Chapter 5

Eigenvalues and Eigenvectors

In Chapter 3 we studied linear maps from one vector space to another vector space. Now we begin our investigation of linear maps from a vector space to itself. Their study constitutes the deepest and most important part of linear algebra. Most of the key results in this area do not hold for infinite-dimensional vector spaces, so we work only on finite-dimensional vector spaces. To avoid trivialities we also want to eliminate the vector space $\{0\}$ from consideration. Thus we make the following assumption:

Recall that $\mathbf{F}$ denotes $\mathbf{R}$ or $\mathbf{C}$. Let's agree that for the rest of the book V will denote a finite-dimensional, nonzero vector space over $\mathbf{F}$.

Invariant Subspaces

In this chapter we develop the tools that will help us understand the structure of operators. Recall that an operator is a linear map from a vector space to itself. Recall also that we denote the set of operators on V by $\mathcal{L}(V)$; in other words, $\mathcal{L}(V) = \mathcal{L}(V, V)$.

Let's see how we might better understand what an operator looks like. Suppose $T \in \mathcal{L}(V)$. If we have a direct sum decomposition

5.1 $$V = U_1 \oplus \cdots \oplus U_m,$$

where each U_j is a proper subspace of V, then to understand the behavior of T, we need only understand the behavior of each $T|_{U_j}$; here $T|_{U_j}$ denotes the restriction of T to the smaller domain U_j. Dealing with $T|_{U_j}$ should be easier than dealing with T because U_j is a smaller vector space than V. However, if we intend to apply tools useful in the study of operators (such as taking powers), then we have a problem: $T|_{U_j}$ may not map U_j into itself; in other words, $T|_{U_j}$ may not be an operator on U_j. Thus we are led to consider only decompositions of the form 5.1 where T maps each U_j into itself.

The notion of a subspace that gets mapped into itself is sufficiently important to deserve a name. Thus, for $T \in \mathcal{L}(V)$ and U a subspace of V, we say that U is ***invariant*** under T if $u \in U$ implies $Tu \in U$. In other words, U is invariant under T if $T|_U$ is an operator on U. For example, if T is the operator of differentiation on $\mathcal{P}_7(\mathbf{R})$, then $\mathcal{P}_4(\mathbf{R})$ (which is a subspace of $\mathcal{P}_7(\mathbf{R})$) is invariant under T because the derivative of any polynomial of degree at most 4 is also a polynomial with degree at most 4.

The most famous unsolved problem in functional analysis is called the invariant subspace problem. It deals with invariant subspaces of operators on infinite-dimensional vector spaces.

Let's look at some easy examples of invariant subspaces. Suppose $T \in \mathcal{L}(V)$. Clearly $\{0\}$ is invariant under T. Also, the whole space V is obviously invariant under T. Must T have any invariant subspaces other than $\{0\}$ and V? Later we will see that this question has an affirmative answer for operators on complex vector spaces with dimension greater than 1 and also for operators on real vector spaces with dimension greater than 2.

If $T \in \mathcal{L}(V)$, then $\operatorname{null} T$ is invariant under T (proof: if $u \in \operatorname{null} T$, then $Tu = 0$, and hence $Tu \in \operatorname{null} T$). Also, $\operatorname{range} T$ is invariant under T (proof: if $u \in \operatorname{range} T$, then Tu is also in $\operatorname{range} T$, by the definition of range). Although $\operatorname{null} T$ and $\operatorname{range} T$ are invariant under T, they do not necessarily provide easy answers to the question about the existence of invariant subspaces other than $\{0\}$ and V because null T may equal $\{0\}$ and range T may equal V (this happens when T is invertible).

We will return later to a deeper study of invariant subspaces. Now we turn to an investigation of the simplest possible nontrivial invariant subspaces—invariant subspaces with dimension 1.

How does an operator behave on an invariant subspace of dimension 1? Subspaces of V of dimension 1 are easy to describe. Take any nonzero vector $u \in V$ and let U equal the set of all scalar multiples of u:

5.2 $$U = \{au : a \in \mathbf{F}\}.$$

Then U is a one-dimensional subspace of V, and every one-dimensional subspace of V is of this form. If $u \in V$ and the subspace U defined by 5.2 is invariant under $T \in \mathcal{L}(V)$, then Tu must be in U, and hence there must be a scalar $\lambda \in \mathbf{F}$ such that $Tu = \lambda u$. Conversely, if u is a nonzero vector in V such that $Tu = \lambda u$ for some $\lambda \in \mathbf{F}$, then the subspace U defined by 5.2 is a one-dimensional subspace of V invariant under T.

*These subspaces are loosely connected to the subject of Herbert Marcuse's well-known book **One-Dimensional Man**.*

The equation

5.3 $$Tu = \lambda u,$$

which we have just seen is intimately connected with one-dimensional invariant subspaces, is important enough that the vectors u and scalars λ satisfying it are given special names. Specifically, a scalar $\lambda \in \mathbf{F}$ is called an ***eigenvalue*** of $T \in \mathcal{L}(V)$ if there exists a nonzero vector $u \in V$ such that $Tu = \lambda u$. We must require u to be nonzero because with $u = 0$ every scalar $\lambda \in \mathbf{F}$ satisfies 5.3. The comments above show that T has a one-dimensional invariant subspace if and only if T has an eigenvalue.

*The regrettable word **eigenvalue** is half-German, half-English. The German adjective **eigen** means own in the sense of characterizing some intrinsic property. Some mathematicians use the term **characteristic value** instead of eigenvalue.*

The equation $Tu = \lambda u$ is equivalent to $(T - \lambda I)u = 0$, so λ is an eigenvalue of T if and only if $T - \lambda I$ is not injective. By 3.21, λ is an eigenvalue of T if and only if $T - \lambda I$ is not invertible, and this happens if and only if $T - \lambda I$ is not surjective.

