

Contents

5 Eigenvalues and Diagonalization	1
5.1 Eigenvalues, Eigenvectors, and The Characteristic Polynomial	1
5.1.1 Eigenvalues and Eigenvectors	2
5.1.2 Eigenvalues and Eigenvectors of Matrices	3
5.1.3 Eigenspaces	6
5.2 Diagonalization	9
5.3 Applications of Diagonalization	14
5.3.1 Transition Matrices and Incidence Matrices	14
5.3.2 Systems of Linear Differential Equations	16
5.3.3 Non-Diagonalizable Matrices and the Jordan Canonical Form	19

5 Eigenvalues and Diagonalization

In this chapter, we will discuss eigenvalues and eigenvectors: these are “characteristic values” (and “characteristic vectors”) associated to a linear operator $T : V \rightarrow V$ that will allow us to study T in a particularly convenient way. Our ultimate goal is to describe methods for finding a basis for V such that the associated matrix for T has an especially simple form.

We will first describe diagonalization, the procedure for (trying to) find a basis such that the associated matrix for T is a diagonal matrix, and characterize the linear operators that are diagonalizable. Then we will discuss a few applications of diagonalization, including the Cayley-Hamilton theorem that any matrix satisfies its characteristic polynomial, and close with a brief discussion of non-diagonalizable matrices.

5.1 Eigenvalues, Eigenvectors, and The Characteristic Polynomial

- Suppose that we have a linear transformation $T : V \rightarrow V$ from a (finite-dimensional) vector space V to itself. We would like to determine whether there exists a basis β of V such that the associated matrix $[T]_{\beta}^{\beta}$ is a diagonal matrix.
 - Ultimately, our reason for asking this question is that we would like to describe T in as simple a way as possible, and it is unlikely we could hope for anything simpler than a diagonal matrix.
 - So suppose that $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and the diagonal entries of $[T]_{\beta}^{\beta}$ are $\{\lambda_1, \dots, \lambda_n\}$.
 - Then, by assumption, we have $T(\mathbf{v}_i) = \lambda_i \mathbf{v}_i$ for each $1 \leq i \leq n$: the linear transformation T behaves like scalar multiplication by λ_i on the vector \mathbf{v}_i .
 - Conversely, if we were able to find a basis β of V such that $T(\mathbf{v}_i) = \lambda_i \mathbf{v}_i$ for some scalars λ_i , with $1 \leq i \leq n$, then the associated matrix $[T]_{\beta}^{\beta}$ would be a diagonal matrix.
 - This suggests we should study vectors \mathbf{v} such that $T(\mathbf{v}) = \lambda \mathbf{v}$ for some scalar λ .

5.1.1 Eigenvalues and Eigenvectors

- **Definition:** If $T : V \rightarrow V$ is a linear transformation, a nonzero vector \mathbf{v} with $T(\mathbf{v}) = \lambda\mathbf{v}$ is called an eigenvector of T , and the corresponding scalar λ is called an eigenvalue of T .
 - **Important note:** We do not consider the zero vector $\mathbf{0}$ an eigenvector. (The reason for this convention is to ensure that if \mathbf{v} is an eigenvector, then its corresponding eigenvalue λ is unique.)
 - **Terminology notes:** The term “eigenvalue” derives from the German “eigen”, meaning “own” or “characteristic”. The terms characteristic vector and characteristic value are occasionally used in place of “eigenvector” and “eigenvalue”. When V is a vector space of functions, we often use the word eigenfunction in place of “eigenvector”.
- Here are a few examples of linear transformations and eigenvectors:
 - **Example:** If $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the map with $T(x, y) = \langle 2x + 3y, x + 4y \rangle$, then the vector $\mathbf{v} = \langle 3, -1 \rangle$ is an eigenvector of T with eigenvalue 1, since $T(\mathbf{v}) = \langle 3, -1 \rangle = \mathbf{v}$.
 - **Example:** If $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the map with $T(x, y) = \langle 2x + 3y, x + 4y \rangle$, the vector $\mathbf{w} = \langle 1, 1 \rangle$ is an eigenvector of T with eigenvalue 5, since $T(\mathbf{w}) = \langle 5, 5 \rangle = 5\mathbf{w}$.
 - **Example:** If $T : M_{2 \times 2}(\mathbb{R}) \rightarrow M_{2 \times 2}(\mathbb{R})$ is the transpose map, then the matrix $\begin{bmatrix} 1 & 1 \\ 1 & 3 \end{bmatrix}$ is an eigenvector of T with eigenvalue 1.
 - **Example:** If $T : M_{2 \times 2}(\mathbb{R}) \rightarrow M_{2 \times 2}(\mathbb{R})$ is the transpose map, then the matrix $\begin{bmatrix} 0 & -2 \\ 2 & 0 \end{bmatrix}$ is an eigenvector of T with eigenvalue -1 .
 - **Example:** If $T : P(\mathbb{R}) \rightarrow P(\mathbb{R})$ is the map with $T(f(x)) = xf'(x)$, then for any integer $n \geq 0$, the polynomial x^n is an eigenfunction of T with eigenvalue n , since $T(x^n) = x \cdot nx^{n-1} = nx^n$.
 - **Example:** If V is the space of infinitely-differentiable functions and $D : V \rightarrow V$ is the differentiation operator, the function $f(x) = e^{rx}$ is an eigenfunction with eigenvalue r , for any real number r , since $D(e^{rx}) = re^{rx}$.
 - **Example:** If $T : V \rightarrow V$ is any linear transformation and \mathbf{v} is a nonzero vector in $\ker(T)$, then \mathbf{v} is an eigenvector of V with eigenvalue 0. In fact, the eigenvectors with eigenvalue 0 are precisely the nonzero vectors in $\ker(T)$.
- Finding eigenvectors is a generalization of computing the kernel of a linear transformation, but, in fact, we can reduce the problem of finding eigenvectors to that of computing the kernel of a related linear transformation:
- **Proposition (Eigenvalue Criterion):** If $T : V \rightarrow V$ is a linear transformation, the nonzero vector \mathbf{v} is an eigenvector of T with eigenvalue λ if and only if \mathbf{v} is in $\ker(\lambda I - T)$, where I is the identity transformation on V .
 - This criterion reduces the computation of eigenvectors to that of computing the kernel of a collection of linear transformations.
 - **Proof:** Assume $\mathbf{v} \neq \mathbf{0}$. Then \mathbf{v} is an eigenvector of T with eigenvalue $\lambda \iff T(\mathbf{v}) = \lambda\mathbf{v} \iff (\lambda I)\mathbf{v} - T(\mathbf{v}) = \mathbf{0} \iff (\lambda I - T)(\mathbf{v}) = \mathbf{0} \iff \mathbf{v}$ is in the kernel of $\lambda I - T$.
- We will remark that some linear operators may have no eigenvectors at all.
- **Example:** If $I : P(\mathbb{R}) \rightarrow P(\mathbb{R})$ is the integration operator $I(p) = \int_0^x p(t) dt$, show that I has no eigenvectors.
 - Suppose that $I(p) = \lambda p$, so that $\int_0^x p(t) dt = \lambda p(x)$.
 - Then, differentiating both sides with respect to x and applying the fundamental theorem of calculus yields $p(x) = \lambda p'(x)$.
 - If p had positive degree n , then $\lambda p'(x)$ would have degree at most $n - 1$, so it could not equal $p(x)$.
 - Thus, p must be a constant polynomial. But the only constant polynomial with $I(p) = \lambda p$ is the zero polynomial, which is by definition not an eigenvector. Thus, I has no eigenvectors.

