf{w} \in H_\mathbf{v}$ can be written $\mathbf{w} = c_0\mathbf{v} + \dots + c_{k-1}T^{k-1}\mathbf{v} = p(T)\mathbf{v}$ where $p(x) = c_0 + c_1 x + \dots c_{k-1}x^{k-1} \in \mathbf{F}[x]$ is a polynomial of degree less than k. Our choice of k implies that $T^k\mathbf{v} \in H_\mathbf{v}$. Hence,

$$T\mathbf{w} = T \circ p(T)\mathbf{v} = p(T)(T\mathbf{v}) = c_0 T\mathbf{v} + \dots + c_{k-1}T^k\mathbf{v} \in H_\mathbf{v},$$

too, since $H_\mathbf{v}$ is closed with respect to linear combinations. This shows that $H_\mathbf{v}$ is T-invariant.

Moreover, taking $\mathbf{w} = T^k\mathbf{v}$, and setting $p_\mathbf{v}(x) = x^k - p(x)$, we see that $p_\mathbf{v}(T)\mathbf{v} = T^k\mathbf{v} - p(T)\mathbf{v} = \mathbf{0}$. We will show that p is both the minimal and the characteristic polynomial of the restricted operator $T|_{H_\mathbf{v}}$.

To see that $p_\mathbf{v}$ is the minimal polynomial, first observe that if $p \in \mathbf{F}[x]$ is any non-zero polynomial with $\deg p < \deg p_\mathbf{v} = k$, then $p(T)\mathbf{v}$ is a non-trivial linear combination of the basis vectors $\mathbf{v}, T\mathbf{v}, \dots, T^{k-1}\mathbf{v}$ for $H_\mathbf{v}$. Hence $p(T)\mathbf{v} \neq \mathbf{0}$ and in particular $p(T)|_{H_\mathbf{v}}$ does not vanish (i.e. is not the zero operator). On the other hand, if $\mathbf{w} = p(T)\mathbf{v}$ is any vector in $H_\mathbf{v}$, then

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

So $p_\mathbf{v}$ is a monic polynomial with minimal degree among polynomials $p \in \mathbf{F}[x]$ such that $p(T)|_{H_\mathbf{v}} = 0$; i.e. $p_\mathbf{v}$ is the minimal polynomial of $T|_{H_\mathbf{v}}$

Let us now rewrite $p_\mathbf{v}(x) = x^k + c_{k-1}x^{k-1} + \dots + c_0$ (this amounts to reversing the signs of the scalars c_j at the beginning of the proof). It remains to show that $p_\mathbf{v}(x) = \det(x\mathrm{id} - T)$. To this end, we ask the reader to verify that the matrix of T relative to the basis $\{\mathbf{v}, T\mathbf{v}, \dots, T^{k-1}\mathbf{v}\}$ is

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

We further ask that the reader use e.g. cofactor expansion about the last column of the matrix $xI - A$ to compute (this takes some concentration but probably less concentration than following someone elses version of the computation) that

$$\det(xI - A) = p_\mathbf{v}$$

as asserted. □

We will call any subspace $H \subset V$ *cyclic* if it is the cyclic subspace associated to some vector $\mathbf{v} \in V$. We call the operator T cyclic if $V = H_{\mathbf{v}}$ is itself a cyclic subspace. Though we will not prove it here, most linear operators are cyclic.

Exercise 4.1 *Verify the following.*

1. *For any $\mathbf{v} \in V$, the set $\mathcal{I}_{\mathbf{v}} := \{p \in \mathbf{F}[x] : p(T)\mathbf{v} = \mathbf{0}\}$ is an ideal generated by $p_{\mathbf{v}}$.*

2. *$p_{\mathbf{v}}|p_{min}$ for any $\mathbf{v} \in V$.*

3. *A one dimensional subspace is cyclic if and only if it is spanned by an eigenvector.*

4. *A diagonalizable linear operator is cyclic if and only if it has no repeated eigenvalues (in the 'if' direction, it helps to use a fact about Vandermonde determinants).*

5. *If the minimal polynomial of T is irreducible, then T is cyclic.*

The fact that $p_{min} = p_{char}$ when T is cyclic suggests that there might be a close relationship between the minimal and characteristic polynomials for more general T. The next result describes this relationship in general.

Theorem 4.2 *For any linear operator $T : V \to V$ on a finite dimenionsal vector space V we have $p_{min}|p_{char}$. Conversely, any irreducible factor p of p_{char} must also divide p_{min}.*

To put it slightly differently, the prime decompositions of p_{min} and p_{char} have the same irreducible factors, but the multiplicities of the factors of p_{char} can be larger. As an immediate consequence of the first assertion in Theorem 4.2 and the fourth item in Propostion 2.8 we obtain the well-known *Cayley-Hamilton Theorem.*

Corollary 4.3 $p_{char}(T) = 0$.

Proof. Given any $\mathbf{v} \in V$, Theorems 3.4 and 4.1 together imply that that $p_{char} = p_{\mathbf{v}}q$ for some $q \in \mathbf{F}[x]$. Thus

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

This proves for any non-zero $\mathbf{v} \in V$ that $p_{char}(T)\mathbf{v} = \mathbf{0}$; i.e. $p_{char}(T)$ is the zero operator. Hence $p_{char} \in \mathcal{I}_T$ which means that $p_{min}|p_{char}$.