Suppose $T \in \mathcal{L}(V)$ and $\lambda \in \mathbf{F}$ is an eigenvalue of T. A vector $u \in V$ is called an ***eigenvector*** of T (corresponding to λ) if $Tu = \lambda u$. Because 5.3 is equivalent to $(T - \lambda I)u = 0$, we see that the set of eigenvectors of T corresponding to λ equals $\operatorname{null}(T - \lambda I)$. In particular, the set of eigenvectors of T corresponding to λ is a subspace of V.

Some texts define eigenvectors as we have, except that 0 *is declared not to be an eigenvector. With the definition used here, the set of eigenvectors corresponding to a fixed eigenvalue is a subspace.*

Let's look at some examples of eigenvalues and eigenvectors. If $a \in \mathbf{F}$, then aI has only one eigenvalue, namely, a, and every vector is an eigenvector for this eigenvalue.

For a more complicated example, consider the operator $T \in \mathcal{L}(\mathbf{F}^2)$ defined by

5.4 $$T(w, z) = (-z, w).$$

If $\mathbf{F} = \mathbf{R}$, then this operator has a nice geometric interpretation: T is just a counterclockwise rotation by 90° about the origin in $\mathbf{R}^2$. An operator has an eigenvalue if and only if there exists a nonzero vector in its domain that gets sent by the operator to a scalar multiple of itself. The rotation of a nonzero vector in $\mathbf{R}^2$ obviously never equals a scalar multiple of itself. Conclusion: if $\mathbf{F} = \mathbf{R}$, the operator T defined by 5.4 has no eigenvalues. However, if $\mathbf{F} = \mathbf{C}$, the story changes. To find eigenvalues of T, we must find the scalars λ such that

$$T(w, z) = \lambda(w, z)$$

has some solution other than $w = z = 0$. For T defined by 5.4, the equation above is equivalent to the simultaneous equations

5.5 $$-z = \lambda w, \quad w = \lambda z.$$

Substituting the value for w given by the second equation into the first equation gives

$$-z = \lambda^2 z.$$

Now z cannot equal 0 (otherwise 5.5 implies that $w = 0$; we are looking for solutions to 5.5 where (w, z) is not the 0 vector), so the equation above leads to the equation

$$-1 = \lambda^2.$$

The solutions to this equation are $\lambda = i$ or $\lambda = -i$. You should be able to verify easily that i and $-i$ are eigenvalues of T. Indeed, the eigenvectors corresponding to the eigenvalue i are the vectors of the form $(w, -wi)$, with $w \in \mathbf{C}$, and the eigenvectors corresponding to the eigenvalue $-i$ are the vectors of the form (w, wi), with $w \in \mathbf{C}$.

Now we show that nonzero eigenvectors corresponding to distinct eigenvalues are linearly independent.

5.6 Theorem: *Let $T \in \mathcal{L}(V)$. Suppose $\lambda_1, \dots, \lambda_m$ are distinct eigenvalues of T and $v_1, \dots, v_m$ are corresponding nonzero eigenvectors. Then $(v_1, \dots, v_m)$ is linearly independent.*

PROOF: Suppose $(v_1, \dots, v_m)$ is linearly dependent. Let k be the smallest positive integer such that

5.7 $$v_k \in \operatorname{span}(v_1, \dots, v_{k-1});$$

the existence of k with this property follows from the linear dependence lemma (2.4). Thus there exist $a_1, \dots, a_{k-1} \in \mathbf{F}$ such that

5.8 $$v_k = a_1 v_1 + \cdots + a_{k-1} v_{k-1}.$$

Apply T to both sides of this equation, getting

$$\lambda_k v_k = a_1 \lambda_1 v_1 + \cdots + a_{k-1} \lambda_{k-1} v_{k-1}.$$

Multiply both sides of 5.8 by λ_k and then subtract the equation above, getting

$$0 = a_1(\lambda_k - \lambda_1) v_1 + \cdots + a_{k-1}(\lambda_k - \lambda_{k-1}) v_{k-1}.$$

Because we chose k to be the smallest positive integer satisfying 5.7, $(v_1, \dots, v_{k-1})$ is linearly independent. Thus the equation above implies that all the a's are 0 (recall that λ_k is not equal to any of $\lambda_1, \dots, \lambda_{k-1}$). However, this means that v_k equals 0 (see 5.8), contradicting our hypothesis that all the v's are nonzero. Therefore our assumption that $(v_1, \dots, v_m)$ is linearly dependent must have been false. ■

The corollary below states that an operator cannot have more distinct eigenvalues than the dimension of the vector space on which it acts.

5.9 Corollary: *Each operator on V has at most $\dim V$ distinct eigenvalues.*

PROOF: Let $T \in \mathcal{L}(V)$. Suppose that $\lambda_1, \dots, \lambda_m$ are distinct eigenvalues of T. Let $v_1, \dots, v_m$ be corresponding nonzero eigenvectors. The last theorem implies that $(v_1, \dots, v_m)$ is linearly independent. Thus $m \le \dim V$ (see 2.6), as desired. ■

Polynomials Applied to Operators

The main reason that a richer theory exists for operators (which map a vector space into itself) than for linear maps is that operators can be raised to powers. In this section we define that notion and the key concept of applying a polynomial to an operator.

If $T \in \mathcal{L}(V)$, then TT makes sense and is also in $\mathcal{L}(V)$. We usually write T^2 instead of TT. More generally, if m is a positive integer, then T^m is defined by

$$T^m = \underbrace{T \dots T}_{m \text{ times}}.$$

For convenience we define T^0 to be the identity operator I on V.

Recall from Chapter 3 that if T is an invertible operator, then the inverse of T is denoted by T^{-1}. If m is a positive integer, then we define T^{-m} to be $(T^{-1})^m$.

You should verify that if T is an operator, then

$$T^m T^n = T^{m+n} \quad \text{and} \quad (T^m)^n = T^{mn},$$

where m and n are allowed to be arbitrary integers if T is invertible and nonnegative integers if T is not invertible.

