The Goal of Linear Algebra
Northwestern University, Summer 2015

The purpose of these notes is to outline what I consider to be the main goal of an abstract linear algebra course, namely:

Main Goal. To understand how properties of linear operators can be reflected in ways of decomposing a vector space into a sum of subspaces and in the types of matrices which can be used to describe that operator relative to a well-chosen basis.

We will see this idea come up again and again, and viewing many of the results we'll see in terms of this one unifying framework will help to understand that they "really" mean. In particular, the various ways of characterizing diagonalizability, the existence of upper-triangular matrices representing operators on finite-dimensional complex vector spaces, the real and complex Spectral Theorems, and the existence of Jordan forms are all results which fit into this framework, and are all in the book.

Here we give more examples showing other ways in which this goal manifests itself in this course. These examples are not explicitly given in the book; there are exercises in the book which touch on them a bit, but they are not phrased in the context of the main goal I suggested above.

Projections

Suppose that $P : V \to V$ is a linear operator on V such that $P^2 = P$, meaning that applying P twice is the same as applying P once. Such an operator is called a *projection*.

Proposition. *If $P \in \mathcal{L}(V)$ is a projection, then $V = \operatorname{null} P \oplus \operatorname{range} P$.*

Proof. For any $v \in V$, we have:

$$P(v - Pv) = Pv - P^2v = Pv - Pv = 0$$

where we use the fact that P is a projection to say $P^2v = Pv$. Thus for any $v \in V$, $v - Pv \in \operatorname{null} P$. Hence

$$v = (v - Pv) + Pv$$

expresses any $v \in V$ as the sum of an element $v - Pv \in \operatorname{null} P$ and an element $Pv \in \operatorname{range} P$, so $V = \operatorname{null} P + \operatorname{range} P$.

To show that this is a direct sum, suppose that $x \in \operatorname{null} P \cap \operatorname{range} P$, so that $x \in \operatorname{null} P$ and $x \in \operatorname{range} P$. Since $x \in \operatorname{range} P$, there is some $u \in V$ such that $x = Pu$, and since $x \in \operatorname{null} P$ we have $Px = 0$. Thus

$$0 = Px = P^2u = Pu = x,$$

so $x = 0$ and therefore $\operatorname{null} P \cap \operatorname{range} P = \{0\}$. We conclude that $V = \operatorname{null} P \oplus \operatorname{range} P$.

Alternatively, to see that this is a direct sum, suppose that $v = x + Pu$ expresses an element $v \in V$ as a sum of some $x \in \operatorname{null} P$ and some $Pu \in \operatorname{range} P$. Then

$$Pv = P(x + Pu) = Px + P^2u = Pu$$

where we use that $Px = 0$ since $x \in \operatorname{null} P$ and the fact that $P^2 = P$. This shows that in the expression $v = x + Pu$, Pu must be Pv, and hence in turn x must be $v - Pv$. Thus the sum we have above:

$$v = (v - Pv) + Pv$$

is the only possible way of writing $v \in V$ as a sum of elements of $\operatorname{null} P$ and $\operatorname{range} P$, so the sum $V = \operatorname{null} P + \operatorname{range} P$ is a direct sum as claimed. $\square$

Thus, the property that P is a projection (i.e. $P^2 = P$) is reflected in the decomposition $V = \operatorname{null} P \oplus \operatorname{range} P$. Note, however, that the converse of the above result does not hold: having $V = \operatorname{null} P \oplus \operatorname{range} P$ does not ensure that P is a projection—this will only be the case if in addition we know that the restriction of P to $\operatorname{range} P$ is the identity.

Using this decomposition, we can now determine a "nice" way of representing projections via matrices. Take a basis $u_1, \ldots, u_k$ for $\operatorname{null} P$ and a basis $Pw_1, \ldots, Pw_\ell$ for $\operatorname{range} P$. Since $V = \operatorname{null} P \oplus \operatorname{range} P$, the u's and w's together give a basis for V. The $(k+\ell) \times (k+\ell)$ matrix of P relative to this basis can be thought of as consisting of four blocks

$$\mathcal{M}(P) = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

as follows. Using the decomposition $V = \operatorname{null} P \oplus \operatorname{range} P$, we can imagine applying P only to elements of $\operatorname{null} P$ or only to elements of $\operatorname{range} P$, and we can imagine writing the resulting vectors as a sum of an element of $\operatorname{null} P$ and an element of $\operatorname{range} P$. Then,

- A is the $k \times k$ matrix which describes the $(\operatorname{null} P)$-component of the behavior of P on $\operatorname{null} P$,
- B is the $k \times \ell$ matrix which describes the $(\operatorname{null} P)$-component of the behavior of P on $\operatorname{range} P$,
- C is the $\ell \times k$ matrix describing the $(\operatorname{range} P)$-component of the behavior of P on $\operatorname{null} P$, and
- D the $\ell \times \ell$ matrix describe the $(\operatorname{range} P)$-component of the behavior of P on $\operatorname{range} P$.