We prove the second assertion in the theorem by induction on $\dim V$. If $\dim V = 1$, then $T = \lambda\mathrm{id}$ for some $\lambda \in \mathbf{F}$, and one verifies the assertion readily. So assume the assertion is true whenever $\dim V < n$, and consider the case $\dim V = n$. Let $p \in \mathbf{F}[x]$ be an irreducible factor of p_{char} and let $\mathbf{v} \in V$ be a non-zero vector. Since $p_{char} = p_{\mathbf{v}}\tilde{p}_{char}$ where $\tilde{p}_{char}$ is the characteristic polynomial for the induced operator $\tilde{T}$ on $V/H_{\mathbf{v}}$, we have that $p|p_{\mathbf{v}}$ or $p|\tilde{p}_{char}$.

In the first case, we have from the exercise above that $p_{\mathbf{v}}|p_{min}$. Since divisibility is transitive, it follows that $p|p_{min}$ as desired. In the second case, the fact that $\mathbf{v} \neq \mathbf{0}$ implies that $\dim V/H_{\mathbf{v}} = \dim V - \dim H_{\mathbf{v}} \leq n - 1$. So our inductive hypothesis and $p|\tilde{p}$ implies that p divides the minimal polynomial $\tilde{p}_{min}$ of $\tilde{T}$. But we saw in ???? that $\tilde{p}_{min}|p_{min}$, so transitivity again gives $p|p_{min}$. This completes the induction step and the proof □

In closing this section we note an interesting consequence of the Theorem 4.2, whose proof we leave as an exercise for the reader.

Corollary 4.4 *Let $p \in \mathbf{F}[x]$ be any polynomial. Then $p(T)$ is invertible if and only p and p_{char} are relatively prime.*

5 Direct Sums

Definition 5.1 *We say that the subspaces $H_1, \ldots, H_k \subset V$ 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 $\mathbf{v}_1 = \cdots = \mathbf{v}_k = \mathbf{0}$.

A sum of independent subspaces is called a *direct sum*, and one generally signifies a direct sum by writing $H_1 \oplus \cdots \oplus H_k$ instead of $H_1 + \cdots + H_k$. Independent subspaces are like independent vectors, except that a collection of independent subspaces can (repeatedly!) include the trivial subspace, whereas an independent set of vectors cannot include the zero vector.

Exercise 5.1 *Verify the following:*

1. *$H_1, \ldots, H_k$ are independent if and only if any $\mathbf{v} \in H_1 + \cdots + H_k$ can be uniquely expressed $\mathbf{v} = \mathbf{v}_1 + \cdots + \mathbf{v}_k$ with $\mathbf{v}_j \in H_j$.*
2. *A set of non-zero vectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\} \subset V$ is linearly independent if and only if* $\operatorname{span} \mathbf{v}_1, \ldots, \operatorname{span} \mathbf{v}_k$ *are independent subspaces.*
3. *Subspaces $H, H' \subset V$ are complementary if and only if $V = H \oplus H'$.*

It will be useful to us below to be able to verify independence of a collection of subspaces inductively. So to this end we prove

Proposition 5.2 *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. □

When $k = 3$, Proposition 5.2 says that $(H_1 \oplus H_2) \oplus H_3 = H_1 \oplus H_2 \oplus H_3$, so $\oplus$ is associative (and clearly also commutative).

Exercise 5.2 *Let $H_1, \ldots, H_k \subset V$ be subspaces satisfying $H_i \cap H_j = \{\mathbf{0}\}$ when $i \neq j$. Does it follow that the subspaces are independent? Prove or give a counterexample.*

A collection of subspaces whose direct sum is V is analogous to a basis for V. The next proposition amplifies this analogy.

Proposition 5.3 *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 (hence H_j is non-trivial). Then these bases are pairwise disjoint and $\mathcal{B} := \mathcal{B}_1 \cup \cdots \cup \mathcal{B}_k$ 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. If two of the bases $\mathcal{B}_j$ and $\mathcal{B}_k$ have a vector $\mathbf{v}$ in common, then $\mathbf{v} \in H_j$ and $-\mathbf{v} \in H_k$. Since $\mathbf{v} + (-\mathbf{v}) = \mathbf{0}$, it follows that H_j and H_k are not independent, contrary to assumption. So the bases are pairwise disjoint.