If $T \in \mathcal{L}(V)$ and $p \in \mathcal{P}(\mathbf{F})$ is a polynomial given by

$$p(z) = a_0 + a_1 z + a_2 z^2 + \cdots + a_m z^m$$

for $z \in \mathbf{F}$, then $p(T)$ is the operator defined by

$$p(T) = a_0 I + a_1 T + a_2 T^2 + \cdots + a_m T^m.$$

For example, if p is the polynomial defined by $p(z) = z^2$ for $z \in \mathbf{F}$, then $p(T) = T^2$. This is a new use of the symbol p because we are applying it to operators, not just elements of $\mathbf{F}$. If we fix an operator $T \in \mathcal{L}(V)$, then the function from $\mathcal{P}(\mathbf{F})$ to $\mathcal{L}(V)$ given by $p \mapsto p(T)$ is linear, as you should verify.

If p and q are polynomials with coefficients in $\mathbf{F}$, then pq is the polynomial defined by

$$(pq)(z) = p(z)q(z)$$

for $z \in \mathbf{F}$. You should verify that we have the following nice multiplicative property: if $T \in \mathcal{L}(V)$, then

$$(pq)(T) = p(T)q(T)$$

for all polynomials p and q with coefficients in $\mathbf{F}$. Note that any two polynomials in T commute, meaning that $p(T)q(T) = q(T)p(T)$, because

$$p(T)q(T) = (pq)(T) = (qp)(T) = q(T)p(T).$$

Upper-Triangular Matrices

Now we come to one of the central results about operators on complex vector spaces.

5.10 Theorem: *Every operator on a finite-dimensional, nonzero, complex vector space has an eigenvalue.*

Compare the simple proof of this theorem given here with the standard proof using determinants. With the standard proof, first the difficult concept of determinants must be defined, then an operator with 0 determinant must be shown to be not invertible, then the characteristic polynomial needs to be defined, and by the time the proof of this theorem is reached, no insight remains about why it is true.

PROOF: Suppose V is a complex vector space with dimension $n > 0$ and $T \in \mathcal{L}(V)$. Choose $v \in V$ with $v \neq 0$. Then

$$(v, Tv, T^2v, \dots, T^nv)$$

cannot be linearly independent because V has dimension n and we have $n + 1$ vectors. Thus there exist complex numbers $a_0, \dots, a_n$, not all 0, such that

$$0 = a_0v + a_1Tv + \cdots + a_nT^nv.$$

Let m be the largest index such that $a_m \neq 0$. Because $v \neq 0$, the coefficients $a_1, \dots, a_m$ cannot all be 0, so $0 < m \leq n$. Make the a's the coefficients of a polynomial, which can be written in factored form (see 4.8) as

$$a_0 + a_1z + \cdots + a_nz^n = c(z - \lambda_1)\dots(z - \lambda_m),$$

where c is a nonzero complex number, each $\lambda_j \in \mathbf{C}$, and the equation holds for all $z \in \mathbf{C}$. We then have

$$\begin{aligned} 0 &= a_0v + a_1Tv + \cdots + a_nT^nv \\ &= (a_0I + a_1T + \cdots + a_nT^n)v \\ &= c(T - \lambda_1I)\dots(T - \lambda_mI)v, \end{aligned}$$

which means that $T - \lambda_jI$ is not injective for at least one j. In other words, T has an eigenvalue. ∎

Recall that in Chapter 3 we discussed the matrix of a linear map from one vector space to another vector space. This matrix depended on a choice of a basis for each of the two vector spaces. Now that we are studying operators, which map a vector space to itself, we need only one basis. In addition, now our matrices will be square arrays, rather than the more general rectangular arrays that we considered earlier. Specifically, let $T \in \mathcal{L}(V)$. Suppose $(v_1, \dots, v_n)$ is a basis of V. For each $k = 1, \dots, n$, we can write

$$Tv_k = a_{1,k}v_1 + \cdots + a_{n,k}v_n,$$

The k^{th} column of the matrix is formed from the coefficients used to write Tv_k as a linear combination of the v's.

where $a_{j,k} \in \mathbf{F}$ for $j = 1, \dots, n$. The n-by-n matrix

5.11
$$\begin{bmatrix} a_{1,1} & \dots & a_{1,n} \\ \vdots & & \vdots \\ a_{n,1} & \dots & a_{n,n} \end{bmatrix}$$

is called the ***matrix*** of T with respect to the basis $(v_1, \dots, v_n)$; we denote it by $\mathcal{M}(T, (v_1, \dots, v_n))$ or just by $\mathcal{M}(T)$ if the basis $(v_1, \dots, v_n)$ is clear from the context (for example, if only one basis is in sight).

If T is an operator on $\mathbf{F}^n$ and no basis is specified, you should assume that the basis in question is the standard one (where the j^{th} basis vector is 1 in the j^{th} slot and 0 in all the other slots). You can then think of the j^{th} column of $\mathcal{M}(T)$ as T applied to the j^{th} basis vector.

A central goal of linear algebra is to show that given an operator $T \in \mathcal{L}(V)$, there exists a basis of V with respect to which T has a reasonably simple matrix. To make this vague formulation ("reasonably simple" is not precise language) a bit more concrete, we might try to make $\mathcal{M}(T)$ have many 0's.

If V is a complex vector space, then we already know enough to show that there is a basis of V with respect to which the matrix of T has 0's everywhere in the first column, except possibly the first entry. In other words, there is a basis of V with respect to which the matrix of T looks like

We often use $*$ to denote matrix entries that we do not know about or that are irrelevant to the questions being discussed.

$$\begin{bmatrix} \lambda & & \\ 0 & * & \\ \vdots & & \\ 0 & & \end{bmatrix};$$

here the $*$ denotes the entries in all the columns other than the first column. To prove this, let λ be an eigenvalue of T (one exists by 5.10) and let v be a corresponding nonzero eigenvector. Extend (v) to a basis of V. Then the matrix of T with respect to this basis has the form above. Soon we will see that we can choose a basis of V with respect to which the matrix of T has even more 0's.

The ***diagonal*** of a square matrix consists of the entries along the straight line from the upper left corner to the bottom right corner. For example, the diagonal of the matrix 5.11 consists of the entries $a_{1,1}, a_{2,2}, \dots, a_{n,n}$.

