Math 350 Fall 2011
Notes about Diagonalization and Invariant subspaces

1. Diagonalization of a matrix

In this section, let A be an $n \times n$ matrix with entries in a field F.

Definition 1. *The* characteristic polynomial *of A is by definition the polynomial*

$$f(t) = \det(A - tI),$$

where I is the $n \times n$ identity matrix.

It is thus a polynomial of degree n in a single variable t with coefficients in F.

Question 1. *Consider the following 2×2 real matrix*

$$\begin{pmatrix} 1 & -4 \\ 0 & -1 \end{pmatrix}.$$

Compute its characteristic polynomial.

Definition 2. *λ is said to be an* eigenvalue *of A if it is a zero of the characteristic polynomial of A in F.*

So $\lambda \in F$ is an eigenvalue of A, if and only if $\det(A - \lambda I) = 0$. It then follows that the nullspace of $A - \lambda I$ is non-zero; so there exists a non-zero vector $v \in F^n$ such that $(A - \lambda I)v = 0$, i.e. $Av = \lambda v$.

Definition 3. *A vector $v \in F^n$ is said to be an* eigenvector *of A, if $v \neq 0$ and $Av = \lambda v$ for some $\lambda \in F$.*

In fact if $v \neq 0$ and $Av = \lambda v$, then v is said to be an eigenvector of A *corresponding to the eigenvalue* λ. Equivalently, if $v \neq 0$, then v is an eigenvector of A corresponding to λ, if and only if v is in the nullspace of $A - \lambda I$.

Question 2. *Find all eigenvalues and eigenvectors of the matrix in Question 1.*

Definition 4. *A is said to be* diagonalizable *(over F) if there exists an invertible matrix P, and a diagonal matrix D, both of which are $n \times n$ and have entries in F, such that*

$$A = PDP^{-1}.$$

If A is diagonalizable, a representation of A in the form PDP^{-1}, where P is invertible and D is diagonal, is called a diagonalization *of A.*

In other words, A is diagonalizable if and only if it is similar to a diagonal matrix (over F).

Question 3. *Show that the matrix in Question 1 over $\mathbb{R}$ is diagonalizable over $\mathbb{R}$, and diagonalize it.*

Theorem 1. *A is diagonalizable over F, if and only if A has n linearly independent eigenvectors in F^n.*

Proof. Suppose $A = PDP^{-1}$ for some invertible matrix P and a diagonal matrix D. Then $AP = PD$, so if $v_1, \dots, v_n$ are the columns of P, and $\lambda_1, \dots, \lambda_n$ are the entries on the diagonal of D, then $Av_i = \lambda_i v_i$ for all $i = 1, \dots, n$. In particular, $v_1, \dots, v_n$ are eigenvectors of A. They are linearly independent since P is invertible.

Conversely, suppose $v_1, \dots, v_n$ are linearly independent eigenvectors of A, say corresponding to eigenvalues $\lambda_1, \dots, \lambda_n \in F$. Let P be the matrix whose i-th row is v_i, and D be the diagonal matrix whose i-th entry on the diagonal is λ_i. Then P is invertible, since its columns are linearly independent, and we have $AP = PD$, since $Av_i = \lambda_i v_i$ for $i = 1, 2, \dots, n$. It follows that $A = PDP^{-1}$, which implies that A is diagonalizable over F. □

Definition 5. *A set of n linearly independent eigenvectors of A in F^n is said to be an* eigenbasis *of F^n associated to A.*

The above theorem can now be reformulated as

Theorem 2. *A is diagonalizable over F, if and only if there exists an eigenbasis of F^n associated to A.*

Definition 6. *The span of all eigenvectors of A corresponding to an eigenvalue λ of A is called the* eigenspace *of A corresponding to λ.*

Thus any eigenspace of A is in particular a subspace of F^n. In general one should think about eigenspaces instead of eigenvectors as much as possible, since eigenspaces are *canonical* objects associated to a matrix, while eigenvectors are not.