- Computing eigenvectors of general linear transformations on infinite-dimensional spaces can be quite difficult.
 - For example, if V is the space of infinitely-differentiable functions, then computing the eigenvectors of the map $T : V \rightarrow V$ with $T(f) = f'' + xf'$ requires solving the differential equation $f'' + xf' = \lambda f$ for an arbitrary λ .
 - It is quite hard to solve that particular differential equation for a general λ (at least, without resorting to using an infinite series expansion to describe the solutions), and the solutions for most values of λ are non-elementary functions.
- In the finite-dimensional case, however, we can recast everything using matrices.
- Proposition: Suppose V is a finite-dimensional vector space with ordered basis β and that $T : V \rightarrow V$ is linear. Then \mathbf{v} is an eigenvector of T with eigenvalue λ if and only if $[\mathbf{v}]_\beta$ is an eigenvector of left-multiplication by $[T]_\beta^\beta$ with eigenvalue λ .
 - Proof: Note that $\mathbf{v} \neq \mathbf{0}$ if and only if $[\mathbf{v}]_\beta \neq \mathbf{0}$, so now assume $\mathbf{v} \neq \mathbf{0}$.
 - Then \mathbf{v} is an eigenvector of T with eigenvalue $\lambda \iff T(\mathbf{v}) = \lambda\mathbf{v} \iff [T(\mathbf{v})]_\beta = [\lambda\mathbf{v}]_\beta \iff [T]_\beta^\beta[\mathbf{v}]_\beta = \lambda[\mathbf{v}]_\beta \iff [\mathbf{v}]_\beta$ is an eigenvector of left-multiplication by $[T]_\beta^\beta$ with eigenvalue λ .

5.1.2 Eigenvalues and Eigenvectors of Matrices

- We will now study eigenvalues and eigenvectors of matrices. For convenience, we restate the definition for this setting:
- Definition: For A an $n \times n$ matrix, a nonzero vector \mathbf{x} with $A\mathbf{x} = \lambda\mathbf{x}$ is called¹ an eigenvector of A , and the corresponding scalar λ is called an eigenvalue of A .

◦ Example: If $A = \begin{bmatrix} 2 & 3 \\ 1 & 4 \end{bmatrix}$, the vector $\mathbf{x} = \begin{bmatrix} 3 \\ -1 \end{bmatrix}$ is an eigenvector of A with eigenvalue 1, because

$$A\mathbf{x} = \begin{bmatrix} 2 & 3 \\ 1 & 4 \end{bmatrix} \begin{bmatrix} 3 \\ -1 \end{bmatrix} = \begin{bmatrix} 3 \\ -1 \end{bmatrix} = \mathbf{x}.$$

◦ Example: If $A = \begin{bmatrix} 2 & -4 & 5 \\ 2 & -2 & 5 \\ 2 & 1 & 2 \end{bmatrix}$, the vector $\mathbf{x} = \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ is an eigenvector of A with eigenvalue 4, because

$$A\mathbf{x} = \begin{bmatrix} 2 & -4 & 5 \\ 2 & -2 & 5 \\ 2 & 1 & 2 \end{bmatrix} \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix} = \begin{bmatrix} 4 \\ 8 \\ 8 \end{bmatrix} = 4\mathbf{x}.$$

- Eigenvalues and eigenvectors can involve complex numbers, even if the matrix A only has real-number entries. We will always work with complex numbers unless specifically indicated otherwise.

◦ Example: If $A = \begin{bmatrix} 6 & 3 & -2 \\ -2 & 0 & 0 \\ 6 & 4 & 2 \end{bmatrix}$, the vector $\mathbf{x} = \begin{bmatrix} 1-i \\ 2i \\ 2 \end{bmatrix}$ is an eigenvector of A with eigenvalue $1+i$, because

$$A\mathbf{x} = \begin{bmatrix} 6 & 3 & -2 \\ -2 & 0 & 0 \\ 6 & 4 & -2 \end{bmatrix} \begin{bmatrix} 1-i \\ 2i \\ 2 \end{bmatrix} = \begin{bmatrix} 2 \\ -2+2i \\ 2+2i \end{bmatrix} = (1+i)\mathbf{x}.$$

- It may at first seem that a given matrix may have many eigenvectors with many different eigenvalues. But in fact, any $n \times n$ matrix can only have a few eigenvalues, and there is a simple way to find them all using determinants:
- Proposition (Computing Eigenvalues): If A is an $n \times n$ matrix, the scalar λ is an eigenvalue of A if and only if $\det(\lambda I - A) = 0$.

¹Technically, such a vector \mathbf{x} is a “right eigenvector” of A : this stands in contrast to a vector \mathbf{y} with $\mathbf{y}A = \lambda\mathbf{y}$, which is called a “left eigenvector” of A . We will only consider right-eigenvectors in our discussion: we do not actually lose anything by ignoring left-eigenvectors, because a left-eigenvector of A is the same as the transpose of a right-eigenvector of A^T .