A matrix is called ***upper triangular*** if all the entries below the diagonal equal 0. For example, the 4-by-4 matrix

$$\begin{bmatrix} 6 & 2 & 7 & 5 \\ 0 & 6 & 1 & 3 \\ 0 & 0 & 7 & 9 \\ 0 & 0 & 0 & 8 \end{bmatrix}$$

is upper triangular. Typically we represent an upper-triangular matrix in the form

$$\begin{bmatrix} \lambda_1 & & * \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix};$$

the 0 in the matrix above indicates that all entries below the diagonal in this n-by-n matrix equal 0. Upper-triangular matrices can be considered reasonably simple—for n large, an n-by-n upper-triangular matrix has almost half its entries equal to 0.

The following proposition demonstrates a useful connection between upper-triangular matrices and invariant subspaces.

5.12 Proposition: *Suppose $T \in \mathcal{L}(V)$ and $(v_1, \dots, v_n)$ is a basis of V. Then the following are equivalent:*

(a) *the matrix of T with respect to $(v_1, \dots, v_n)$ is upper triangular;*

(b) $Tv_k \in \operatorname{span}(v_1, \dots, v_k)$ *for each $k = 1, \dots, n$;*

(c) $\operatorname{span}(v_1, \dots, v_k)$ *is invariant under T for each $k = 1, \dots, n$.*

Proof: The equivalence of (a) and (b) follows easily from the definitions and a moment's thought. Obviously (c) implies (b). Thus to complete the proof, we need only prove that (b) implies (c). So suppose that (b) holds. Fix $k \in \{1, \dots, n\}$. From (b), we know that

$$
\begin{aligned}
Tv_1 &\in \operatorname{span}(v_1) \subset \operatorname{span}(v_1, \ldots, v_k);\\
Tv_2 &\in \operatorname{span}(v_1, v_2) \subset \operatorname{span}(v_1, \ldots, v_k);\\
&\vdots\\
Tv_k &\in \operatorname{span}(v_1, \ldots, v_k).
\end{aligned}
$$

Thus if v is a linear combination of $(v_1, \ldots, v_k)$, then

$$Tv \in \operatorname{span}(v_1, \ldots, v_k).$$

In other words, $\operatorname{span}(v_1, \ldots, v_k)$ is invariant under T, completing the proof. ■

Now we can show that for each operator on a complex vector space, there is a basis of the vector space with respect to which the matrix of the operator has only 0's below the diagonal. In Chapter 8 we will improve even this result.

This theorem does not hold on real vector spaces because the first vector in a basis with respect to which an operator has an upper-triangular matrix must be an eigenvector of the operator. Thus if an operator on a real vector space has no eigenvalues (we have seen an example on $\mathbf{R}^2$*), then there is no basis with respect to which the operator has an upper-triangular matrix.*

5.13 Theorem: *Suppose V is a complex vector space and $T \in \mathcal{L}(V)$. Then T has an upper-triangular matrix with respect to some basis of V.*

PROOF: We will use induction on the dimension of V. Clearly the desired result holds if $\dim V = 1$.

Suppose now that $\dim V > 1$ and the desired result holds for all complex vector spaces whose dimension is less than the dimension of V. Let λ be any eigenvalue of T (5.10 guarantees that T has an eigenvalue). Let

$$U = \operatorname{range}(T - \lambda I).$$

Because $T - \lambda I$ is not surjective (see 3.21), $\dim U < \dim V$. Furthermore, U is invariant under T. To prove this, suppose $u \in U$. Then

$$Tu = (T - \lambda I)u + \lambda u.$$

Obviously $(T - \lambda I)u \in U$ (from the definition of U) and $\lambda u \in U$. Thus the equation above shows that $Tu \in U$. Hence U is invariant under T, as claimed.

Thus $T|_U$ is an operator on U. By our induction hypothesis, there is a basis $(u_1, \ldots, u_m)$ of U with respect to which $T|_U$ has an upper-triangular matrix. Thus for each j we have (using 5.12)

5.14 $$Tu_j = (T|_U)(u_j) \in \operatorname{span}(u_1, \ldots, u_j).$$

Extend $(u_1, \dots, u_m)$ to a basis $(u_1, \dots, u_m, v_1, \dots, v_n)$ of V. For each k, we have

$$Tv_k = (T - \lambda I)v_k + \lambda v_k.$$

The definition of U shows that $(T - \lambda I)v_k \in U = \operatorname{span}(u_1, \dots, u_m)$. Thus the equation above shows that

5.15 $$Tv_k \in \operatorname{span}(u_1, \dots, u_m, v_1, \dots, v_k).$$

From 5.14 and 5.15, we conclude (using 5.12) that T has an upper-triangular matrix with respect to the basis $(u_1, \dots, u_m, v_1, \dots, v_n)$. ■

How does one determine from looking at the matrix of an operator whether the operator is invertible? If we are fortunate enough to have a basis with respect to which the matrix of the operator is upper triangular, then this problem becomes easy, as the following proposition shows.

5.16 Proposition: *Suppose $T \in \mathcal{L}(V)$ has an upper-triangular matrix with respect to some basis of V. Then T is invertible if and only if all the entries on the diagonal of that upper-triangular matrix are nonzero.*

PROOF: Suppose $(v_1, \dots, v_n)$ is a basis of V with respect to which T has an upper-triangular matrix

5.17 $$\mathcal{M}(T, (v_1, \dots, v_n)) = \begin{bmatrix} \lambda_1 & & & * \\ & \lambda_2 & & \\ & & \ddots & \\ 0 & & & \lambda_n \end{bmatrix}.$$

We need to prove that T is not invertible if and only if one of the λ_k's equals 0.