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 5.4 *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. The previous proposition gives

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

The first equality follows from the fact that $\cup \mathcal{B}_j$ is a basis for V; the second equality follows from the fact that the bases $\mathcal{B}_j$ are mutually disjoint. □

Now we return to the linear transformation $T : V \to V$ introduced at the

Definition 5.5 *A T-*invariant decomposition *of V is a collection of invariant subspaces H_j, $1 \leq j \leq k$ such that $V = \bigoplus_{j=1}^{k} H_j$. The decomposition is* proper *if $H_j \neq V$ for any j.*

Theorem 5.6 (Block diagonalization) *Let $V = H_1 \oplus \cdots \oplus H_k$ be a T-invariant decomposition of V into non-trivial subspaces H_j, 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*

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

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

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 $\mathcal{B} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ where $\mathcal{B}_1 = \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ and $\mathcal{B}_2 = \{\mathbf{v}_{k+1}, \dots, \mathbf{v}_n\}$. Hence for any $\mathbf{v} \in V$, we have

$$[\mathbf{v}]_{\mathcal{B}} = \begin{pmatrix} [\mathbf{v}]_{\mathcal{B}_1} \\ [\mathbf{v}]_{\mathcal{B}_2} \end{pmatrix},$$

where the top or bottom component vanishes if $\mathbf{v} \in H_1$ or $\mathbf{v} \in H_2$, respectively. So since $H_1 = \operatorname{span} \mathcal{B}_1$ and $H_2 = \operatorname{span} \mathcal{B}_2$ are T-invariant, the jth column of $[T]_{\mathcal{B}}$ is given by

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

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 proper invariant decomposition $V = H_1 \oplus \cdots \oplus H_k$ into non-trivial subspaces H_j that 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. Before moving on, though, it is worth recasting the content of this section in more 'functorial' terms. We leave the reader to verify the following fact, which is a more or less immediate consequence of definitons.

Proposition 5.7 *Let $H_1, \ldots, H_k \subset V$ be subspaces. Then $\sigma(\mathbf{v}_1, \ldots, \mathbf{v}_k) = \mathbf{v}_1 + \cdots + \mathbf{v}_k$ defines a linear transformation $\sigma : H_1 \times \cdots \times H_k \to V$. In addition,*

- *σ is injective if and only if $H_1, \ldots, H_k$ are independent.*
- *σ is surjective if and only if $H_1 + \cdots + H_k = V$.*

Hence $V = H_1 \oplus \cdots \oplus H_k$ if and only if σ is an isomorphism.

This proposition allows one to infer Corollary 5.4 and Proposition 5.3 from corresponding facts about product vector spaces. Chasing definitions also allows one to give a functorial (or 'coordinate-free') version of Theorem 5.6.

Proposition 5.8 *Let $H_1, \ldots, H_k \subset V$ and σ be as in Proposition 5.7. Suppose additionally that $T : V \to V$ is a linear operator. If each of the subspaces H_j is T-invariant, then there is a well-defined linear operator $S : H_1 \times \cdots \times H_k \to H_1 \times \cdots \times H_k$ given by $S(\mathbf{v}_1, \ldots, \mathbf{v}_k) = (T\mathbf{v}_1, \ldots, T\mathbf{v}_k)$ and satisfying $\sigma \circ S = T \circ \sigma$. If $V = H_1 \oplus \cdots \oplus H_k$, then $T = \sigma \circ S \circ \sigma^{-1}$ is similar to S.*

6 Primary Decomposition

Having dwelt on cyclic subspaces in order to prove Theorem 4.2, we return to our other means of identifying invariant subspaces.

Definition 6.1 *Let $p \in \mathbf{F}[x]$ be any irreducible polynomial. The* primary subspace *associated to p and T is*

$$H_p := \{\mathbf{v} \in V : p(T)^k \mathbf{v} = \mathbf{0} \textit{ for some } k \in \mathbf{N}\} = \bigcup_{k \in \mathbf{N}} \ker p(T)^k$$

We leave the reader to verify that

Proposition 6.2 *H_p is a T-invariant subspace. It is non-trivial if and only if $p|p_{min}$.*

The main goal of this section is to establish the following important connection between factors of the minimal/characteristic polynomials and decompositions of V into invariant subspaces.

Theorem 6.3 (Primary Decomposition Theorem) *Let $p_1, \ldots, p_\ell$ be the distinct irreducible factors of p_{min}. Then we have a T-invariant decomposition*

$$V = H_{p_1} \oplus \cdots \oplus H_{p_\ell}$$

of V into the corresponding primary subspaces. The minimal polynomial of the restriction $T|_{H_{p_j}}$ is the largest power of $p_j^{r_j}$ dividing p_{min}, and the characteristic polynomial is the largest power $p_j^{m_j}$ dividing p_{char}. Hence $\dim H_{p_j} = m_j \deg p_j$.

In contrast with the cyclic decomposition theorem to be stated and proven later, the primary decomposition of a linear operator is canonical, completely determined by the field $\mathbf{F}$, the vector space V and the linear operator T. Proving Theorem 6.3 requires a preliminary result that is interesting all by itself.

Lemma 6.4 *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. $\square$

Proof of Theorem 6.3. The unique factorization theorem for polynomials allows us to write $p_{min} = p_1^{r_1} \dots p_\ell^{r_\ell}$, where the factors p_j are distinct irreducible polynomials and the multiplicities r_j are positive integers. Theorem 4.2 tells us that $p_{char} = p_1^{m_1} \dots p_\ell^{m_\ell}$, where $m_j \geq r_j$ satisfy $\sum m_j \deg p_j = \dim V$. We will prove our assertions by induction on the number ℓ of distinct irreducible factors.