2. Diagonalization of a linear map

In this and subsequent sections, V is a vector space over F of dimension $n < \infty$, and $T: V \to V$ is a linear map from V *into itself*.

Recall that given any ordered basis α of V, one can find a matrix representation $[T]_\alpha$ of T with respect to α. This matrix is an $n \times n$ matrix with entries in F, and if β is a different ordered basis of V, then $[T]_\beta$ and $[T]_\alpha$ are similar to each other. Since the determinant of any two similar matrices are the same, we can make the following definition:

Definition 7. *The determinant of a linear map $T: V \to V$ is defined to be the determinant of $[T]_\alpha$, where α is any ordered basis of V. We denote this by* $\det(T)$.

This definition is independent of the choice of α. With this we define:

Definition 8. *The* characteristic polynomial *of a linear map $T: V \to V$ is defined to be the polynomial $f(t) = \det(T - tI)$, where $I: V \to V$ is the identity map.*

In other words, the characteristic polynomial of T is given by

$$f(t) = \det([T]_\alpha - tI),$$

where α is any ordered basis of V and I is the identity matrix. This is independent of the choice of α.

Question 4. *Let V be the space of real polynomials of degree ≤ 2. It is a vector space over $\mathbb{R}$. Let $T\colon V \to V$ be the linear map*

$$T(p(x)) = p'(x).$$

Find the characteristic polynomial of T.

Definition 9. *λ is said to be an* eigenvalue *of T if it is a zero of the characteristic polynomial of T in F.*

So $\lambda \in F$ is an eigenvalue of T, if and only if $\det(T - \lambda I) = 0$. It then follows that the nullspace of $T - \lambda I$ is non-zero; so there exists a non-zero vector $v \in V$ such that[1] $(T - \lambda I)v = 0$, i.e. $Tv = \lambda v$. Conversely, if there is a non-zero vector $v \in V$ such that $Tv = \lambda v$, then λ is an eigenvalue of T.

Definition 10. *A vector $v \in V$ is said to be an* eigenvector *of T, if $v \neq 0$ and $Tv = \lambda v$ for some $\lambda \in F$.*

In fact if $v \neq 0$ and $Tv = \lambda v$, then v is said to be an eigenvector of T *corresponding to the eigenvalue* λ. Equivalently, if $v \neq 0$, then v is an eigenvector of T corresponding to λ, if and only if v is in the nullspace of $T - \lambda I$.

Question 5. *Find all the eigenvalues and eigenvectors of T where T is the linear map in Question 4.*

Definition 11. *A linear map $T\colon V \to V$ is said to be* diagonalizable*, if there exists an ordered basis α of V such that the matrix representation $[T]_\alpha$ of T is a diagonal matrix.*

It follows that such an ordered basis of V consists only of eigenvectors of T; in fact $\alpha = \{v_1, \dots, v_n\}$, and $[T]_\alpha$ is a diagonal matrix with entries $\lambda_1, \dots, \lambda_n$ on the diagonal, then $Tv_i = \lambda_i v_i$ for all $i = 1, \dots, n$, and thus α consists of eigenvectors of T.

Conversely, if α is an ordered basis of T consisting only of eigenvectors of T, say $\alpha = \{v_1, \dots, v_n\}$, and $Tv_i = \lambda_i v_i$ for all $i = 1, \dots, n$ where $\lambda_1, \dots, \lambda_n \in F$, then the matrix representation $[T]_\alpha$ of T is a diagonal matrix with entries $\lambda_1, \dots, \lambda_n$, and T is diagonalizable. Thus we have proved:

Theorem 3. *A linear map $T\colon V \to V$ is diagonalizable, if and only if there exists a basis of V that consists only of eigenvectors of T.*

Such a basis is usually called an *eigenbasis* of V associated to T. Since in a vector space of dimension n, n vectors form a basis if and only if they are linearly independent, we have

Theorem 4. *A linear map $T\colon V \to V$ is diagonalizable, if and only if there exists n linearly independent eigenvectors of T, where $n = \dim(V)$.*

[1] From now on we write Tv for $T(v)$ for the sake of brevity.