First we will prove that if one of the λ_k's equals 0, then T is not invertible. If $\lambda_1 = 0$, then $Tv_1 = 0$ (from 5.17) and hence T is not invertible, as desired. So suppose that $1 < k \le n$ and $\lambda_k = 0$. Then, as can be seen from 5.17, T maps each of the vectors $v_1, \dots, v_{k-1}$ into $\operatorname{span}(v_1, \dots, v_{k-1})$. Because $\lambda_k = 0$, the matrix representation 5.17 also implies that $Tv_k \in \operatorname{span}(v_1, \dots, v_{k-1})$. Thus we can define a linear map

$$S\colon \operatorname{span}(v_1, \dots, v_k) \to \operatorname{span}(v_1, \dots, v_{k-1})$$

by $Sv = Tv$ for $v \in \operatorname{span}(v_1, \dots, v_k)$. In other words, S is just T restricted to $\operatorname{span}(v_1, \dots, v_k)$.

Note that $\operatorname{span}(v_1, \dots, v_k)$ has dimension k and $\operatorname{span}(v_1, \dots, v_{k-1})$ has dimension $k - 1$ (because $(v_1, \dots, v_n)$ is linearly independent). Because $\operatorname{span}(v_1, \dots, v_k)$ has a larger dimension than $\operatorname{span}(v_1, \dots, v_{k-1})$, no linear map from $\operatorname{span}(v_1, \dots, v_k)$ to $\operatorname{span}(v_1, \dots, v_{k-1})$ is injective (see 3.5). Thus there exists a nonzero vector $v \in \operatorname{span}(v_1, \dots, v_k)$ such that $Sv = 0$. Hence $Tv = 0$, and thus T is not invertible, as desired.

To prove the other direction, now suppose that T is not invertible. Thus T is not injective (see 3.21), and hence there exists a nonzero vector $v \in V$ such that $Tv = 0$. Because $(v_1, \dots, v_n)$ is a basis of V, we can write

$$v = a_1 v_1 + \cdots + a_k v_k,$$

where $a_1, \dots, a_k \in \mathbf{F}$ and $a_k \neq 0$ (represent v as a linear combination of $(v_1, \dots, v_n)$ and then choose k to be the largest index with a nonzero coefficient). Thus

$$\begin{aligned}
0 &= Tv \\
0 &= T(a_1 v_1 + \cdots + a_k v_k) \\
&= (a_1 T v_1 + \cdots + a_{k-1} T v_{k-1}) + a_k T v_k.
\end{aligned}$$

The last term in parentheses is in $\operatorname{span}(v_1, \dots, v_{k-1})$ (because of the upper-triangular form of 5.17). Thus the last equation shows that $a_k T v_k \in \operatorname{span}(v_1, \dots, v_{k-1})$. Multiplying by $1/a_k$, which is allowed because $a_k \neq 0$, we conclude that $T v_k \in \operatorname{span}(v_1, \dots, v_{k-1})$. Thus when $T v_k$ is written as a linear combination of the basis $(v_1, \dots, v_n)$, the coefficient of v_k will be 0. In other words, λ_k in 5.17 must be 0, completing the proof. ■

Powerful numeric techniques exist for finding good approximations to the eigenvalues of an operator from its matrix.

Unfortunately no method exists for exactly computing the eigenvalues of a typical operator from its matrix (with respect to an arbitrary basis). However, if we are fortunate enough to find a basis with respect to which the matrix of the operator is upper triangular, then the problem of computing the eigenvalues becomes trivial, as the following proposition shows.

5.18 Proposition: *Suppose $T \in \mathcal{L}(V)$ has an upper-triangular matrix with respect to some basis of V. Then the eigenvalues of T consist precisely of the entries on the diagonal of that upper-triangular matrix.*

PROOF: Suppose $(v_1, \ldots, v_n)$ is a basis of V with respect to which T has an upper-triangular matrix

$$\mathcal{M}(T, (v_1, \ldots, v_n)) = \begin{bmatrix} \lambda_1 & & & * \\ & \lambda_2 & & \\ & & \ddots & \\ 0 & & & \lambda_n \end{bmatrix}.$$

Let $\lambda \in \mathbf{F}$. Then

$$\mathcal{M}(T - \lambda I, (v_1, \ldots, v_n)) = \begin{bmatrix} \lambda_1 - \lambda & & & * \\ & \lambda_2 - \lambda & & \\ & & \ddots & \\ 0 & & & \lambda_n - \lambda \end{bmatrix}.$$

Hence $T - \lambda I$ is not invertible if and only if λ equals one of the λ_j's (see 5.16). In other words, λ is an eigenvalue of T if and only if λ equals one of the λ_j's, as desired. ■

Diagonal Matrices

A ***diagonal matrix*** is a square matrix that is 0 everywhere except possibly along the diagonal. For example,

$$\begin{bmatrix} 8 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 5 \end{bmatrix}$$

is a diagonal matrix. Obviously every diagonal matrix is upper triangular, although in general a diagonal matrix has many more 0's than an upper-triangular matrix.

An operator $T \in \mathcal{L}(V)$ has a diagonal matrix

$$\begin{bmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{bmatrix}$$

with respect to a basis $(v_1, \ldots, v_n)$ of V if and only

$$\begin{aligned} Tv_1 &= \lambda_1 v_1 \\ &\vdots \\ Tv_n &= \lambda_n v_n; \end{aligned}$$

this follows immediately from the definition of the matrix of an operator with respect to a basis. Thus an operator $T \in \mathcal{L}(V)$ has a diagonal matrix with respect to some basis of V if and only if V has a basis consisting of eigenvectors of T.

If an operator has a diagonal matrix with respect to some basis, then the entries along the diagonal are precisely the eigenvalues of the operator; this follows from 5.18 (or you may want to find an easier proof that works only for diagonal matrices).

Unfortunately not every operator has a diagonal matrix with respect to some basis. This sad state of affairs can arise even on complex vector spaces. For example, consider $T \in \mathcal{L}(\mathbf{C}^2)$ defined by

5.19 $$T(w, z) = (z, 0).$$

As you should verify, 0 is the only eigenvalue of this operator and the corresponding set of eigenvectors is the one-dimensional subspace $\{(w, 0) \in \mathbf{C}^2 : w \in \mathbf{C}\}$. Thus there are not enough linearly independent eigenvectors of T to form a basis of the two-dimensional space $\mathbf{C}^2$. Hence T does not have a diagonal matrix with respect to any basis of $\mathbf{C}^2$.