If $\ell = 1$, there is nothing to do except invoke Theorem 4.2 to see that when $p_{min} = p_1^{r_1}$, one also has $p_{char} = p_1^{m_1}$ for some $m_1 \geq r_1$.

Assuming the theorem is true when $\ell = k - 1$, we proceed to the case $\ell = k$. The polynomials $p = p_1^{r_1} \dots p_{\ell-1}^{r_1}$ and $q = p_\ell^{r_\ell}$ are relatively prime, so Lemma 6.4 tells us that

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

The minimal polynomial for $T|_{H_{p_\ell}}$ divides p_{min} and is therefore equal to p_ℓ^r for some $r < r_\ell$. To see that $r = r_\ell$, Note that if $\mathbf{v} \in V$ is any vector, then $\mathbf{v} = \mathbf{u} + \mathbf{w}$ where $\mathbf{u} \in \ker p(T)$ and $\mathbf{w} \in \ker p_\ell^r(T)$. Thus

$$p(T)p_\ell^r(T)\mathbf{v} = p(T)p_\ell^r(T)\mathbf{u} + p(T)p_\ell^r(T)\mathbf{w} = p_\ell^r(T)\mathbf{0} + p(T)\mathbf{0} = \mathbf{0}.$$

Hence $r \deg p_\ell + \deg p = \deg p_\ell^r p \geq \deg p_{min} = r_\ell \deg p_\ell + \deg p$, which means $r \geq r_\ell$, too. We conclude that $r = r_\ell$ as desired.

One shows similarly, that p is the minimal polynomial of $T|_{\ker p(T)}$. Our inductive hypothesis therefore implies that

$$\ker p(T) = H_{p_1} \oplus \cdots \oplus H_{p_{\ell-1}},$$

with the minimal and characteristic polynomials of $T|_{H_{p_j}}$ as described in the Theorem. Putting together the previous two paragraphs completes the inductive step and the proof. □

It is instructive to consider the implications of Theorem 6.3 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

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

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 5.6, 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 6.5 (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.

The interested reader can verify a couple of further facts describing the relationship between the factors in the primary decomposition and more general invariant subspaces of V.

Exercise 6.1 *Let p_j be the irreducible factors of p_{min} and $H \subset V$ be any invariant subspace. Then*

1. *The primary decomposition of the restricted operator $T|_H$ is $H = \bigoplus_{j=1}^{\ell} H \cap H_{p_j}$, where some of the intersections might be trivial.*
2. *The primary decomposition of the induced operator $\tilde{T}$ on V/H is $V/H = \bigoplus_{j=1}^{\ell} \widetilde{H_{p_j}}$, where $\tilde{H}_{p_j} := \{\tilde{\mathbf{v}} : \mathbf{v} \in H_{p_j}\}$ might be trivial for some j.*

7 Cyclic Decomposition

In the wake of the Primary Decomposition Theorem, one might reasonably ask if we can do better. That is, can we further decompose the primary subspaces H_{p_j} into smaller invariant subspaces. Certainly this is possible in some cases. Suppose for instance $p_j(x) = x - \lambda$ appears as factor of p_{min} with multiplicity $r_j = 1$. Then H_{p_j} is just the eigenspace for the eigenvalue λ. If, moreover, $m := \dim H_{p_j} > 1$, then any basis $\{\mathbf{v}_1, \dots, \mathbf{v}_m\}$ for H_{p_j} gives us a decomposition

$$H_{p_j} = H_{\mathbf{v}_1} \oplus \cdots \oplus H_{\mathbf{v}_m}$$

into the one dimensional cyclic subspaces generated by the eigenvectors $\mathbf{v}_i$. Observe that this decomposition is not uniquely determined, since a different choice of basis results in a different decomposition. Nor are the subspaces $H_{\mathbf{v}_j}$ realizable as $\ker p(T)$ for some $p \in \mathbf{F}[x]$. After all, p would have to divide the minimal polynomial for $T|_{H_{p_j}}$ which is $x - \lambda$, which implies $p(x) = x - \lambda$.

Definition 7.1 *Suppose that $\mathbf{v}_1, \dots, \mathbf{v}_k \in V$ are vectors such that*

- $V = \bigoplus_{j=1}^{k} H_{\mathbf{v}_j}$
- *for each j, there is a positive integer s_j and an irreducible $p_j \in \mathbf{F}[x]$ such that $p_{\mathbf{v}_j} = p_j^{s_j}$.*

Then we say that the subspaces $H_{\mathbf{v}_j}$ give a cyclic decomposition of V relative to T. *We call the subspaces $H_{\mathbf{v}_j}$ the* factors *in the decomposition.*

The requirement that $p_{\mathbf{v}_j}$ is a prime power implies that $H_{\mathbf{v}_j} \subset H_{p_j}$ is contained in a primary subspace. This turns out to guarantee that the decomposition in the definition is as fine as possible.