By $(\operatorname{null} P)$-component I mean the element of $\operatorname{null} P$ needed to write a vector as a sum according to $V = \operatorname{null} P + \operatorname{range} P$, and similarly for $(\operatorname{range} P)$-component.

To be clear, A and C encode what happens when applying P to the terms of the given basis coming from $\operatorname{null} P$. Since each u_i is in $\operatorname{null} P$, $Pu_i = 0$ for all i so A and C should both be zero; in other words, the $(\operatorname{null} P)$-component and $(\operatorname{range} P)$-component of $Pu_i = 0$ are both zero. Now, B and D encode what happens when applying P to one of the Pw_i's coming from $\operatorname{range} P$. We have

$$P(Pw_i) = Pw_i$$

for each i since $P^2 = P$, meaning that P sends something in $\operatorname{range} P$ to itself. Thus, the $(\operatorname{null} P)$-component of any Pw_i is zero since we don't need a contribution from $\operatorname{null} P$ when writing Pw_i according to the decomposition $V = \operatorname{null} P \oplus \operatorname{range} P$, and the $(\operatorname{range} P)$-component of Pw_i is Pw_i itself. Hence B should be the zero matrix and D the identity since P behaves like the identity when restricted to $\operatorname{range} P$ because it sends a basis vector in this range to itself.

Thus, with respect to the basis chosen according to the decomposition $V = \operatorname{null} P \oplus \operatorname{range} P$, the matrix of P looks like

$$\mathcal{M}(P) = \begin{pmatrix} 0 & 0 \\ 0 & I \end{pmatrix}.$$

Conversely, an operator such that there is some basis relative to which it has the above form must in fact be a projection, so having a matrix of this form representing an operator also reflects the fact that it is a projection. Note that a matrix of the above form satisfies:

$$\begin{pmatrix} 0 & 0 \\ 0 & I \end{pmatrix}^2 = \begin{pmatrix} 0 & 0 \\ 0 & I \end{pmatrix},$$

which mimics the projection property $P^2 = P$.

Summary. To summarize, in relation to our main goal: the property $P^2 = P$ of a projection is reflected in the direct sum decomposition $V = \operatorname{null} P \oplus \operatorname{range} P$ and also in the form $\left(\begin{smallmatrix} 0 & 0 \\ 0 & I \end{smallmatrix}\right)$ a matrix representing P can have.

Why is such a P called a projection? Because when writing $V = \operatorname{range} P \oplus \operatorname{null} P$, P is what tells us the result of *projecting* a vector in V onto the subspace range P; in other words, using the notation from the homework $P_{U,W}$ for the map which projects a vector in $V = U \oplus W$ onto U, P here is precisely the projection $P = P_{\operatorname{range} P, \operatorname{null} P}$.

Block Diagonal Matrices

Recall that a subspace U of a vector space V is T-*invariant*, where T is some linear operator on V, if $Tu \in U$ for all $u \in U$. Suppose that

$$V = U \oplus W$$

where U and W are each T-invariant subspaces of V. Taking bases $u_1, \ldots, u_m$ for U and $w_1, \ldots, w_k$ for W gives a basis $u_1, \ldots, u_m, w_1, \ldots, w_k$ for V. We want to know what the matrix of T relative to this basis looks like?

As before, we think of this matrix as consisting of four smaller pieces:

$$\mathcal{M}(T) = \begin{pmatrix} A & B \\ C & D \end{pmatrix},$$

where A is the $m \times m$ matrix characterizing the contribution from U needed when expressing the result of applying T to a vector in U, C is the $k \times m$ matrix characterizing the contribution from W needed when expressing the result of applying T to a vector in U, B tells us the contribution from U needed when applying T to something in W, and D tells us the contribution from W needed when applying T to something in W. The first column of this matrix is supposed to encode the coefficients in

$$Tu_1 = a_{11}u_1 + \cdots + a_{m1}u_m + c_{11}w_1 + \cdots + c_{k1}w_k.$$

But $u_1 \in U$ and U is T-invariant, so $Tu_1 \in U$ and hence Tu_1 should be expressible solely in terms of $u_1, \ldots, u_m$. Thus all the c_{j1} coefficients must be zero, so the first column of $\mathcal{M}(T)$ will consist of zeroes once we get the the "C" piece. Indeed, we claim that C must be the zero matrix since the columns to which this contributes give the coefficients of the w_j's needed when expressing Tu_i, and the point is that these coefficients are all zero since $Tu_i \in U$ due to the invariance of U, so no contribution from W is needed.