The next proposition shows that if an operator has as many distinct eigenvalues as the dimension of its domain, then the operator has a diagonal matrix with respect to some operator. However, some operators with fewer eigenvalues also have diagonal matrices (in other words, the converse of the next proposition is not true). For example, the operator T defined on the three-dimensional space $\mathbf{F}^3$ by

$$T(z_1, z_2, z_3) = (4z_1, 4z_2, 5z_3)$$

has only two eigenvalues (4 and 5), but this operator has a diagonal matrix with respect to the standard basis.

Later we will find other conditions that imply that certain operators have a diagonal matrix with respect to some basis (see 7.9 and 7.13).

5.20 Proposition: *If $T \in \mathcal{L}(V)$ has $\dim V$ distinct eigenvalues, then T has a diagonal matrix with respect to some basis of V.*

PROOF: Suppose that $T \in \mathcal{L}(V)$ has $\dim V$ distinct eigenvalues $\lambda_1, \dots, \lambda_{\dim V}$. For each j, let $v_j \in V$ be a nonzero eigenvector corresponding to the eigenvalue λ_j. Because nonzero eigenvectors corresponding to distinct eigenvalues are linearly independent (see 5.6), $(v_1, \dots, v_{\dim V})$ is linearly independent. A linearly independent list of

$\dim V$ vectors in V is a basis of V (see 2.17); thus $(v_1, \ldots, v_{\dim V})$ is a basis of V. With respect to this basis consisting of eigenvectors, T has a diagonal matrix. ■

We close this section with a proposition giving several conditions on an operator that are equivalent to its having a diagonal matrix with respect to some basis.

5.21 Proposition: *Suppose $T \in \mathcal{L}(V)$. Let $\lambda_1, \ldots, \lambda_m$ denote the distinct eigenvalues of T. Then the following are equivalent:*

For complex vector spaces, we will extend this list of equivalences later (see Exercises 16 and 23 in Chapter 8).

(a) *T has a diagonal matrix with respect to some basis of V;*

(b) *V has a basis consisting of eigenvectors of T;*

(c) *there exist one-dimensional subspaces $U_1, \ldots, U_n$ of V, each invariant under T, such that*

$$V = U_1 \oplus \cdots \oplus U_n;$$

(d) *$V = \operatorname{null}(T - \lambda_1 I) \oplus \cdots \oplus \operatorname{null}(T - \lambda_m I)$;*

(e) *$\dim V = \dim \operatorname{null}(T - \lambda_1 I) + \cdots + \dim \operatorname{null}(T - \lambda_m I)$.*

PROOF: We have already shown that (a) and (b) are equivalent.

Suppose that (b) holds; thus V has a basis $(v_1, \ldots, v_n)$ consisting of eigenvectors of T. For each j, let $U_j = \operatorname{span}(v_j)$. Obviously each U_j is a one-dimensional subspace of V that is invariant under T (because each v_j is an eigenvector of T). Because $(v_1, \ldots, v_n)$ is a basis of V, each vector in V can be written uniquely as a linear combination of $(v_1, \ldots, v_n)$. In other words, each vector in V can be written uniquely as a sum $u_1 + \cdots + u_n$, where each $u_j \in U_j$. Thus $V = U_1 \oplus \cdots \oplus U_n$. Hence (b) implies (c).

Suppose now that (c) holds; thus there are one-dimensional subspaces $U_1, \ldots, U_n$ of V, each invariant under T, such that

$$V = U_1 \oplus \cdots \oplus U_n.$$

For each j, let v_j be a nonzero vector in U_j. Then each v_j is an eigenvector of T. Because each vector in V can be written uniquely as a sum $u_1 + \cdots + u_n$, where each $u_j \in U_j$ (so each u_j is a scalar multiple of v_j), we see that $(v_1, \ldots, v_n)$ is a basis of V. Thus (c) implies (b).

At this stage of the proof we know that (a), (b), and (c) are all equivalent. We will finish the proof by showing that (b) implies (d), that (d) implies (e), and that (e) implies (b).

Suppose that (b) holds; thus V has a basis consisting of eigenvectors of T. Thus every vector in V is a linear combination of eigenvectors of T. Hence

5.22 $$V = \operatorname{null}(T - \lambda_1 I) + \cdots + \operatorname{null}(T - \lambda_m I).$$

To show that the sum above is a direct sum, suppose that

$$0 = u_1 + \cdots + u_m,$$

where each $u_j \in \operatorname{null}(T - \lambda_j I)$. Because nonzero eigenvectors corresponding to distinct eigenvalues are linearly independent, this implies (apply 5.6 to the sum of the nonzero vectors on the right side of the equation above) that each u_j equals 0. This implies (using 1.8) that the sum in 5.22 is a direct sum, completing the proof that (b) implies (d).

That (d) implies (e) follows immediately from Exercise 17 in Chapter 2.

Finally, suppose that (e) holds; thus

5.23 $$\dim V = \dim \operatorname{null}(T - \lambda_1 I) + \cdots + \dim \operatorname{null}(T - \lambda_m I).$$

Choose a basis of each $\operatorname{null}(T - \lambda_j I)$; put all these bases together to form a list $(v_1, \ldots, v_n)$ of eigenvectors of T, where $n = \dim V$ (by 5.23). To show that this list is linearly independent, suppose

$$a_1 v_1 + \cdots + a_n v_n = 0,$$

where $a_1, \ldots, a_n \in \mathbf{F}$. For each $j = 1, \ldots, m$, let u_j denote the sum of all the terms $a_k v_k$ such that $v_k \in \operatorname{null}(T - \lambda_j I)$. Thus each u_j is an eigenvector of T with eigenvalue λ_j, and

$$u_1 + \cdots + u_m = 0.$$

Because nonzero eigenvectors corresponding to distinct eigenvalues are linearly independent, this implies (apply 5.6 to the sum of the nonzero vectors on the left side of the equation above) that each u_j equals 0. Because each u_j is a sum of terms $a_k v_k$, where the v_k's were chosen to be a basis of $\operatorname{null}(T - \lambda_j I)$, this implies that all the a_k's equal 0. Thus $(v_1, \ldots, v_n)$ is linearly independent and hence is a basis of V (by 2.17). Thus (e) implies (b), completing the proof. ∎

Invariant Subspaces on Real Vector Spaces

We know that every operator on a complex vector space has an eigenvalue (see 5.10 for the precise statement). We have also seen an example showing that the analogous statement is false on real vector spaces. In other words, an operator on a nonzero real vector space may have no invariant subspaces of dimension 1. However, we now show that an invariant subspace of dimension 1 or 2 always exists.