Definition 7.2 *An invariant subspace $H \subset V$ is* irreducible *if there is no proper invariant decomposition $H = H_1 \oplus H_2$.*

Proposition 7.3 *For any $\mathbf{v} \in V$, the subspace $H_{\mathbf{v}}$ is irreducible if and only if $p_{\mathbf{v}} = p^s$ for $p \in \mathbf{F}[x]$ irreducible and $s \geq 0$.*

Proof. If $p_{\mathbf{v}}$ is not a prime power, then we can write $p_{\mathbf{v}} = pq$, where p and q are non-constant relatively prime polynomials. Lemma 6.4 then gives us an invariant decomposition

$$H_{\mathbf{v}} = (\ker p(T) \cap H_{\mathbf{v}}) \oplus (\ker q(T) \cap H_{\mathbf{v}}).$$

Since $p_{\mathbf{v}}$ is the minimal polynomial associated to $T|_{H_{\mathbf{v}}}$, it follows that neither factor in the decomposition is trivial, and therefore neither equals H. So H is reducible.

Suppose instead that $p_{\mathbf{v}} = p^s$ is a prime power. If $H_1, H_2 \subset H_{\mathbf{v}}$ are invariant subspaces, then the minimal polynomials p_j of $T|_{H_j}$ must divide p^s; i.e. $p_j = p^{s_j}$ for $s_j \leq s$. Since $p_{\mathbf{v}}$ is also the characteristic polynomial of $T|_{H_{\mathbf{v}}}$, we have $\dim H_{\mathbf{v}} = s \deg p$. So if $H_1, H_2 \neq H_{\mathbf{v}}$, we see that $s_j \deg p \leq \dim H_j < \dim H_{\mathbf{v}} = s \deg p$; i.e. $s' = \max\{s_1, s_2\} < s$. Now if $H = H_1 + H_2$ and $\mathbf{w} \in H$ is any vector, we write $\mathbf{w} = \mathbf{w}_1 + \mathbf{w}_2$ where $\mathbf{w}_j \in H_j$ and see that

$$p^{s'}(T)\mathbf{w} = p^{s'}(T)\mathbf{w}_1 + p^{s'}(T)\mathbf{w}_2 = \mathbf{0} + \mathbf{0}.$$

That is, $p^{s'}(T) = 0$ on $H_{\mathbf{v}}$, contradicting the fact that p^s is the minimal polynomial of $T|_H$. We conclude that $H_{\mathbf{v}}$ is irreducible. □

Theorem 7.4 (Cyclic decomposition theorem) *Any linear operator $T : V \to V$ on a finite dimensional vector space V admits a cyclic decomposition. Any two cyclic decompositions*

$$\bigoplus_{j=1}^{k} H_{\mathbf{v}_j} = V = \bigoplus_{j=1}^{k} H_{\mathbf{w}_j}$$

of V relative to T are equivalent (after reordering factors if necessary) in the sense that each has the same number of factors and there is an isomorphism $\phi : V \to V$ such that $\phi(\mathbf{v}_j) = \mathbf{w}_j$ and $\phi \circ T = T \circ \phi$. Hence $\phi(H_{\mathbf{v}_j}) = H_{\mathbf{w}_j}$ and $p_{\mathbf{v}_j} = p_{\mathbf{w}_j}$.

The rest of this section will be devoted to the proof of Theorem 7.4. The proof requires some warm-up discussion. Since each subspace $H_{\mathbf{v}_j}$ in a cyclic decomposition lies inside some primary subspace H_{p_j}, and since Theorem 6.3 tells us that V is a direct sum of primary subspaces, it suffices to work with the restriction of T to a single primary subspace. That is, we assume henceforth that the minimal polynomial for T has the form $p_{min} = p^r$ for some $r \geq 1$, and therefore V is equal to the primary subspace associated to p. We call r the *order* of T.

From here, there are two central ideas behind our arguments. The first is that since $p(T)$ is nilpotent on V, it is easier to work as much as possible in terms of $p(T)$ rather than T. The second is to work inductively, reducing the order of T by considering the induced operator $\tilde{T} : \tilde{V} \to \tilde{V}$ where $\tilde{V} := V/\ker p(T)$. The reader should check definitions to verify that the order of $\tilde{T}$ is $r-1$. The remainder of our preliminary discussion is aimed at making it possible for us to use the two main ideas laid out here.

Let $\mathbf{v} \in V$ be any vector. The minimal polynomial $p_{\mathbf{v}}$ of $T|_{H_{\mathbf{v}}}$ must divide p_{min}; hence $p_{\mathbf{v}} = p^s$ for some $s \leq r$. That is, s is the smallest non-negative integer such that $p(T)^s\mathbf{v} = \mathbf{0}$. We call s the *order* of $\mathbf{v}$. Observe that s is positive if and only if $\mathbf{v} \neq \mathbf{0}$, that $s \geq 1$ if and only if $\mathbf{v} \notin \ker p(T)$, and that the order of $p(T)^k s$ is $s - k$ for all $0 \leq k \leq s$.

Lemma 7.5 *For any $\mathbf{v} \in V$, the cyclic subspace $H_\mathbf{v}$ is invariant under $p(T)$. If $\mathbf{v}$ has order $s \geq 1$, then*

$$p(T)^{s-1} H_\mathbf{v} = H_\mathbf{v} \cap \ker p(T) = H_{p(T)^{s-1}\mathbf{v}}.$$

In particular, if $s \geq 2$, then any $\mathbf{w} \in H_\mathbf{v} \cap \ker p(T)$ can be written $\mathbf{w} = p(T)\mathbf{u}$ for some $\mathbf{u} \in H_\mathbf{v} - \ker p(T)$.