- Proof: Suppose λ is an eigenvalue with associated nonzero eigenvector \mathbf{x} .
 - Then $A\mathbf{x} = \lambda\mathbf{x}$, or as we observed earlier, $(\lambda I - A)\mathbf{x} = \mathbf{0}$.
 - But from our results on invertible matrices, the matrix equation $(\lambda I - A)\mathbf{x} = \mathbf{0}$ has a nonzero solution for \mathbf{x} if and only if the matrix $\lambda I - A$ is not invertible, which is in turn equivalent to saying that $\det(\lambda I - A) = 0$.
- When we expand the determinant $\det(tI - A)$, we will obtain a polynomial of degree n in the variable t , as can be verified by an easy induction.
 - Definition: For an $n \times n$ matrix A , the degree- n polynomial $p(t) = \det(tI - A)$ is called the characteristic polynomial of A , and its roots are precisely the eigenvalues of A .
 - Some authors instead define the characteristic polynomial as the determinant of the matrix $A - tI$ rather than $tI - A$. We define it this way because then the coefficient of t^n will always be 1, rather than $(-1)^n$.
 - To find the eigenvalues of a matrix, we need only find the roots of its characteristic polynomial.
 - When searching for roots of polynomials of small degree, the following case of the rational root test is often helpful.
 - Proposition: Suppose the polynomial $p(t) = t^n + \dots + b$ has integer coefficients and leading coefficient 1. Then any rational number that is a root of $p(t)$ must be an integer that divides b .
 - The proposition cuts down on the amount of trial and error necessary for finding rational roots of polynomials, since we only need to consider integers that divide the constant term.
 - Of course, a generic polynomial will not have a rational root, so to compute eigenvalues in practice one generally needs to use some kind of numerical approximation procedure to find roots. (But we will arrange the examples so that the polynomials will factor nicely.)
 - Example: Find the eigenvalues of $A = \begin{bmatrix} 3 & 1 \\ 2 & 4 \end{bmatrix}$.
 - First we compute the characteristic polynomial $\det(tI - A) = \begin{vmatrix} t-3 & -1 \\ -2 & t-4 \end{vmatrix} = t^2 - 7t + 10$.
 - The eigenvalues are then the zeroes of this polynomial. Since $t^2 - 7t + 10 = (t-2)(t-5)$ we see that the zeroes are $t = 2$ and $t = 5$, meaning that the eigenvalues are 2 and 5.
 - Example: Find the eigenvalues of $A = \begin{bmatrix} 1 & 4 & \sqrt{3} \\ 0 & 3 & -8 \\ 0 & 0 & \pi \end{bmatrix}$.
 - Observe that $\det(tI - A) = \begin{vmatrix} t-1 & -4 & -\sqrt{3} \\ 0 & t-3 & 8 \\ 0 & 0 & t-\pi \end{vmatrix} = (t-1)(t-3)(t-\pi)$ since the matrix is upper-triangular. Thus, the eigenvalues are 1, 3, π .
 - The idea from the example above works in generality:
 - Proposition (Eigenvalues of Triangular Matrix): The eigenvalues of an upper-triangular or lower-triangular matrix are its diagonal entries.
 - Proof: If A is an $n \times n$ upper-triangular (or lower-triangular) matrix, then so is $tI - A$.
 - Then by properties of determinants, $\det(tI - A)$ is equal to the product of the diagonal entries of $tI - A$.
 - Since these diagonal entries are simply $t - a_{i,i}$ for $1 \leq i \leq n$, the eigenvalues are $a_{i,i}$ for $1 \leq i \leq n$, which are simply the diagonal entries of A .

- It can happen that the characteristic polynomial has a repeated root. In such cases, it is customary to note that the associated eigenvalue has “multiplicity” and include the eigenvalue the appropriate number of extra times when listing them.

- For example, if a matrix has characteristic polynomial $t^2(t-1)^3$, we would say the eigenvalues are 0 with multiplicity 2, and 1 with multiplicity 3. We would list the eigenvalues as $\lambda = 0, 0, 1, 1, 1$.

- Example: Find the eigenvalues of $A = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 3 & 0 \\ 0 & 0 & 0 \end{bmatrix}$.

- By expanding along the bottom row we see $\det(tI - A) = \begin{vmatrix} t-1 & 1 & 0 \\ -1 & t-3 & 0 \\ 0 & 0 & t \end{vmatrix} = t \begin{vmatrix} t-1 & 1 \\ -1 & t-3 \end{vmatrix} = t(t^2 - 4t + 4) = t(t-2)^2$.

- Thus, the characteristic polynomial has a single root $t = 0$ and a double root $t = 2$, so A has an eigenvalue 0 of multiplicity 1 and an eigenvalue 2 of multiplicity 2. As a list, the eigenvalues are $\lambda = \boxed{0, 2, 2}$.

- Example: Find the eigenvalues of $A = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}$.

- Since A is upper-triangular, the eigenvalues are the diagonal entries, so A has an eigenvalue 1 of multiplicity 3. As a list, the eigenvalues are $\lambda = \boxed{1, 1, 1}$.

- Note also that the characteristic polynomial may have non-real numbers as roots, even if the entries of the matrix are real.

- Since the characteristic polynomial will have real coefficients, any non-real eigenvalues will come in complex conjugate pairs. Furthermore, the eigenvectors for these eigenvalues will also necessarily contain non-real entries.

- Example: Find the eigenvalues of $A = \begin{bmatrix} 1 & 1 \\ -2 & 3 \end{bmatrix}$.

- First we compute the characteristic polynomial $\det(tI - A) = \begin{vmatrix} t-1 & -1 \\ 2 & t-3 \end{vmatrix} = t^2 - 4t + 5$.

- The eigenvalues are then the zeroes of this polynomial. By the quadratic formula, the roots are $\frac{4 \pm \sqrt{-4}}{2} = 2 \pm i$, so the eigenvalues are $\boxed{2 + i, 2 - i}$.

- Example: Find the eigenvalues of $A = \begin{bmatrix} -1 & 2 & -4 \\ 3 & -2 & 1 \\ 4 & -4 & 4 \end{bmatrix}$.