5.24 Theorem: *Every operator on a finite-dimensional, nonzero, real vector space has an invariant subspace of dimension* 1 *or* 2*.*

PROOF: Suppose V is a real vector space with dimension $n > 0$ and $T \in \mathcal{L}(V)$. Choose $v \in V$ with $v \neq 0$. Then

$$(v, Tv, T^2v, \dots, T^nv)$$

cannot be linearly independent because V has dimension n and we have $n+1$ vectors. Thus there exist real numbers $a_0, \dots, a_n$, not all 0, such that

$$0 = a_0v + a_1Tv + \cdots + a_nT^nv.$$

Make the a's the coefficients of a polynomial, which can be written in factored form (see 4.14) as

$$\begin{aligned} &a_0 + a_1x + \cdots + a_nx^n \\ &\quad = c(x - \lambda_1)\dots(x - \lambda_m)(x^2 + \alpha_1x + \beta_1)\dots(x^2 + \alpha_Mx + \beta_M), \end{aligned}$$

Here either m or M might equal 0.

where c is a nonzero real number, each λ_j, α_j, and β_j is real, $m+M \geq 1$, and the equation holds for all $x \in \mathbf{R}$. We then have

$$\begin{aligned} 0 &= a_0v + a_1Tv + \cdots + a_nT^nv \\ &= (a_0I + a_1T + \cdots + a_nT^n)v \\ &= c(T - \lambda_1I)\dots(T - \lambda_mI)(T^2 + \alpha_1T + \beta_1I)\dots(T^2 + \alpha_MT + \beta_MI)v, \end{aligned}$$

which means that $T - \lambda_jI$ is not injective for at least one j or that $(T^2 + \alpha_jT + \beta_jI)$ is not injective for at least one j. If $T - \lambda_jI$ is not injective for at least one j, then T has an eigenvalue and hence a one-dimensional invariant subspace. Let's consider the other possibility. In other words, suppose that $(T^2 + \alpha_jT + \beta_jI)$ is not injective for some j. Thus there exists a nonzero vector $u \in V$ such that

5.25 $$T^2u + \alpha_j Tu + \beta_j u = 0.$$

We will complete the proof by showing that $\operatorname{span}(u, Tu)$, which clearly has dimension 1 or 2, is invariant under T. To do this, consider a typical element of $\operatorname{span}(u, Tu)$ of the form $au + bTu$, where $a, b \in \mathbf{R}$. Then

$$\begin{aligned} T(au + bTu) &= aTu + bT^2u \\ &= aTu - b\alpha_j Tu - b\beta_j u, \end{aligned}$$

where the last equality comes from solving for T^2u in 5.25. The equation above shows that $T(au + bTu) \in \operatorname{span}(u, Tu)$. Thus $\operatorname{span}(u, Tu)$ is invariant under T, as desired. ■

We will need one new piece of notation for the next proof. Suppose U and W are subspaces of V with

$$V = U \oplus W.$$

Each vector $v \in V$ can be written uniquely in the form

$$v = u + w,$$

$P_{U,W}$ is often called the ***projection*** *onto U with null space W.*

where $u \in U$ and $w \in W$. With this representation, define $P_{U,W} \in \mathcal{L}(V)$ by

$$P_{U,W}v = u.$$

You should verify that $P_{U,W}v = v$ if and only if $v \in U$. Interchanging the roles of U and W in the representation above, we have $P_{W,U}v = w$. Thus $v = P_{U,W}v + P_{W,U}v$ for every $v \in V$. You should verify that ${P_{U,W}}^2 = P_{U,W}$; furthermore $\operatorname{range} P_{U,W} = U$ and $\operatorname{null} P_{U,W} = W$.

We have seen an example of an operator on $\mathbf{R}^2$ with no eigenvalues. The following theorem shows that no such example exists on $\mathbf{R}^3$.

5.26 Theorem: *Every operator on an odd-dimensional real vector space has an eigenvalue.*

PROOF: Suppose V is a real vector space with odd dimension. We will prove that every operator on V has an eigenvalue by induction (in steps of size 2) on the dimension of V. To get started, note that the desired result obviously holds if $\dim V = 1$.

Now suppose that $\dim V$ is an odd number greater than 1. Using induction, we can assume that the desired result holds for all operators on all real vector spaces with dimension 2 less than $\dim V$. Suppose $T \in \mathcal{L}(V)$. We need to prove that T has an eigenvalue. If it does, we are done. If not, then by 5.24 there is a two-dimensional subspace U of V that is invariant under T. Let W be any subspace of V such that

$$V = U \oplus W;$$

2.13 guarantees that such a W exists.

Because W has dimension 2 less than $\dim V$, we would like to apply our induction hypothesis to $T|_W$. However, W might not be invariant under T, meaning that $T|_W$ might not be an operator on W. We will compose with the projection $P_{W,U}$ to get an operator on W. Specifically, define $S \in \mathcal{L}(W)$ by

$$Sw = P_{W,U}(Tw)$$

for $w \in W$. By our induction hypothesis, S has an eigenvalue λ. We will show that this λ is also an eigenvalue for T.

