1 Conjugacy

Definition 1 *Let V and W be isomorphic vector spaces over F. Let $T : V \to V$ and $S : W \to W$ be two linear transformations. We say T and S are* conjugate *(a.k.a.* similar*) and write $T \sim S$ if there is an isomorphism $\phi : V \to W$ such that $S = \phi \circ T \circ \phi^{-1}$, i.e. if the following commutative diagram holds:*

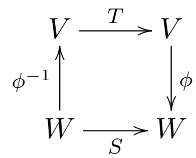

Conjugacy of linear transformations is reflexive ($T \sim T$), symmetric ($S \sim T$ implies $T \sim S$) and transitive ($R \sim S$ and $S \sim T$ imply $R \sim T$).

Definition 2 *Let $A, A' \in M_n(F)$. We say A and A' are* conjugate *(a.k.a.* similar*) if there is an invertible matrix P such that $A' = PAP^{-1}$.*

Conjugacy of matrices is reflexive ($A \sim A$), symmetric ($A \sim B$ implies $B \sim A$) and transitive ($A \sim B$ and $B \sim C$ imply $A \sim C$).

Proposition 3 *Let $A, A' \in M_n(F)$; let $T : F^n \to F^n$ be defined by $T(\mathbf{v}) = A\mathbf{v}$ and let $T' : F^n \to F^n$ be defined by $T'(\mathbf{v}) = A'\mathbf{v}$. Then $A \sim A'$ if and only if $T \sim T'$.*

By definition, if two matrices represent the same linear transformation with respect to different bases of V, then they are conjugate (let P be the change of basis matrix between the two bases). The next result is the converse: it says that if two matrices are conjugate, then they are two matrices representing the same linear transformation, relative to different bases:

Proposition 4 *Let V be a finite-dimensional vector space over F and let $T : V \to V$ be linear. Suppose there is a basis of V such that the matrix of T relative to that basis is $A \in M_n(F)$. Let $A' \in M_n(F)$ be some other matrix. Then $A \sim A'$ if and only if there is a basis of V such that the matrix of T relative to that basis is A'.*

Proposition 5 *If $A' = PAP^{-1}$, then for any $n \in \{0, 1, 2, ...\}$, $(A')^n = PA^nP^{-1}$.*

2 Eigenvalues and eigenvectors

One way of studying a linear transformation is to describe its invariant subspaces. The one-dimensional invariant subspaces of T are associated to eigenvalues and eigenvectors of T; in particular every one-dimensional invariant subspace of T is the span of some eigenvector of T.

Definition 6 *Let $T : V \to V$ be a linear transformation. We say a scalar λ is an* eigenvalue *of T if there exists a nonzero vector $\mathbf{v} \in V$ such that $T(\mathbf{v}) = \lambda\mathbf{v}$. The vector $\mathbf{v}$ is called an* eigenvector *of T corresponding to λ.*

The eigenvector(s) associated to a particular eigenvalue λ are not unique (in particular, any nonzero multiple of an eigenvector associated to λ is an eigenvector associated to λ). More generally:

Definition 7 *Let $T : V \to V$ be a linear transformation. If λ is an eigenvalue of T, the set V_λ of eigenvectors of T corresponding to λ is a subspace of V called the* $\lambda-$eigenspace *of T.*

We can also talk about eigenvalues and eigenvectors of a matrix: given $A \in M_n(F)$ (in particular, A must be square), A defines a linear transformation $T : F^n \to F^n$ by setting $T(\mathbf{x}) = A\mathbf{x}$. The eigenvalues of this T are called the eigenvalues of A (and similarly for eigenvectors and eigenspaces).

Definition 8 *Let $A \in M_n(F)$. We say λ is an* eigenvalue *of A if there exists a nonzero vector $\mathbf{v} \in F^n$ such that $A\mathbf{v} = \lambda\mathbf{v}$. The vector $\mathbf{v}$ is called an* eigenvector *of T corresponding to λ. The set V_λ of eigenvectors of A corresponding to λ is a subspace of F^n called the* $\lambda-$eigenspace *of A.*

Example: Let $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$. Suppose $\mathbf{v} = (v_1, v_2)$ is an eigenvector of A. Then $A\mathbf{v} = (-v_2, v_1) = \lambda(v_1, v_2)$ for some constant λ. Thus we have $-v_2 = \lambda v_1$ and $v_1 = \lambda v_2$ so by substituting the second equation into the first, we obtain $-v_2 = \lambda^2 v_2$, i.e. $\lambda^2 = -1$. So if we require the eigenvalues of A to be real, this A has no eigenvalues and thus no eigenvectors. But if we allow complex eigenvalues and eigenvectors, we obtain (from the equation $\lambda^2 = -1$) $\lambda = i$ and $\lambda = -i$ as eigenvalues. (Whether or not we want to allow complex eigenvalues depends on the context.)

We find eigenvectors associated to $\lambda = i$ as follows: set $A\mathbf{v} = i\mathbf{v}$ and solve for $\mathbf{v} = (v_1, v_2)$ to obtain $-v_2 = iv_1$ and $v_1 = iv_2$. Thus $\mathbf{v} \in V_i$ if and only if $\mathbf{v} \in Span(i, 1)$, i.e. $(i, 1)$ is an eigenvector associated to $\lambda = i$.