- By expanding along the top row,

$$\begin{aligned} \det(tI - A) &= \begin{vmatrix} t+1 & -2 & 4 \\ -3 & t+2 & -1 \\ -4 & 4 & t-4 \end{vmatrix} \\ &= (t+1) \begin{vmatrix} t+2 & -1 \\ 4 & t-4 \end{vmatrix} + 2 \begin{vmatrix} -3 & -1 \\ -4 & t-4 \end{vmatrix} + 4 \begin{vmatrix} -3 & t+2 \\ -4 & 4 \end{vmatrix} \\ &= (t+1)(t^2 - 2t - 4) + 2(-3t + 8) + 4(4t - 4) \\ &= t^3 - t^2 + 4t - 4. \end{aligned}$$

- To find the roots, we wish to solve the cubic equation $t^3 - t^2 + 4t - 4 = 0$.
- By the rational root test, if the polynomial has a rational root then it must be an integer dividing -4 : that is, one of $\pm 1, \pm 2, \pm 4$. Testing the possibilities reveals that $t = 1$ is a root, and then we get the factorization $(t-1)(t^2 + 4) = 0$.
- The roots of the quadratic are $t = \pm 2i$, so the eigenvalues are $\boxed{1, 2i, -2i}$.

5.1.3 Eigenspaces

- Using the characteristic polynomial, we can find all the eigenvalues of a matrix A without actually determining the associated eigenvectors. However, we often also want to find the eigenvectors associated to each eigenvalue.
- We might hope that there is a straightforward way to describe all the eigenvectors, and (conveniently) there is: the set of all eigenvectors with a particular eigenvalue λ has a vector space structure.
- **Proposition** (Eigenspaces): If $T : V \rightarrow V$ is linear, then for any fixed value of λ , the set E_λ of vectors in V satisfying $T(\mathbf{v}) = \lambda\mathbf{v}$ is a subspace of V . This space E_λ is called the eigenspace associated to the eigenvalue λ , or more simply the λ -eigenspace.

- Notice that E_λ is precisely the set of eigenvectors with eigenvalue λ , along with the zero vector.
- The eigenspaces for a matrix A are defined in the same way: E_λ is the space of vectors \mathbf{v} such that $A\mathbf{v} = \lambda\mathbf{v}$.
- **Proof:** By definition, E_λ is the kernel of the linear transformation $\lambda I - T$, and is therefore a subspace of V .

- **Example:** Find the 1-eigenspaces, and their dimensions, for $A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$.

- For the 1-eigenspace of A , we want to find all vectors with $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$.

- Clearly, all vectors satisfy this equation, so the 1-eigenspace of A is the set of all vectors $\begin{bmatrix} a \\ b \end{bmatrix}$, and has dimension 2.

- For the 1-eigenspace of B , we want to find all vectors with $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$, or equivalently, $\begin{bmatrix} a + b \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$.

- The vectors satisfying the equation are those with $b = 0$, so the 1-eigenspace of B is the set of vectors of the form $\begin{bmatrix} a \\ 0 \end{bmatrix}$, and has dimension 1.

- Notice that the characteristic polynomial of each matrix is $(t - 1)^2$, since both matrices are upper-triangular, and they both have a single eigenvalue $\lambda = 1$ of multiplicity 2. Nonetheless, the matrices do not have the same eigenvectors, and the dimensions of their 1-eigenspaces are different.

- In the finite-dimensional case, we would like to compute a basis for the λ -eigenspace: this is equivalent to solving the system $(\lambda I - A)\mathbf{v} = \mathbf{0}$, which we can do by row-reducing the matrix $\lambda I - A$.

- **Example:** Find all eigenvalues, and a basis for each eigenspace, for the matrix $A = \begin{bmatrix} 2 & 2 \\ 3 & 1 \end{bmatrix}$.

- We have $tI - A = \begin{bmatrix} t - 2 & -2 \\ -3 & t - 1 \end{bmatrix}$, so $p(t) = \det(tI - A) = (t - 2)(t - 1) - (-2)(-3) = t^2 - 3t - 4$.

- Since $p(t) = t^2 - 3t - 4 = (t - 4)(t + 1)$, the eigenvalues are $\lambda = -1, 4$.

- For $\lambda = -1$, we want to find the nullspace of $\begin{bmatrix} -1 - 2 & -2 \\ -3 & -1 - 1 \end{bmatrix} = \begin{bmatrix} -3 & -2 \\ -3 & -2 \end{bmatrix}$. By row-reducing we find the row-echelon form is $\begin{bmatrix} -3 & -2 \\ 0 & 0 \end{bmatrix}$, so the (-1) -eigenspace is 1-dimensional and is spanned by

$$\begin{bmatrix} -2 \\ 3 \end{bmatrix}.$$

- For $\lambda = 4$, we want to find the nullspace of $\begin{bmatrix} 4-2 & -2 \\ -3 & 4-1 \end{bmatrix} = \begin{bmatrix} 2 & -2 \\ -3 & 3 \end{bmatrix}$. By row-reducing we find the row-echelon form is $\begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix}$, so the 4-eigenspace is 1-dimensional and is spanned by $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

- Example: Find all eigenvalues, and a basis for each eigenspace, for the matrix $A = \begin{bmatrix} 1 & 0 & 1 \\ -1 & 1 & 3 \\ -1 & 0 & 3 \end{bmatrix}$.

- First, we have $tI - A = \begin{bmatrix} t-1 & 0 & -1 \\ 1 & t-1 & -3 \\ 1 & 0 & t-3 \end{bmatrix}$, so $p(t) = (t-1) \cdot \begin{vmatrix} t-1 & -3 \\ 0 & t-3 \end{vmatrix} + (-1) \cdot \begin{vmatrix} 1 & t-1 \\ 1 & 0 \end{vmatrix} = (t-1)^2(t-3) + (t-1)$.

- Since $p(t) = (t-1) \cdot [(t-1)(t-3) + 1] = (t-1)(t-2)^2$, the eigenvalues are $\lambda = 1, 2, 2$.

- For $\lambda = 1$ we want to find the nullspace of $\begin{bmatrix} 1-1 & 0 & -1 \\ 1 & 1-1 & -3 \\ 1 & 0 & 1-3 \end{bmatrix} = \begin{bmatrix} 0 & 0 & -1 \\ 1 & 0 & -3 \\ 1 & 0 & -3 \end{bmatrix}$. This matrix's

reduced row-echelon form is $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$, so the 1-eigenspace is 1-dimensional and spanned by $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$.

- For $\lambda = 2$ we want to find the nullspace of $\begin{bmatrix} 2-1 & 0 & -1 \\ 1 & 2-1 & -3 \\ 1 & 0 & 2-3 \end{bmatrix} = \begin{bmatrix} 1 & 0 & -1 \\ 1 & 1 & -3 \\ 1 & 0 & -1 \end{bmatrix}$. This matrix's re-

duced row-echelon form is $\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -2 \\ 0 & 0 & 0 \end{bmatrix}$, so the 2-eigenspace is 1-dimensional and spanned by $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$.