Proof. The first item in Proposition 2.8 tells us that $H_\mathbf{v}$ is $p(T)$-invariant. Since $H_\mathbf{v} \subset \ker p(T)^s$ and $p(T)^{s-1}H_\mathbf{v}$ is spanned by vectors $p(T)^{s-1}T^j\mathbf{v}$, $j \geq 0$, we have $H_{p(T)^{s-1}\mathbf{v}} = p(T)^{s-1}H_\mathbf{v} \subset \ker p(T) \cap H_\mathbf{v}$. To establish the remaining inclusion, let $\mathbf{w} \in H_\mathbf{v} \cap \ker p(T)$ be any vector. By defintion of $H_\mathbf{v}$, we have $w = q(T)\mathbf{v}$ for some $q \in \mathbf{F}[x]$. We can therefore use $p(T)\mathbf{w} = p(T)q(T)\mathbf{v} = \mathbf{0}$ to infer that $p^s | qp$, i.e. $q = ap^{s-1}$ for some $a \in \mathbf{F}[x]$. Finally, taking $\mathbf{u} = a(T)\mathbf{v} \in H_\mathbf{v}$, we see that $\mathbf{w} = p(T)^{s-1}\mathbf{u} \in p(T)^{s-1}H_\mathbf{v}$. We conclude that $H_\mathbf{v} \cap \ker p(T)^{s-1} = p(T)H_\mathbf{v}$ as desired. □

The next three lemmas will allow us to prove the existence of a cyclic decomposition for V. Given a subspace $H \subset V$, we set $\tilde{H} := \{\tilde{\mathbf{v}} \in V/\ker p(T) : \mathbf{v} \in H\}$ be the image of H in the $\tilde{V}$ under the quotient map. Note in particular that $\widetilde{H_\mathbf{v}} = H_{\tilde{\mathbf{v}}}$ for any $\mathbf{v} \in V$.

Lemma 7.6 *Given vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k \in V - \ker p(T)$, the associated cyclic subspaces $H_{\mathbf{v}_j}$ are independent if and only if their images $\widetilde{H_{\mathbf{v}_j}} \subset \tilde{V}$ under the quotient map are independent.*

Proof. We leave the reader to check that if $H_{\mathbf{v}_j}$ are independent in V, then $\widetilde{H_{\mathbf{v}_j}}$ are independent in $\tilde{V}$. For this, one can even allow $\tilde{\mathbf{v}}_j \in \ker p(T)$.

Assume instead then that the subspaces $\widetilde{H_{\mathbf{v}_j}}$ are independent in $\tilde{V}$. Given $\mathbf{w}_j \in H_{\mathbf{v}_j}$ satisfying $\sum \mathbf{w}_j = \mathbf{0}$, it follows that $\sum \tilde{\mathbf{w}}_j = \tilde{\mathbf{0}}$. Thus $\tilde{\mathbf{w}}_j = \tilde{\mathbf{0}}$ for each j, i.e. $\mathbf{w}_j \in \ker p(T)$. By Lemma 7.5 we may write $\mathbf{w}_j = p(T)\mathbf{u}_j$ where $\mathbf{u}_j \in H_{\mathbf{v}_j} - \ker p(T)$. The fact that $\sum \mathbf{w}_j = \mathbf{0}$ may then be recast as $\sum \tilde{\mathbf{u}}_j = \tilde{\mathbf{0}}$. Independence of the subspaces $\widetilde{H_{\mathbf{v}_j}}$ then gives $\mathbf{u}_j = \ker p(T)$ for each j, too; hence $\mathbf{w}_j = p(T)\mathbf{u}_j = \mathbf{0}$. This proves that the subspaces $H_{\mathbf{v}_j}$ are independent in V. □

Lemma 7.7 *Let $H \subset V$ be any invariant subspace. Then there exist vectors $\mathbf{v}_j \in V$, $1 \leq j \leq k$ such that*

$$\ker p(T) = (H \cap \ker p(T)) \oplus H_{\mathbf{v}_1} \oplus \cdots \oplus H_{\mathbf{v}_k}.$$

Proof. If $\ker p(T) \subset H$, there is nothing to prove. So suppose we have a non-zero $\mathbf{v} \in \ker p(T) - H$. We claim that $H_\mathbf{v} \cap H = \{\mathbf{0}\}$ is trivial. Hence $H_\mathbf{v}$ and H are independent subspaces, and either $H \oplus H_\mathbf{v} = \ker p(T)$, or we can repeat this process with $H \oplus H_\mathbf{v}$ in place of H. Since $\dim(H \oplus H_\mathbf{v}) \cap \ker p(T) = (\dim H \oplus \ker p(T)) + \deg p$ increases with every repetition, we will eventually arrive at the desired sequence of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k \in \ker p(T)$.