Question 6. *Show that the linear map T in Question 4 is not diagonalizable.*

Definition 12. *If $T: V \to V$ is diagonalizable, then a* diagonalization *of T is by definition an ordered eigenbasis of V associated to T, together with the matrix representation $[T]_\alpha$.*

In other words, it is a set of n linearly independent eigenvectors of T, and the diagonal matrix that consists of their corresponding eigenvalues.

Question 7. *Let V be the space of all real polynomials of degree ≤ 1. It is a vector space over $\mathbb{R}$. Let $S: V \to V$ be defined by*

$$S(p(x)) = 2(x-1)p'(x)$$

for all $p(x) \in V$. Show that S is diagonalizable, and find a diagonalization of S.

Finally, comparing what we had in these two sections, we have:

Theorem 5. *Suppose A is an $n \times n$ matrix with entries in F. Then A is diagonalizable, if and only if the linear map $T: F^n \to F^n$ associated to A is diagonalizable.*

We also have:

Definition 13. *The span of all eigenvectors of T corresponding to an eigenvalue λ of T is called the* eigenspace *of T corresponding to λ.*

Thus any eigenspace of T is in particular a subspace of V. Again in general one should think about eigenspaces instead of eigenvectors as much as possible when one considers linear maps.

3. Invariant subspaces

Again, in this section, $T: V \to V$ is a linear map from V into itself, where V is a vector space over a field F, and $\dim(V) = n < \infty$.

Definition 14. *W is said to be an* invariant subspace *associated to T if it is a subspace of V, and $T(w) \in W$ for all $w \in W$.*

Sometimes we also say that W is a T-*invariant* subspace of V.

Question 8. *For example, let $T: \mathbb{R}^2 \to \mathbb{R}^2$ be given by*

$$T\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2 & -4 \\ 0 & -1 \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix}.$$

Show that the x-axis is a T-invariant subspace of $\mathbb{R}^2$, while the y-axis is not.

If $T: V \to V$ is a linear map, then the following are all trivial T-invariant subspaces of V: $\{0\}$, V, the nullspace of T, and the range of T. Also, any eigenspace of T, i.e. the nullspace of $T - \lambda I$ for any λ, is a T-invariant subspace of V; so is the range of $T - \lambda I$ for any λ.

Definition 15. *If $T\colon V \to V$ is linear and W is a T-invariant subspace of V, one can define the restriction of T to W by*

$$T|_W : W \to W$$
$$T|_W(w) = T(w) \quad \text{for all } w \in W.$$

The restriction of T to W is a linear map from W to W. One important fact is:

Theorem 6. *If W is a T-invariant subspace of V, then the characteristic polynomial of $T|_W$ divides the characteristic polynomial of T; more precisely, if $f(t)$ and $g(t)$ are the characteristic polynomials of T and $T|_W$ respectively, then there is a polynomial $q(t)$ with coefficients in F such that*

$$f(t) = g(t)q(t).$$

Proof. Suppose $\alpha = \{v_1, \dots, v_m\}$ is an ordered basis of W. Extend it to an ordered basis $\beta = \{v_1, \dots, v_n\}$ of V. Then the matrix representation $[T]_\beta$ of T is of the form

$$\begin{pmatrix} A & B \\ 0 & C \end{pmatrix}$$

where A is an $m \times m$ matrix, B is an $m \times (n-m)$ matrix, C is an $(n-m) \times (n-m)$ matrix and 0 is the $(n-m) \times m$ zero matrix. In fact A is then the matrix representation $[T|_W]_\alpha$. It follows that if $f(t)$ and $g(t)$ are the characteristic polynomials of T and $T|_W$ respectively, then

$$f(t) = \det\left(\begin{pmatrix} A & B \\ 0 & C \end{pmatrix} - tI\right) = \det\begin{pmatrix} A - tI & B \\ 0 & C - tI \end{pmatrix},$$

so

$$f(t) = \det(A - tI)\det(C - tI) = g(t)\det(C - tI),$$

and letting $q(t) = \det(C - tI)$ we have the claim. □

Question 9. *For example, if W is the x-axis, and T is the linear map in Question 8, find the characteristic polynomial of $T|_W$, and show that it divides the characteristic polynomial of T.*