Similarly, we find eigenvectors associated to $\lambda = -i$ as follows: set $A\mathbf{v} = -i\mathbf{v}$ and solve for $\mathbf{v} = (v_1, v_2)$ to obtain $-v_2 = -iv_1$ and $v_1 = -iv_2$. Thus $\mathbf{v} \in V_{-i}$ if and only if $\mathbf{v} \in Span(i, -1)$, i.e. $(i, -1)$ is an eigenvector associated to $\lambda = -i$.

Theorem 9 (Properties of eigenvalues) *let $A \in M_n(F)$. Then:*

1. *0 is an eigenvalue of A if and only if A is not invertible;*
2. *$\lambda \in F$ is an eigenvalue of A if and only if $N(A-\lambda I) \neq \{\mathbf{0}\}$ if and only if $\det(A-\lambda I) = 0$;*
3. *if $F = \mathbb{C}$, then A has at least one complex eigenvalue;*
4. *if A is a real matrix and $\lambda \in \mathbb{C}$ is an eigenvalue of A, then so is $\overline{\lambda}$;*
5. *if A is a real matrix and n is odd, then A has a real eigenvalue;*
6. *the eigenvalues of a triangular matrix are its diagonal entries;*
7. *the sum of the eigenvalues of A (counting multiplicities) is the trace of A;*
8. *the product of the eigenvalues of A (counting multiplicities) is the determinant of A;*
9. *if two matrices are conjugate, then they have the same eigenvalues.*

The second statement of the above theorem is of particular importance because it gives a method of computing the eigenvalues of a matrix. Define the *characteristic polynomial* of a square matrix $A \in M_n(F)$ to be $\det(A - \lambda I)$ (think of this as a function of λ). The eigenvalues of A are the roots of the characteristic polynomial. (In the above example, the characteristic polynomial is $\lambda^2 + 1$.)

Remark: When we say "counting multiplicities" in the preceding theorem, we mean the following: suppose the characteristic polynomial of a 3×3 matrix is $-(\lambda - 2)^2(\lambda - 3)$. The eigenvalues of this matrix are therefore $\lambda = 2$ and $\lambda = 3$; we say $\lambda = 2$ has multiplicity 2 since it is a "double root" of the characteristic polynomial. The trace of this matrix would then be the sum of the eigenvalues which is $2 + 2 + 3 = 7$ and similarly the determinant would be $2 \cdot 2 \cdot 3 = 12$ (we count $\lambda = 2$ twice in these calculations since it has multiplicity two). Conceivably, an eigenvalue of an $n \times n$ matrix could have any multiplicity up to n, but *the sum of the multiplicities of the complex eigenvalues of a matrix is always n* (because the sum of the multiplicities of the roots of any degree n polynomial is n, by the Fundamental Theorem of Algebra).

Theorem 10 *Let V be a finite-dimensional vector space over $\mathbb{C}$. Then given any $T : V \to V$, there is an orthonormal basis (with respect to any specified inner product) of V such that the matrix of T relative to that basis is upper-triangular.*

Corollary 11 *Let $A \in M_n(\mathbb{C})$. Then there is a unitary matrix U such that $U^H AU$ is upper triangular.*

Proof: Let $T : \mathbb{C}^n \to \mathbb{C}^n$ be defined by $T(\mathbf{x}) = A\mathbf{x}$. By Theorem 10, there is an orthonormal basis $\mathcal{B}$ of $\mathbb{C}^n$ such that the matrix C of T relative to that basis is upper triangular. Let U be the change of basis matrix from $\mathcal{B}$ to the standard basis $\mathcal{S}$; the columns of U are the vectors of $\mathcal{B}$ which are orthonormal; therefore U is unitary. We have $C = U^{-1}AU = U^H AU$ since U is unitary.

3 Diagonalization

Diagonal matrices are "nice": first, they all commute with each other; second, it is easy to take them to integer powers (just take the power of each element). It is reasonable, therefore, to ask whether or not given a linear transformation $T : V \to V$, we can find a basis of V such that the matrix of T relative to that basis is diagonal. (Theorem 9 says that we can always find a basis such that the matrix is triangular.)

Theorem 12 *Let V be a finite-dimensional vector space over F. Then given any linear $T : V \to V$, the following are equivalent:*

1. *"V has a basis of eigenvectors", i.e. there is a basis of V such that every vector in the basis is an eigenvector of T;*
2. *there is a matrix of T relative to some basis which is diagonal.*

Theorem 13 *Let V be a vector space over F with $\dim V = n < \infty$. Let $T : V \to V$ be linear. Then:*

1. *any set of eigenvectors, each corresponding to a **different** eigenvalue of T, is linearly independent;*
2. *if T has n different eigenvalues, then there is a basis of V such that the matrix of T relative to that basis is diagonal.*

Remark: the converse of the preceding theorem is false. For example, if T has standard matrix I (the identity), then T has only one distinct eigenvalue (namely $\lambda = 1$), but I is diagonal.