To establish the claim, let $W := H \cap H_\mathbf{v}$. As the intersection of invariant subspaces W is itself an invariant subspace of $H_\mathbf{v}$. Let S denote the operator on $H_\mathbf{v}/W$ induced by the restriction $T|_{H_\mathbf{v}}$. The characteristic polynomial of $T|_{H_\mathbf{v}}$ is p, and the characteristic polynomial q of S must therefore divide p. Since p is irreducible, it follows that $q = 1$ or $q = p$. In the first case, we infer that $\dim H_\mathbf{v}/W = 0$ which means that $H_\mathbf{v} = W \subset H$. This contradicts the fact that $\mathbf{v} \in H_{\mathbf{v}} - W$. In the second case, we infer $\dim H_{\mathbf{v}}/W = \deg p = \dim H_{\mathbf{v}}$ and therefore $W = H_{\mathbf{v}} \cap H$ is trivial as claimed. □

The next lemma really doesn't have anything to do with operators at all, but we state it only for the current context.

Lemma 7.8 *Suppose that $H \subset V$ is a subspace and $\tilde{H} = \tilde{V}$ and that $\ker p(T) = (H \cap \ker p(T)) \oplus W$ for some $W \subset \ker p(T)$. Then $V = H \oplus W$.*

Proof. First note that since $W \subset \ker p(T)$, we have that if $H \cap W = (\ker p(T) \cap H) \cap W$ is trivial. So H and W are independent.

Given $\mathbf{v} \in V$, the hypothesis that $\tilde{\mathbf{v}} \in \tilde{H}$ means that $\mathbf{v} = \mathbf{v}' + \mathbf{u}$ where $\mathbf{v}' \in H$ and $\mathbf{u} \in \ker p(T)$. The hypothesis that W complements $H \cap \ker p(T)$ in $\ker p(T)$ allows us to further decompose $\mathbf{u} = \mathbf{v}'' + \mathbf{w}$ where $\mathbf{v}'' \in H$ and $\mathbf{w} \in W$. Hence $\mathbf{v} = (\mathbf{v}' + \mathbf{v}'') + \mathbf{w} \in H + W$. So $V = H \oplus W$. □

The next (and last) lemma will help us prove that different cyclic decompositions of V are equivalent.

Lemma 7.9 *Let $\mathbf{v}, \mathbf{w} \in V$ be vectors with the same order s. Then linear transformation $\phi : H_{\mathbf{v}} \to H_{\mathbf{w}}$ determined by $\phi(T^j\mathbf{v}) = T^j\mathbf{w}$, $0 \le j < s \deg p$ is an isomorphism satisfying $\phi \circ T = T \circ \phi$.*

Proof. The requirement that $\phi(T^j\mathbf{v}) = T^j\mathbf{w}$ for $0 \le j < s \deg p$ amounts to saying that ϕ carries the usual basis for $H_{\mathbf{v}}$ to the corresponding basis for $H_{\mathbf{w}}$. There is exactly one linear transformation accomplishing this task, and it must be an automorphism. So ϕ is a well-defined isomorphism from $H_{\mathbf{v}}$ to $H_{\mathbf{w}}$. We have by construction that $\phi \circ T(T^j\mathbf{v}) = T \circ \phi(T^j\mathbf{v})$ for $0 \le j \le s-2$. The same holds for $j = s-1$ because $p(T)^s\mathbf{v} = p(T)^s\mathbf{w} = \mathbf{0}$—i.e. $p_v = p_w$ means that the final basis vectors for $H_{\mathbf{v}}$ and $H_{\mathbf{w}}$ can be written as linear combinations of their predecessors in exactly the same way. At any rate, since $\phi \circ T = T \circ \phi$ for all vectors in a basis for $H_{\mathbf{v}}$, the same holds for all vectors in $H_{\mathbf{v}}$. □

Proof of Theorem 7.4. As mentioned above, we will work by induction on the order s of T. When $s = 1$, we have $\ker p(T) = V$. Hence the existence of the cyclic decomposition $V = \bigoplus_{j=1}^{k} H_{\mathbf{v}_j}$ follows immediately from Lemma 7.7, and the characteristic polynomials $p_{\mathbf{v}_j}$ for $T|_{H_{\mathbf{v}_j}}$ are all equal to p. Hence the number k of factors in the decomposition is $\dim V/\deg p$ which is in particular independent of the choice of $\mathbf{v}_j$.

If $V = \bigoplus_{j=1}^{k} H_{\mathbf{w}_j}$ is another cyclic decomposition of V, Lemma 7.9 gives us isomorphisms $\phi_j : H_{\mathbf{v}_j} \to H_{\mathbf{w}_j}$ between corresponding factors. We assemble these into an isomorphism $\phi : V \to V$ by using that any $\mathbf{v} \in V$ decomposes uniquely as $\mathbf{v} = \mathbf{u}_1 + \cdots + \mathbf{u}_k$, $\mathbf{u}_j \in H_{\mathbf{v}_j}$ and setting

$$\phi(\mathbf{u}_1 + \cdots + \mathbf{u}_k) = \phi_1(\mathbf{u}_1) + \cdots + \phi_k(\mathbf{u}_k).$$

This completes the case $s = 1$.