Similarly, since W is T-invariant, the result of applying T to something in W is expressible solely in terms of W itself, with no contribution from U needed. This means that the B piece in the matrix above must also be zero. (Concretely, since $Tw_j \in W$ for all $w_j \in W$, writing Tw_j as a linear combination of our given basis requires that the coefficients of all the u_i's be zero, and these coefficients make up the entries of B.) Thus the matrix T relative to this basis must be of the form

$$\mathcal{M}(T) = \begin{pmatrix} A & 0 \\ 0 & D \end{pmatrix}.$$

We can think of A is the matrix of the restriction $T|_U : U \to U$ relative to the basis $u_1, \ldots, u_m$, and D as the matrix of the restriction $T|_W : W \to W$ relative to the basis $w_1, \ldots, w_k$. Such a matrix is called *block diagonal* since it looks "diagonal", only that the diagonal pieces aren't simply single entries but rather smaller sized matrices.

More generally, given a decomposition $V = U_1 \oplus \cdots \oplus U_m$ into T-invariant subspaces $U_1, \ldots, U_m$, the matrix of T relative to a basis for V constructed by taking bases for each U_i will be block diagonal:

$$\mathcal{M}(T) = \begin{pmatrix} A_1 & & \\ & \ddots & \\ & & A_m \end{pmatrix}$$

where A_i is a matrix of size $(\dim U_i) \times (\dim U_i)$. Conversely, the only way to obtain such a matrix relative to a basis is for the vectors in that basis to come from subspaces which are invariant under T, since the form this matrix has will imply that if $u \in U_i$, Tu should be expressible solely in terms of the chosen basis for U_i, which implies that $Tu \in U_i$. The upshot is that block diagonal matrices arise precisely when decomposing a vector space into *invariant* subspaces.

Summary. In relation to our main goal, the property of being able to decompose a vector space into a direct sum $V = U_1 \oplus \cdots \oplus U_m$ of T-invariant subspaces U_i is reflected by the possibility of being able to represent T by a block diagonal matrix.

Nilpotent Operators

A linear operator $T \in \mathcal{L}(V)$ is said to be *nilpotent* if there is some $n \geq 1$ such that $T^n = 0$, where 0 here denotes the zero operator which sends everything to zero.

Proposition. *An operator T on a finite-dimensional complex vector space is nilpotent if and only if 0 is its only eigenvalue.*

Proof using matrices. (The forward direction—that the only possible eigenvalue of a nilpotent operator is zero—is actually true for any vector space, including real or infinite dimensional ones. It is the backwards direction which requires that V be finite-dimensional and complex.)

Suppose T is nilpotent and that λ is an eigenvalue of T. Let $v \neq 0$ be a corresponding eigenvector and let $n \geq 1$ satisfy $T^n = 0$. Note that

$$T^2 v = T(Tv) = T(\lambda v) = \lambda T v = \lambda^2 v,$$

and more generally $T^k v = \lambda^k v$; that is, since T scales v by a factor of λ, applying T repeatedly amounts to scaling by that same factor repeatedly. Thus $T^n v = \lambda^n v$, but on the other hand $T^n v = 0$ since $T^n = 0$. Hence $\lambda^n v = 0$ and since $v \neq 0$, this means $\lambda^n = 0$ so $\lambda = 0$ as claimed.