One important example of construction of invariant subspaces is the following. Suppose $T\colon V \to V$ is a linear map, where V is a finite dimensional vector space over a field F. Suppose $v \in V$, $v \neq 0$. Let k be the largest positive integer such that $\{v, Tv, T^2v, \dots, T^{k-1}v\}$ are linearly independent[2]. Such k exists since V is finite dimensional. Let W be the span of $\{v, Tv, \dots, T^{k-1}v\}$. Then W is a subspace of V, and we claim that W is T-invariant: in fact if $w \in W$, then

$$w = a_0 v + a_1 Tv + \cdots + a_{k-1}T^{k-1}v$$

for some scalars $a_0, \dots, a_{k-1} \in F$, so

$$Tw = a_0 Tv + a_1 T^2 v + \cdots + a_{k-1}T^k v.$$

[2]From now on, we write T^i for the composition of i copies of T. In other words, T^2v is a shorthand of $T(Tv)$, T^3v is a shorthand for $T(T(Tv))$, etc. More generally, if $f(t)$ is a polynomial in a single variable t, say $f(t) = a_0 + a_1 t + \cdots + a_m t^m$, and $T\colon V \to V$ is a linear map, then $f(T)$ is by definition a linear map from V to V such that $f(T)v = a_0 v + a_1 Tv + \cdots + a_m T^m v$ for all $v \in V$.

But T^kv is a linear combination of $v, Tv, \dots, T^{k-1}v$, since $\{v, Tv, \dots, T^{k-1}v, T^kv\}$ is linearly dependent. Thus we have $Tw \in W$, which proves that W is T-invariant.

Now $\{v, Tv, \dots, T^{k-1}v\}$ is a basis of W (why?), so $\dim(W) = k$. If

$$T^kv = -b_0v - b_1Tv - \cdots - b_{k-1}T^{k-1}v$$

for some scalars $b_0, \dots, b_{k-1}$, then the characteristic polynomial of $T|_W$ can be shown to be

$$g(t) = (-1)^k(b_0 + b_1t + \cdots + b_{k-1}t^{k-1} + t^k)$$

(check!), so in particular $g(T)(v) = (-1)^k(b_0v + b_1Tv + \cdots + b_{k-1}T^{k-1}v + T^kv) = 0$.

Now if $f(t)$ is the characteristic polynomial of T, then $g(t)$ divides $f(t)$, i.e. there exists a polynomial $q(t)$ with coefficients in F such that $f(t) = g(t)q(t)$. It follows that $f(T)(v) = q(T)g(T)v = q(T)(0) = 0$. Since this is true for all $v \in V$, we see that $f(T)$ is the zero map from V to V. This proves:

Theorem 7 (Cayley-Hamilton)**.** *If V is a finite dimensional vector space over a field F, $T\colon V \to V$ is a linear map, and $f(t)$ is the characteristic polynomial of T, then the linear map $f(T)\colon V \to V$ is the zero map, i.e. T 'satisfies' the characteristic polynomial of T.*

Corollary 8. *If A is an $n \times n$ matrix with entries in a field F, and $f(t)$ is the characteristic polynomial of A, then $f(A)$ is the zero matrix.*

4. Criteria about diagonalizability

Again in this section, V is a vector space over a field F of dimension $n < \infty$, and $T\colon V \to V$ is a linear map from V into itself. We state some important criteria about the diagonalizability of T.

Theorem 9. *If $v_1, \dots, v_k \in V$ are eigenvectors of T corresponding to the eigenvalues $\lambda_1, \dots, \lambda_k \in F$, and the eigenvalues $\lambda_1, \dots, \lambda_k$ are pairwise distinct, then the set of vectors $\{v_1, \dots, v_k\}$ is linearly independent.*

Proof. The proof is by induction on k. The claim is trivial if $k = 1$. Now when $k = 2$, i.e. if v_1, v_2 are eigenvectors of T corresponding to eigenvalues λ_1 and λ_2, with $\lambda_1 \neq \lambda_2$, then suppose

$$a_1v_1 + a_2v_2 = 0$$

for some scalars a_1, a_2. Applying $T - \lambda_2 I$ to both sides, we get

$$a_1(\lambda_1 - \lambda_2)v_1 = 0,$$

which implies $a_1 = 0$. Hence $a_2 = 0$. So the case $k = 2$ is proved.