Assuming inductively that the theorem is true when the order of T is $s-1$, we consider the case when the order of T is s. By our inductive hypothesis, we $\tilde{\mathbf{v}}_1, \ldots, \tilde{\mathbf{v}}_i \in V$ such that

$\tilde{V} = V/\ker p(T) = \bigoplus_{j=1}^{i} \tilde{H}_{\mathbf{v}_j}$. We can apply Lemmas 7.7 and 7.8 to $H = H_{\mathbf{v}_1} + \cdots + H_{\mathbf{v}_k}$ to get vectors $\mathbf{v}_{i+1}, \ldots, \mathbf{v}_k$ such that

$$V = H \oplus V_{i+1} \oplus \cdots \oplus V_k.$$

By Lemma 7.6, the subspaces $H_{\mathbf{v}_j}$ remain independent in V, so we have in fact that

$$V = \oplus_{j=1}^{k} H_{\mathbf{v}_j}.$$

This completes the proof of existence of a cyclic decomposition.

Now if $V = \oplus_{j=1}^{k} H_{\mathbf{w}_j}$ is another cyclic decomposition, we order the factors in both decompositions from largest to smallest. Then if i' is the largest integer such that order of $\mathbf{w}_{i'}$ exceeds 1, then we get a cyclic decomposition

$$\tilde{V} = \oplus_{j=1}^{i'} \tilde{H}_{\mathbf{w}_j}$$

of the quotient space, independence of the factors proceeding from Lemma 7.6. Our inductive hypothesis then implies that $i' = i$ and that the order of $\mathbf{v}_j$ equals that of $\mathbf{w}_j$ for each $1 \leq j \leq i$. In fact the orders remain the same for $i < j \leq k$, too since $\mathbf{v}_j$ and $\mathbf{w}_j$ both have order 1 for all $j > i$. So finally we can repeat the argument for equivalence from the case $s = 1$ to construct an equivalence isomorphism $\phi : V \to V$ carrying $H_{\mathbf{v}_j}$ to $H_{\mathbf{w}_j}$ for all j. $\square$

8 Jordan Canonical Form

Let us apply the cyclic decomposition theorem to the case of a primary subspace $H = H_{x-\lambda}$ assoicated to a polynomial $p(x) = x - \lambda$ of degree one. Replacing T with $S := T - \lambda\mathrm{id}$, we have that $S|_H$ is nilpotent of some order r. We may therefore decompose H into a direct sum of finitely many S-cyclic subspaces $W = \operatorname{span} \mathcal{B}$, where $\mathcal{B} = \{\mathbf{v}, S\mathbf{v}, \ldots, S^{k-1}\mathbf{v}\}$ for some $\mathbf{v} \in V$, and $S^k\mathbf{v} = \mathbf{0}$. In matrix terms, this becomes

$$[S|_W]_{\mathcal{B}\mathcal{B}} = \begin{bmatrix} 0 & 0 & 0 & \ldots & 0 & 0 \\ 1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & 1 & 0 & \ldots & 0 & 0 \\ 0 & 0 & 1 & \ldots & 0 & 0 \\ & & & \vdots & & \\ 0 & 0 & 0 & \ldots & 1 & 0 \end{bmatrix}.$$

The subspaces W will also be invariant (in fact cyclic) by $T = S + \lambda\mathrm{id}$, and we see that

$$[T|_W]_{\mathcal{B}\mathcal{B}} = \begin{bmatrix} \lambda & 0 & 0 & \ldots & 0 & 0 \\ 1 & \lambda & 0 & \ldots & 0 & 0 \\ 0 & 1 & \lambda & \ldots & 0 & 0 \\ 0 & 0 & 1 & \ldots & 0 & 0 \\ & & & \vdots & & \\ 0 & 0 & 0 & \ldots & 1 & \lambda \end{bmatrix}.$$

Note that if we reverse the order of the vectors in $\mathcal{B}$, then the matrix becomes upper triangular, with '1's above rather than below the main diagonal. This is the more conventional way to present the matrix, but for the sake of internal consistency, we stick with the lower triangular version here.

In the case $\mathbf{F} = \mathbf{C}$, when all irreducible factors of p_{min} have degree one, we can apply these observations to *every* primary subspace and conclude as follows.

Theorem 8.1 (Jordan Canonical Form) *Let $T : V \to V$ be a linear operator on a complex vector space V. Then there is a basis $\mathcal{B} \subset V$ such that $[T]_{\mathcal{B}\mathcal{B}}$ has block diagonal form*

$$\begin{bmatrix} A_1 & & & \\ & A_2 & & \\ & & \ddots & \\ & & & A_k \end{bmatrix}$$

where for each i, there exists an eigenvalue λ_i of T such that

$$A_i = \begin{bmatrix} \lambda_i & 0 & 0 & \dots & 0 & 0 \\ 1 & \lambda_i & 0 & \dots & 0 & 0 \\ 0 & 1 & \lambda_i & \dots & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 \\ & & & \vdots & & \\ 0 & 0 & 0 & \dots & 1 & \lambda_i \end{bmatrix}.$$

Any two such bases for V give rise to similar such matrices, having the same blocks, albeit possibly in different order.