- Example: Find all eigenvalues, and a basis for each eigenspace, for the matrix $A = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}$.

- We have $tI - A = \begin{bmatrix} t & 0 & 0 \\ -1 & t & 1 \\ 0 & -1 & t \end{bmatrix}$, so $p(t) = \det(tI - A) = t \cdot \begin{vmatrix} t & 1 \\ -1 & t \end{vmatrix} = t \cdot (t^2 + 1)$.

- Since $p(t) = t \cdot (t^2 + 1)$, the eigenvalues are $\lambda = 0, i, -i$.

- For $\lambda = 0$ we want to find the nullspace of $\begin{bmatrix} 0 & 0 & 0 \\ -1 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}$. This matrix's reduced row-echelon form is

$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$, so the 0-eigenspace is 1-dimensional and spanned by $\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$.

- For $\lambda = i$ we want to find the nullspace of $\begin{bmatrix} i & 0 & 0 \\ -1 & i & 1 \\ 0 & -1 & i \end{bmatrix}$. This matrix's reduced row-echelon form is

$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -i \\ 0 & 0 & 0 \end{bmatrix}$, so the i -eigenspace is 1-dimensional and spanned by $\begin{bmatrix} 0 \\ i \\ 1 \end{bmatrix}$.

- For $\lambda = -i$ we want to find the nullspace of $\begin{bmatrix} -i & 0 & 0 \\ -1 & -i & 1 \\ 0 & -1 & -i \end{bmatrix}$. This matrix's reduced row-echelon form

is $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & i \\ 0 & 0 & 0 \end{bmatrix}$, so the $(-i)$ -eigenspace is 1-dimensional and spanned by $\begin{bmatrix} 0 \\ -i \\ 1 \end{bmatrix}$.

- Notice that in the example above, with a real matrix having complex-conjugate eigenvalues, the associated eigenvectors were also complex conjugates. This is no accident:

- **Proposition** (Conjugate Eigenvalues): If A is a real matrix and \mathbf{v} is an eigenvector with a complex eigenvalue λ , then the complex conjugate $\bar{\mathbf{v}}$ is an eigenvector with eigenvalue $\bar{\lambda}$. In particular, a basis for the $\bar{\lambda}$ -eigenspace is given by the complex conjugate of a basis for the λ -eigenspace.

- **Proof:** The first statement follows from the observation that the complex conjugate of a product or sum is the appropriate product or sum of complex conjugates, so if A and B are any matrices of compatible sizes for multiplication, we have $\overline{A \cdot B} = \bar{A} \cdot \bar{B}$.
- Thus, if $A\mathbf{v} = \lambda\mathbf{v}$, taking complex conjugates gives $\bar{A}\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$, and since $\bar{A} = A$ because A is a real matrix, we see $A\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$: thus, $\bar{\mathbf{v}}$ is an eigenvector with eigenvalue $\bar{\lambda}$.
- The second statement follows from the first, since complex conjugation does not affect linear independence or dimension.

- **Example:** Find all eigenvalues, and a basis for each eigenspace, for the matrix $A = \begin{bmatrix} 3 & -1 \\ 2 & 5 \end{bmatrix}$.

- We have $tI - A = \begin{bmatrix} t-3 & 1 \\ -2 & t-5 \end{bmatrix}$, so $p(t) = \det(tI - A) = (t-3)(t-5) - (-2)(1) = t^2 - 8t + 17$, so the eigenvalues are $\lambda = 4 \pm i$.

- For $\lambda = 4 + i$, we want to find the nullspace of $\begin{bmatrix} t-3 & 1 \\ -2 & t-5 \end{bmatrix} = \begin{bmatrix} 1+i & 1 \\ -2 & -1+i \end{bmatrix}$. Row-reducing this matrix yields

$$\begin{bmatrix} 1+i & 1 \\ -2 & -1+i \end{bmatrix} \xrightarrow{R_2+(1-i)R_1} \begin{bmatrix} 1+i & 1 \\ 0 & 0 \end{bmatrix}$$

from which we can see that the $(4+i)$ -eigenspace is 1-dimensional and spanned by $\begin{bmatrix} 1 \\ -1-i \end{bmatrix}$.

- For $\lambda = 4 - i$ we can simply take the conjugate of the calculation we made for $\lambda = 4 + i$: thus, the $(4-i)$ -eigenspace is also 1-dimensional and spanned by $\begin{bmatrix} 1 \\ -1+i \end{bmatrix}$.

- We will mention one more result about eigenvalues that can be useful in double-checking calculations:

- **Theorem** (Eigenvalues, Trace, and Determinant): The product of the eigenvalues of A is the determinant of A , and the sum of the eigenvalues of A equals the trace of A .

- Recall that the trace of a matrix is defined to be the sum of its diagonal entries.
- **Proof:** Let $p(t)$ be the characteristic polynomial of A .
- If we expand out the product $p(t) = (t - \lambda_1) \cdot (t - \lambda_2) \cdots (t - \lambda_n)$, we see that the constant term is equal to $(-1)^n \lambda_1 \lambda_2 \cdots \lambda_n$.
- But the constant term is also just $p(0)$, and since $p(t) = \det(tI - A)$ we have $p(0) = \det(-A) = (-1)^n \det(A)$: thus, $\lambda_1 \lambda_2 \cdots \lambda_n = \det(A)$.
- Furthermore, upon expanding out the product $p(t) = (t - \lambda_1) \cdot (t - \lambda_2) \cdots (t - \lambda_n)$, we see that the coefficient of t^{n-1} is equal to $-(\lambda_1 + \cdots + \lambda_n)$.
- If we expand out the determinant $\det(tI - A)$ to find the coefficient of t^{n-1} , we can show (with a little bit of effort) that the coefficient is the negative of the sum of the diagonal entries of A .
- Thus, setting the two expressions equal shows that the sum of the eigenvalues equals the trace of A .

- Example: Find the eigenvalues of the matrix $A = \begin{bmatrix} 2 & 1 & 1 \\ -2 & -1 & -2 \\ 2 & 2 & -3 \end{bmatrix}$, and verify the formulas for trace and determinant in terms of the eigenvalues.