Let $w \in W$ be a nonzero eigenvector for S corresponding to the eigenvalue λ; thus $(S - \lambda I)w = 0$. We would be done if w were an eigenvector for T with eigenvalue λ; unfortunately that need not be true. So we will look for an eigenvector of T in $U + \operatorname{span}(w)$. To do that, consider a typical vector $u + aw$ in $U + \operatorname{span}(w)$, where $u \in U$ and $a \in \mathbf{R}$. We have

$$\begin{aligned}(T - \lambda I)(u + aw) &= Tu - \lambda u + a(Tw - \lambda w)\\ &= Tu - \lambda u + a(P_{U,W}(Tw) + P_{W,U}(Tw) - \lambda w)\\ &= Tu - \lambda u + a(P_{U,W}(Tw) + Sw - \lambda w)\\ &= Tu - \lambda u + aP_{U,W}(Tw).\end{aligned}$$

Note that on the right side of the last equation, $Tu \in U$ (because U is invariant under T), $\lambda u \in U$ (because $u \in U$), and $aP_{U,W}(Tw) \in U$ (from the definition of $P_{U,W}$). Thus $T - \lambda I$ maps $U + \operatorname{span}(w)$ into U. Because $U + \operatorname{span}(w)$ has a larger dimension than U, this means that $(T - \lambda I)|_{U+\operatorname{span}(w)}$ is not injective (see 3.5). In other words, there exists a nonzero vector $v \in U + \operatorname{span}(w) \subset V$ such that $(T - \lambda I)v = 0$. Thus T has an eigenvalue, as desired. ∎

Exercises

1. Suppose $T \in \mathcal{L}(V)$. Prove that if $U_1, \dots, U_m$ are subspaces of V invariant under T, then $U_1 + \cdots + U_m$ is invariant under T.

2. Suppose $T \in \mathcal{L}(V)$. Prove that the intersection of any collection of subspaces of V invariant under T is invariant under T.

3. Prove or give a counterexample: if U is a subspace of V that is invariant under every operator on V, then $U = \{0\}$ or $U = V$.

4. Suppose that $S, T \in \mathcal{L}(V)$ are such that $ST = TS$. Prove that $\operatorname{null}(T - \lambda I)$ is invariant under S for every $\lambda \in \mathbf{F}$.

5. Define $T \in \mathcal{L}(\mathbf{F}^2)$ by
$$T(w, z) = (z, w).$$
Find all eigenvalues and eigenvectors of T.

6. Define $T \in \mathcal{L}(\mathbf{F}^3)$ by
$$T(z_1, z_2, z_3) = (2z_2, 0, 5z_3).$$
Find all eigenvalues and eigenvectors of T.

7. Suppose n is a positive integer and $T \in \mathcal{L}(\mathbf{F}^n)$ is defined by
$$T(x_1, \dots, x_n) = (x_1 + \cdots + x_n, \dots, x_1 + \cdots + x_n);$$
in other words, T is the operator whose matrix (with respect to the standard basis) consists of all 1's. Find all eigenvalues and eigenvectors of T.

8. Find all eigenvalues and eigenvectors of the backward shift operator $T \in \mathcal{L}(\mathbf{F}^\infty)$ defined by
$$T(z_1, z_2, z_3, \dots) = (z_2, z_3, \dots).$$

9. Suppose $T \in \mathcal{L}(V)$ and $\dim \operatorname{range} T = k$. Prove that T has at most $k + 1$ distinct eigenvalues.

10. Suppose $T \in \mathcal{L}(V)$ is invertible and $\lambda \in \mathbf{F} \setminus \{0\}$. Prove that λ is an eigenvalue of T if and only if $\frac{1}{\lambda}$ is an eigenvalue of T^{-1}.

11. Suppose $S, T \in \mathcal{L}(V)$. Prove that ST and TS have the same eigenvalues.

12. Suppose $T \in \mathcal{L}(V)$ is such that every vector in V is an eigenvector of T. Prove that T is a scalar multiple of the identity operator.

13. Suppose $T \in \mathcal{L}(V)$ is such that every subspace of V with dimension $\dim V - 1$ is invariant under T. Prove that T is a scalar multiple of the identity operator.

14. Suppose $S, T \in \mathcal{L}(V)$ and S is invertible. Prove that if $p \in \mathcal{P}(\mathbf{F})$ is a polynomial, then

$$p(STS^{-1}) = Sp(T)S^{-1}.$$

15. Suppose $\mathbf{F} = \mathbf{C}$, $T \in \mathcal{L}(V)$, $p \in \mathcal{P}(\mathbf{C})$, and $a \in \mathbf{C}$. Prove that a is an eigenvalue of $p(T)$ if and only if $a = p(\lambda)$ for some eigenvalue λ of T.

16. Show that the result in the previous exercise does not hold if $\mathbf{C}$ is replaced with $\mathbf{R}$.

17. Suppose V is a complex vector space and $T \in \mathcal{L}(V)$. Prove that T has an invariant subspace of dimension j for each $j = 1, \dots, \dim V$.

18. Give an example of an operator whose matrix with respect to some basis contains only 0's on the diagonal, but the operator is invertible.

These two exercises show that 5.16 fails without the hypothesis that an upper-triangular matrix is under consideration.

19. Give an example of an operator whose matrix with respect to some basis contains only nonzero numbers on the diagonal, but the operator is not invertible.

20. Suppose that $T \in \mathcal{L}(V)$ has $\dim V$ distinct eigenvalues and that $S \in \mathcal{L}(V)$ has the same eigenvectors as T (not necessarily with the same eigenvalues). Prove that $ST = TS$.

21. Suppose $P \in \mathcal{L}(V)$ and $P^2 = P$. Prove that $V = \operatorname{null} P \oplus \operatorname{range} P$.

22. Suppose $V = U \oplus W$, where U and W are nonzero subspaces of V. Find all eigenvalues and eigenvectors of $P_{U,W}$.

23. Give an example of an operator $T \in \mathcal{L}(\mathbf{R}^4)$ such that T has no (real) eigenvalues.

24. Suppose V is a real vector space and $T \in \mathcal{L}(V)$ has no eigenvalues. Prove that every subspace of V invariant under T has even dimension.