Conversely, suppose that λ is the only eigenvalue of T. Since T is an operator on a finite-dimensional complex vector space, there exists a basis of V relative to which $\mathcal{M}(T)$ is upper-triangular. The diagonal entries in this upper-triangular matrix are the eigenvalues of T, so since T only has zero as an eigenvalue, $\mathcal{M}(T)$ must look like:

$$\mathcal{M}(T) = \begin{pmatrix} 0 & * & \cdots & * \\ & 0 & \ddots & \vdots \\ & & \ddots & * \\ & & & 0 \end{pmatrix}$$

where the $*$'s denote some unknown entries which could be anything and blank spaces denote zeroes. The key observation is that multiplying such a matrix by itself will give a similar matrix only with the entries above the diagonal now forced to all be zero; for instance,

$$\begin{pmatrix} 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \\ 0 & 0 & 0 & 0 \end{pmatrix}^2 = \begin{pmatrix} 0 & 0 & * & * \\ 0 & 0 & 0 & * \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Taking a higher power again introduces a new "diagonal" of zeroes located further to the "upper-right":

$$\begin{pmatrix} 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \\ 0 & 0 & 0 & 0 \end{pmatrix}^3 = \begin{pmatrix} 0 & 0 & 0 & * \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

and so on. The fact is that each time we take a higher power of an upper triangular matrix with all diagonal entries equal to 0, we get a matrix with more and more zeroes where the only possible nonzero entries are being "pushed" off towards the upper right.

Thus for $n = \dim V$, the n-th power $\mathcal{M}(T)^n$ of $\mathcal{M}(T)$ will have all zero entries since at this stage all nonzero entries have been pushed off. But $\mathcal{M}(T)^n = \mathcal{M}(T^n)$, so if this is the zero matrix then T^n must be the zero operator, so T is nilpotent. □

Note already how the proof of this proposition ties in with our main goal: the property of T being nilpotent is reflected in the fact that it can be represented by an upper-triangular matrix with all diagonal entries equal to zero.

Proof without using matrices. We now give another proof of the backwards direction of the above proposition, which is really just the matrix proof above only phrased in terms of bases instead. Again, since T is an operator on a finite-dimensional complex vector space, there exists a basis $v_1, \ldots, v_n$ such that

$$Tv_i \in \operatorname{span}(v_1, \ldots, v_i) \text{ for each } i,$$

or equivalently such that $\operatorname{span}(v_1, \ldots, v_i)$ is T-invariant for each i. (As we saw in class or as the book shows, the point is that this is the property which makes $\mathcal{M}(T)$ be upper triangular.) In particular, for each i there are scalars such that

$$Tv_i = a_{i1}v_1 + \cdots + a_{ii}v_i.$$

The scalars a_{ii} form the diagonal of $\mathcal{M}(T)$, and hence are the eigenvalues of T and so are all zero in our scenario. Thus, we get that

$$Tv_i \in \operatorname{span}(v_1, \ldots, v_{i-1}) \text{ for each } i,$$

so that each Tv_i can be expressed as a linear combination solely of the basis vectors *before* v_i itself:

$$Tv_1 = 0,\ Tv_2 \in \operatorname{span}(v_1),\ Tv_3 \in \operatorname{span}(v_1, v_2),\ \text{etc.}$$

where $Tv_1 = 0$ reflects the fact that there are no basis vectors before v_1.

Now, consider $T^2 v_i$ for some i. We have that Tv_i is expressible as

$$Tv_i = b_1v_1 + \cdots + b_{i-1}v_{i-1}.$$

Applying T again gives

$$T^2 v_i = b_1 T v_1 + \cdots + b_{i-1} T v_{i-1}.$$

But each term Tv_j here is expressible solely in terms of the previous vectors, and in particular Tv_{i-1} is expressible in terms of all vectors up to and including v_{i-2}. This implies that the above expression for $T^2 v_i$ can be written in the form

$$T^2 v_i = c_1 v_1 + \cdots + c_{i-2} v_{i-2}$$

for some scalars c_j, so that $T^2 v_i \in \operatorname{span}(v_1, \ldots, v_{i-2})$. And so on, the point is that each application of T gives a vector which is expressible using one fewer v_i. Thus, since Tv_2 is expressible using only v_1, $T^2 v_2$ should be expressible using none of the v_i's, so $T^2 v_2 = 0$; since Tv_3 is expressible using only v_1 and v_2, $T^2 v_3$ is expressible using only v_1, and hence $T^3 v = 0$; and in general $T^{n-1} v_i$ is expressible using only v_1, so $T^n v_i = 0$ for all i. Since T^n sends each basis vector to zero, it must send everything to zero so $T^n = 0$ as claimed. $\square$

Summary. The property $T^n = 0$ of T being nilpotent is reflected by the fact that there is a basis for which $Tv_i \in \operatorname{span}(v_1, \ldots, v_{i-1})$ for each i, and by the fact that T can be represented by an upper triangular matrix with all zeroes (the eigenvalues) down the diagonal.