More generally, suppose the claim is true for a certain value of k. Suppose one has $k + 1$ eigenvectors of T, namely $v_1, \dots, v_{k+1}$, corresponding to eigenvalues $\lambda_1, \dots, \lambda_{k+1}$, such that no two of the λ_i's are equal. Then if $a_1, \dots, a_{k+1}$ are scalars such that

$$a_1v_1 + \cdots + a_{k+1}v_{k+1} = 0,$$

by applying $T - \lambda_{k+1}I$ to both sides, we get

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

By induction hypothesis, $\{v_1, \ldots, v_k\}$ are linearly independent, so

$$a_1(\lambda_1 - \lambda_{k+1}) = a_2(\lambda_2 - \lambda_{k+1}) = \cdots = a_k(\lambda_k - \lambda_{k+1}) = 0.$$

Since $\lambda_1, \ldots, \lambda_k$ are all different from λ_{k+1}, this implies

$$a_1 = \cdots = a_k = 0,$$

so $a_{k+1} = 0$ as well, and that completes the induction. □

Corollary 10. *If $T : V \to V$ has n distinct eigenvalues over F, where $n = \dim(V)$, then T is diagonalizable.*

Corollary 11. *Suppose $V_1, \ldots, V_k$ are eigenspaces of T corresponding to eigenvalues $\lambda_1, \ldots, \lambda_k$, and the eigenvalues $\lambda_1, \ldots, \lambda_k$ are pairwise distinct.*

(a) If $v_1 \in V_1$, ..., $v_k \in V_k$ satisfies

$$v_1 + \cdots + v_k = 0,$$

then

$$v_1 = \cdots = v_k = 0.$$

(b) If $\beta_1, \ldots, \beta_k$ are linearly independent subsets of $V_1, \ldots, V_k$ respectively, then $\beta_1 \cup \beta_2 \cup \cdots \cup \beta_k$ is a linearly independent subset of V.

These are easy corollaries of the previous theorem. Their proofs are left as exercises.

Next is a little piece of terminology from algebra:

Definition 16. *If $p(t)$ is a polynomial of one variable with coefficients in a field F and that has degree n, then we say that $p(t)$* splits *over F if there are scalars $c, a_1, \ldots, a_n \in F$ such that*

$$p(t) = c(t - a_1)(t - a_2)\ldots(t - a_n).$$

In other words, $p(t)$ is said to split over F if and only if $p(t)$ can be completely factorized into a product of linear factors with coefficients in F.

Question 10. *Show that the polynomial $t^2 + 1$ splits over $\mathbb{C}$, but not over $\mathbb{R}$.*

Theorem 12. *Suppose V is a finite dimensional vector space over F, and a linear map $T : V \to V$ is diagonalizable. Then the characteristic polynomial of T splits over F.*

Definition 17. *If λ is an eigenvalue of a linear map $T : V \to V$, its* algebraic multiplicity *is the number of times the factor $(t - \lambda)$ appears in the factorization of the characteristic polynomial of T, and its* geometric multiplicity *is the dimension of the eigenspace corresponding to λ.*

Question 11. *Suppose V is the vector space of real polynomials of degree ≤ 2 on $\mathbb{R}$. Let $T\colon V \to V$ be the linear map defined by*

$$T(p(x)) = 2(x-1)p'(x) + (x-1)^2 p''(x).$$

Find the eigenvalues of T, and their algebraic and geometric multiplicities.