- By expanding along the top row, we can compute

$$\begin{aligned} \det(tI - A) &= (t-2) \begin{vmatrix} t+1 & 2 \\ -2 & t+3 \end{vmatrix} - (-1) \begin{vmatrix} 2 & 2 \\ -2 & t+3 \end{vmatrix} + (-1) \begin{vmatrix} 2 & t+1 \\ -2 & -2 \end{vmatrix} \\ &= (t-2)(t^2 + 4t + 7) + (2t + 10) - (2t - 2) = t^3 + 2t^2 - t - 2. \end{aligned}$$

- To find the eigenvalues, we wish to solve the cubic equation $t^3 + 2t^2 - t - 2 = 0$.
 - By the rational root test, if the polynomial has a rational root then it must be an integer dividing -2 : that is, one of $\pm 1, \pm 2$. Testing the possibilities reveals that $t = 1, t = -1$, and $t = -2$ are each roots, from which we obtain the factorization $(t-1)(t+1)(t+2) = 0$.
 - Thus, the eigenvalues are $t = -2, -1, 1$.
 - We see that $\text{tr}(A) = 2 + (-1) + (-3) = -2$, while the sum of the eigenvalues is $(-2) + (-1) + 1 = -2$.
 - Also, $\det(A) = 2$, and the product of the eigenvalues is $(-2)(-1)(1) = 2$.
- In all of the examples above, the dimension of each eigenspace was less than or equal to the multiplicity of the eigenvalue as a root of the characteristic polynomial. This is true in general:
 - Theorem (Eigenvalue Multiplicity): If λ is an eigenvalue of the matrix A which appears exactly k times as a root of the characteristic polynomial, then the dimension of the eigenspace corresponding to λ is at least 1 and at most k .

- Remark: The number of times that λ appears as a root of the characteristic polynomial is sometimes called the “algebraic multiplicity” of λ , and the dimension of the eigenspace corresponding to λ is sometimes called the “geometric multiplicity” of λ . In this language, the theorem above says that the geometric multiplicity is less than or equal to the algebraic multiplicity.
- Example: If the characteristic polynomial of a matrix is $(t-1)^3(t-3)^2$, then the eigenspace for $\lambda = 1$ is at most 3-dimensional, and the eigenspace for $\lambda = 3$ is at most 2-dimensional.
- Proof: The statement that the eigenspace has dimension at least 1 is immediate, because (by assumption) λ is a root of the characteristic polynomial and therefore has at least one nonzero eigenvector associated to it.
- For the other statement, observe that the dimension of the λ -eigenspace is the dimension of the solution space of the homogeneous system $(\lambda I - A)\mathbf{x} = \mathbf{0}$. (Equivalently, it is the dimension of the nullspace of the matrix $\lambda I - A$.)
- If λ appears k times as a root of the characteristic polynomial, then when we put the matrix $\lambda I - A$ into its reduced row-echelon form B , we claim that B must have at most k rows of all zeroes.
- Otherwise, the matrix B (and hence $\lambda I - A$ too, since the nullity and rank of a matrix are not changed by row operations) would have 0 as an eigenvalue more than k times, because B is in echelon form and therefore upper-triangular.
- But the number of rows of all zeroes in a square matrix in reduced row-echelon form is the same as the number of nonpivotal columns, which is the number of free variables, which is the dimension of the solution space.
- So, putting all the statements together, we see that the dimension of the eigenspace is at most k .

5.2 Diagonalization

- Let us now return to our original question that motivated our discussion of eigenvalues and eigenvectors in the first place: given a linear operator $T : V \rightarrow V$ on a vector space V , can we find a basis β of V such that the associated matrix $[T]_{\beta}^{\beta}$ is a diagonal matrix?

- **Definition:** A linear operator $T : V \rightarrow V$ on a finite-dimensional vector space V is diagonalizable if there exists a basis β of V such that the associated matrix $[T]_{\beta}^{\beta}$ is a diagonal matrix.
 - We can also formulate essentially the same definition for matrices: if A is an $n \times n$ matrix, then A is the associated matrix of the linear transformation T given by left-multiplication by A .
 - We then would like to say that A is diagonalizable when T is diagonalizable.
 - By our results on change of basis, this is equivalent to saying that there exists an invertible matrix Q , namely the change-of-basis matrix $Q = [I]_{\gamma}^{\beta}$, for which $Q^{-1}AQ = [I]_{\gamma}^{\beta}[T]_{\gamma}^{\gamma}[I]_{\beta}^{\gamma} = [T]_{\beta}^{\beta}$ is a diagonal matrix.
- **Definition:** An $n \times n$ matrix A is diagonalizable if there exists an invertible $n \times n$ matrix Q for which $Q^{-1}AQ$ is a diagonal matrix.
 - Recall that we say two $n \times n$ matrices A and B are similar if there exists an invertible $n \times n$ matrix Q such that $B = Q^{-1}AQ$.
- Our goal is to study and then characterize diagonalizable linear transformations, which (per the above discussion) is equivalent to characterizing diagonalizable matrices.
- **Proposition** (Characteristic Polynomials and Similarity): If A and B are similar, then they have the same characteristic polynomial, determinant, trace, and eigenvalues (and their eigenvalues have the same multiplicities).
 - **Proof:** Suppose $B = Q^{-1}AQ$. For the characteristic polynomial, we simply compute $\det(tI - B) = \det(Q^{-1}(tI)Q - Q^{-1}AQ) = \det(Q^{-1}(tI - A)Q) = \det(Q^{-1}) \det(tI - A) \det(Q) = \det(tI - A)$.
 - The determinant and trace are both coefficients (up to a factor of ± 1) of the characteristic polynomial, so they are also equal.
 - Finally, the eigenvalues are the roots of the characteristic polynomial, so they are the same and occur with the same multiplicities for A and B .
- The eigenvectors for similar matrices are also closely related:
- **Proposition** (Eigenvectors and Similarity): If $B = Q^{-1}AQ$, then \mathbf{v} is an eigenvector of B with eigenvalue λ if and only if $Q\mathbf{v}$ is an eigenvector of A with eigenvalue λ .
 - **Proof:** Since Q is invertible, $\mathbf{v} = \mathbf{0}$ if and only if $Q\mathbf{v} = \mathbf{0}$. Now assume $\mathbf{v} \neq \mathbf{0}$.
 - First suppose \mathbf{v} is an eigenvector of B with eigenvalue λ . Then $A(Q\mathbf{v}) = Q(Q^{-1}AQ)\mathbf{v} = Q(B\mathbf{v}) = Q(\lambda\mathbf{v}) = \lambda(Q\mathbf{v})$, meaning that $Q\mathbf{v}$ is an eigenvector of A with eigenvalue λ .
 - Conversely, if $Q\mathbf{v}$ is an eigenvector of A with eigenvalue λ . Then $B\mathbf{v} = Q^{-1}A(Q\mathbf{v}) = Q^{-1}\lambda(Q\mathbf{v}) = \lambda(Q^{-1}Q\mathbf{v}) = \lambda\mathbf{v}$, so \mathbf{v} is an eigenvector of B with eigenvalue λ .
- **Corollary:** If $B = Q^{-1}AQ$, then the eigenspaces for B have the same dimensions as the eigenspaces for A .
- As we have essentially worked out already, diagonalizability is equivalent to the existence of a basis of eigenvectors:
- **Theorem** (Diagonalizability): A linear operator $T : V \rightarrow V$ is diagonalizable if and only if there exists a basis β of V consisting of eigenvectors of T .
 - **Proof:** First suppose that V has a basis of eigenvectors $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ with respective eigenvalues $\lambda_1, \dots, \lambda_n$. Then by hypothesis, $T(\mathbf{v}_i) = \lambda_i\mathbf{v}_i$, and so $[T]_{\beta}^{\beta}$ is the diagonal matrix with diagonal entries $\lambda_1, \dots, \lambda_n$.
 - Conversely, suppose T is diagonalizable and let $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be a basis such that $[T]_{\beta}^{\beta}$ is a diagonal matrix whose diagonal entries are $\lambda_1, \dots, \lambda_n$. Then by hypothesis, each \mathbf{v}_i is nonzero and $T(\mathbf{v}_i) = \lambda_i\mathbf{v}_i$, so each \mathbf{v}_i is an eigenvector of T .
- Although the result above does give a characterization of diagonalizable matrices, it is not entirely obvious how to determine whether a basis of eigenvectors exists.