We can ask a similar question about matrices: when is a matrix similar to a diagonal matrix?

Definition 14 *Let $A \in M_n(F)$. We say A is* diagonalizable *if A is conjugate to a diagonal matrix, i.e. if there exists an invertible $P \in M_n(F)$ such that $P^{-1}AP$ is a diagonal matrix.*

Theorem 15 *Let $A \in M_n(F)$. Then:*

1. *any set of eigenvectors, each corresponding to a* ***different*** *eigenvalue of A, is linearly independent;*
2. *A is diagonalizable if and only if there is a basis of F^n consisting of eigenvectors of A;*
3. *A is diagonalizable if and only if it has n linearly independent eigenvectors;*
4. *if A has n different eigenvalues, then A is diagonalizable (the converse is false);*
5. *A is diagonalizable if and only if $A = P\Lambda P^{-1}$ where Λ is diagonal, in which case the entries of Λ are the eigenvalues of A and the columns of P are the associated eigenvectors of A.*

Not every matrix is diagonalizable. The prototypical example of a non-diagonalizable matrix is $A = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$, which is not diagonalizable no matter what λ is. The reason is that the only eigenvalue of this matrix is λ, but setting $A\mathbf{x} = \lambda\mathbf{x}$ gives $\mathbf{x} \in Span(1,0)$. So there is only one linearly independent eigenvector of A, hence A is not diagonalizable by the above theorem.

In contrast, consider $A = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}$. Again, the only eigenvalue of this matrix is λ, but setting $A\mathbf{x} = \lambda\mathbf{x}$ gives $\mathbf{x} \in Span((1,0),(0,1))$. So there are two linearly independent eigenvectors of A corresponding to λ (namely $(1,0)$ and $(0,1)$), hence A is diagonalizable since its eigenvectors form a basis of F^2. Of course we knew A was diagonalizable just by looking at it (it is diagonal), but this discussion explains the theory a bit more.

Application: Suppose $A \in M_n(F)$ is diagonalizable. Then we can write $A = P\Lambda P^{-1}$ where Λ is the matrix of eigenvalues of A and the columns of P are the corresponding eigenvectors of A. Therefore for any $n \in \{0,1,2,...\}$, we have $A^n = P\Lambda^n P^{-1}$ (which is easy to figure since Λ is diagonal). Furthermore, we have $e^A = Pe^\Lambda P^{-1}$ (also easy to figure) and $e^{At} = Pe^{\Lambda t}P^{-1}$ (also easy to figure); these formulas are useful for solving differential equations.

3.1 Non-diagonalizable matrices

Inevitably, someone will ask how to compute powers or exponentials of a non-diagonalizable matrix. This is not an impossible task, but it requires a bit more theory. Here is a theorem, related to what is called the "Jordan canonical form" of a matrix (look up Jordan forms on Wikipedia if you want to read about them):

Theorem 16 *Given any $A \in M_n(F)$, we can write $A = D + N$ where $D \in M_n(F)$ is diagonalizable, $N \in M_n(F)$ is* nilpotent *(this means $N^r = 0$ for some $r \geq 0$), and $DN = ND$.*

Having done this, we can write $D = P\Lambda P^{-1}$ where Λ is again the diagonal matrix of eigenvalues of A (these coincide with the eigenvalues of D) and P is the matrix of eigenvectors of D (which have something to do with eigenvectors of A but I won't be specific here). Then, we have:

$$\begin{aligned} A^n &= (D+N)^n \\ &= D^n + nD^{n-1}N + \frac{n(n-1)}{2}D^{n-2}N^2 + ... + nDN^{n-1} + N^n \\ &= P\Lambda^n P^{-1} + nP\Lambda^{n-1}P^{-1}N + \frac{n(n-1)}{2}P\Lambda^{n-2}P^{-1}N^2 + ... + nP\Lambda P^{-1}N^{n-1} + N^n \end{aligned}$$

(notice that in the second line we use the fact that D and N commute). Therefore, if n is big enough, most of these terms are zero (since $N^r = 0$ for some r). In particular, for $n \geq r$ we have

$$\begin{aligned} A^n &= P\Lambda^n P^{-1} + nP\Lambda^{n-1}P^{-1}N + \frac{n(n-1)}{2}P\Lambda^{n-2}P^{-1}N^2 + ... + \frac{n!}{r!(n-r)!}P\Lambda^{r-1}P^{-1}N^{r-1} \\ &= \sum_{j=0}^{r-1} \frac{n!}{j!(n-j)!}P\Lambda^{n-j}P^{-1}N^j \end{aligned}$$

which is (relatively) easy to figure. Exponentials work similarly:

$$e^A = e^{D+N} = e^D e^N \text{ since } D \text{ and } N \text{ commute.}$$

Now e^D can be computed as above since D is diagonalizable; e^N can be computed from the definition, since if $N^r = 0$ we have

$$e^N = I + N + \frac{N^2}{2} + \frac{N^3}{3!} + ... + \frac{N^{r-1}}{(r-1)!}$$

(in particular, this is a finite sum). Then multiply the matrices e^D and e^N to get e^A.