From Theorem 12, we have:

Lemma 1. *Suppose $T\colon V \to V$ is a linear map and $\dim(V) = n$. If T is diagonalizable, then the sum of the algebraic multiplicities of the eigenvalues of T is equal to n.*

This is because if T is diagonalizable over F, then by Theorem 12, the characteristic polynomial of T splits, meaning that the characteristic polynomial of T has exactly n zeroes (counting multiplicities) in F. Thus the sum of the algebraic multiplicities of the eigenvalues of T is n, as desired.

The main result of this section is:

Theorem 13. *The algebraic multiplicity of any eigenvalue of T is bigger than or equal to its geometric multiplicity. Furthermore, T is diagonalizable, if and only if equality holds for every eigenvalue of T, i.e. if and only if the algebraic multiplicity of every eigenvalue of T is equal to its geometric multiplicity.*

Proof. If λ is an eigenvalue of T and M, m are its algebraic multiplicity and geometric multiplicities respectively, we first show $m \leq M$. To do so, let W be the eigenspace of T corresponding to λ, and $T|_W$ be the restriction of T to W. Let also $f(t)$ be the characteristic polynomial of T, and $g(t)$ be the characteristic polynomial of $T|_W$. Then by Theorem 6, $g(t)$ divides $f(t)$. But $\dim(W) = m$, and by computing explicitly the characteristic polynomial of $T|_W$, we get

$$g(t) = (\lambda - t)^m.$$

Now M is the number of factors of $(\lambda - t)$ that arises in the factorization of $f(t)$. So from $(\lambda - t)^m$ divides $f(t)$, we conclude that $m \leq M$, which is the desired claim.

Next, suppose the algebraic multiplicity of every eigenvalue of T is equal to its geometric multiplicity. Then let $\lambda_1, \dots, \lambda_k$ be a listing of the distinct eigenvalues of T, and $M_1, \dots, M_k$ be their corresponding algebraic or geometric multiplicities. By Lemma 1, we get

$$M_1 + M_2 + \cdots + M_k = n.$$

Now let $V_1, \dots, V_k$ be the eigenspaces of T corresponding to $\lambda_1, \dots, \lambda_k$. Then their dimensions are $M_1, \dots, M_k$ respectively. Let $\beta_1, \dots, \beta_k$ be a basis of $V_1, \dots, V_k$ respectively. Then $\beta := \beta_1 \cup \cdots \cup \beta_k$ is linearly independent by Corollary 11(b). Furthermore, β has $M_1 + \cdots + M_k = n$ vectors, and all elements of β are all eigenvectors of T. Thus T has n linearly independent eigenvectors, which shows that T is diagonalizable.

Now suppose T is diagonalizable. Then there exists an eigenbasis β of V associated to T. Let $\lambda_1, \dots, \lambda_k$ be a listing of the distinct eigenvalues of T, and $V_1, \dots, V_k$

be the eigenspaces of T corresponding to $\lambda_1, \ldots, \lambda_k$ respectively. For $i = 1, \ldots, k$, let $\beta_i = \beta \cap V_i$, and n_i be the number of elements of β_i. Then β is the disjoint union of $\beta_1, \ldots, \beta_k$, since every vector in β is an eigenvector, and thus belongs to exactly one of the β_i's. It follows that

$$n = \sum_{i=1}^{k} n_i. \tag{1}$$

If now m_i is the geometric multiplicity of λ_i, and M_i be the algebraic multiplicity of λ_i, then

$$n_i \le m_i \le M_i \quad \text{for all } i = 1, \ldots, k. \tag{2}$$

The first inequality holds since β_i is a set of linearly independent vectors in V_i, where $m_i = \dim(V_i)$ and n_i = number of elements of β_i; the second inequality holds by the first part of our theorem. On the other hnad,

$$\sum_{i=1}^{k} M_i = n \tag{3}$$

by Lemma 1. From (1), (2) and (3), we conclude that $m_i = M_i$ for all $i = 1, \ldots, k$. This proves that the algebraic and geometric multiplicities are the same for each eigenvalue of T, and we conclude the proof of the theorem. □

Question 12. *Show that the linear map T in Question 11 is diagonalizable using Theorem 13.*