- It turns out that we can essentially check this property on each eigenspace.
 - As we already proved, the dimension of the λ -eigenspace of A is less than or equal to the multiplicity of λ as a root of the characteristic polynomial.
 - But since the characteristic polynomial has degree n , that means the sum of the dimensions of the λ -eigenspaces is at most n , and can equal n only when each eigenspace has dimension equal to the multiplicity of its corresponding eigenvalue.
 - Our goal is to show that the converse holds as well: if each eigenspace has the proper dimension, then the matrix will be diagonalizable.
- We first need an intermediate result about linear independence of eigenvectors having distinct eigenvalues:
 - **Theorem** (Independent Eigenvectors): If $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ are eigenvectors of T associated to distinct eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$, then $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ are linearly independent.
 - **Proof:** We induct on n .
 - The base case $n = 1$ is trivial, since by definition an eigenvector cannot be the zero vector.
 - Now suppose $n \geq 2$ and that we had a linear dependence $a_1\mathbf{v}_1 + \dots + a_n\mathbf{v}_n = \mathbf{0}$ for eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ having distinct eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$,
 - Applying T to both sides yields $\mathbf{0} = T(\mathbf{0}) = T(a_1\mathbf{v}_1 + \dots + a_n\mathbf{v}_n) = a_1(\lambda_1\mathbf{v}_1) + \dots + a_n(\lambda_n\mathbf{v}_n)$.
 - But now if we scale the original dependence by λ_1 and subtract this new relation (to eliminate \mathbf{v}_1), we obtain $a_2(\lambda_2 - \lambda_1)\mathbf{v}_2 + a_3(\lambda_3 - \lambda_1)\mathbf{v}_3 + \dots + a_n(\lambda_n - \lambda_1)\mathbf{v}_n = \mathbf{0}$.
 - By the inductive hypothesis, all coefficients of this dependence must be zero, and so since $\lambda_k \neq \lambda_1$ for each k , we conclude that $a_2 = \dots = a_n = 0$. Then $a_1\mathbf{v}_1 = \mathbf{0}$ implies $a_1 = 0$ also, so we are done.
 - **Theorem** (Diagonalizability Criterion): An $n \times n$ matrix is diagonalizable (over the complex numbers) if and only if for each eigenvalue λ , the dimension of the λ -eigenspace is equal to the multiplicity of λ as a root of the characteristic polynomial.
 - **Proof:** If the $n \times n$ matrix A is diagonalizable, then by our previous theorem on diagonalizability, V has a basis β of eigenvectors for A .
 - For any eigenvalue λ_i of A , let b_i be the number of elements of β having eigenvalue λ_i , and let d_i be the multiplicity of λ_i as a root of the characteristic polynomial.
 - Then $\sum_i b_i = n$ since β is a basis of V , and $\sum_i d_i = n$ by our results about the characteristic polynomial, and $b_i \leq d_i$ as we proved before. Thus, $n = \sum_i b_i \leq \sum_i d_i = n$, so $n_i = d_i$ for each i .
 - For the other direction, suppose that all eigenvalues of A lie in the scalar field of V , and that $b_i = d_i$ for all i . Then let β be the union of bases for each eigenspace of A : by hypothesis, β contains $\sum_i b_i = \sum_i d_i = n$ vectors, so to conclude it is a basis of the n -dimensional vector space V , we need only show that it is linearly independent.
 - Explicitly, let $\beta_i = \{\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,j_i}\}$ be a basis of the λ_i -eigenspace for each i , so that $\beta = \{\mathbf{v}_{1,1}, \mathbf{v}_{1,2}, \dots, \mathbf{v}_{k,j}\}$ and $A\mathbf{v}_{i,j} = \lambda_i\mathbf{v}_{i,j}$ for each pair (i, j) .
 - Suppose we have a dependence $a_{1,1}\mathbf{v}_{1,1} + \dots + a_{k,j}\mathbf{v}_{k,j} = \mathbf{0}$. Let $\mathbf{w}_i = \sum_j a_{i,j}\mathbf{v}_{i,j}$, and observe that \mathbf{w}_i has $A\mathbf{w}_i = \lambda_i\mathbf{w}_i$, and that $\mathbf{w}_1 + \mathbf{w}_2 + \dots + \mathbf{w}_k = \mathbf{0}$.
 - If any of the \mathbf{w}_i were nonzero, then we would have a nontrivial linear dependence between eigenvectors of A having distinct eigenvalues, which is impossible by the previous theorem.
 - Therefore, each $\mathbf{w}_i = \mathbf{0}$, meaning that $a_{i,1}\mathbf{v}_{i,1} + \dots + a_{i,j_i}\mathbf{v}_{i,j_i} = \mathbf{0}$. But then since β_i is linearly independent, all of the coefficients $a_{i,j}$ must be zero. Thus, β is linearly independent and therefore is a basis for V .
 - **Corollary:** If A is an $n \times n$ matrix with n distinct eigenvalues, then A is diagonalizable.
 - **Proof:** Every eigenvalue must occur with multiplicity 1 as a root of the characteristic polynomial, since there are n eigenvalues and the sum of their multiplicities is also n .
 - Then the dimension of each eigenspace is equal to 1 (since it is always between 1 and the multiplicity), so by the theorem above, A is diagonalizable.

- The proof of the diagonalizability theorem gives an explicit procedure for determining both diagonalizability and the diagonalizing matrix. To determine whether a linear transformation T (or matrix A) is diagonalizable, and if so how to find a basis β such that $[T]_\beta^\beta$ is diagonal (or a matrix Q with $Q^{-1}AQ$ diagonal), follow these steps:

- Step 1: Find the characteristic polynomial and eigenvalues of T (or A).
- Step 2: Find a basis for each eigenspace of T (or A).
- Step 3a: Determine whether T (or A) is diagonalizable. If each eigenspace is “nondefective” (i.e., its dimension is equal to the number of times the corresponding eigenvalue appears as a root of the characteristic polynomial) then T is diagonalizable, and otherwise, T is not diagonalizable.
- Step 3b: For a diagonalizable linear transformation T , take β to be a basis of eigenvectors for T . For a diagonalizable matrix A , the diagonalizing matrix Q can be taken to be the matrix whose columns are a basis of eigenvectors of A .

- Example: For $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $T(x, y) = \langle -2y, 3x + 5y \rangle$, determine whether T is diagonalizable and if so, find a basis β such that $[T]_\beta^\beta$ is diagonal.

- The associated matrix A for T relative to the standard basis is $A = \begin{bmatrix} 0 & -2 \\ 3 & 5 \end{bmatrix}$.
- For the characteristic polynomial, we compute $\det(tI - A) = t^2 - 5t + 6 = (t - 2)(t - 3)$, so the eigenvalues are therefore $\lambda = 2, 3$. Since the eigenvalues are distinct we know that T is diagonalizable.
- A short calculation yields that $\langle 1, -1 \rangle$ is a basis for the 2-eigenspace, and that $\langle -2, 3 \rangle$ is a basis for the 3-eigenspace.
- Thus, for $\beta = \boxed{\{\langle 1, -1 \rangle, \langle -2, 3 \rangle\}}$, we can see that $[T]_\beta^\beta = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}$ is diagonal.

- Example: For $A = \begin{bmatrix} 1 & -1 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix}$, determine whether there exists a diagonal matrix D and an invertible matrix Q with $D = Q^{-1}AQ$, and if so, find them.

- We compute $\det(tI - A) = (t - 1)^3$ since $tI - A$ is upper-triangular, and the eigenvalues are $\lambda = 1, 1, 1$.
- The 1-eigenspace is then the nullspace of $I - A = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$, which (since the matrix is already in row-echelon form) is 1-dimensional and spanned by $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$.
- Since the eigenspace for $\lambda = 1$ is 1-dimensional but the eigenvalue appears 3 times as a root of the characteristic polynomial, the matrix A is not diagonalizable and there is no such Q .

- Example: For $A = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 2 & 0 \\ 0 & 2 & 1 \end{bmatrix}$, determine whether there exists a diagonal matrix D and an invertible matrix Q with $D = Q^{-1}AQ$, and if so, find them.

- We compute $\det(tI - A) = (t - 1)^2(t - 2)$, so the eigenvalues are $\lambda = 1, 1, 2$.
- A short calculation yields that $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$ is a basis for the 1-eigenspace and that $\begin{bmatrix} -1 \\ 1 \\ 2 \end{bmatrix}$ is a basis for the 2-eigenspace.
- Since the eigenspaces both have the proper dimensions, A is diagonalizable, and we can take $D = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix}$ with $Q = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 1 \\ 0 & 1 & 2 \end{bmatrix}$.

◦ To check: we have $Q^{-1} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & -2 & 1 \\ 0 & 1 & 0 \end{bmatrix}$, so $Q^{-1}AQ = \begin{bmatrix} 1 & 1 & 0 \\ 0 & -2 & 1 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & -1 & 0 \\ 0 & 2 & 0 \\ 0 & 2 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 1 \\ 0 & 1 & 2 \end{bmatrix} =$

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix} = D.$$

◦ Remark: We could (for example) also take $D = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$ if we wanted, and the associated conjugating matrix could have been $Q = \begin{bmatrix} -1 & 1 & 0 \\ 1 & 0 & 0 \\ 2 & 0 & 1 \end{bmatrix}$ instead. There is no particular reason to care much about which diagonal matrix we want as long as we make sure to arrange the eigenvectors in the correct order. We could also have used any other bases for the eigenspaces to construct Q .

- Knowing that a matrix is diagonalizable can be very computationally useful.
 - For example, if A is diagonalizable with $D = Q^{-1}AQ$, then it is very easy to compute any power of A .
 - Explicitly, since we can rearrange to write $A = QDQ^{-1}$, then $A^k = (QDQ^{-1})^k = Q(D^k)Q^{-1}$, since the conjugate of the k th power is the k th power of a conjugate.
 - But since D is diagonal, D^k is simply the diagonal matrix whose diagonal entries are the k th powers of the diagonal entries of D .

• Example: If $A = \begin{bmatrix} -2 & -6 \\ 3 & 7 \end{bmatrix}$, find a formula for the k th power A^k , for k a positive integer.

◦ First, we (try to) diagonalize A . Since $\det(tI - A) = t^2 - 5t + 4 = (t - 1)(t - 4)$, the eigenvalues are 1 and 4. Since these are distinct, A is diagonalizable.

◦ Computing the eigenvectors of A yields that $\begin{bmatrix} -2 \\ 1 \end{bmatrix}$ is a basis for the 1-eigenspace, and $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ is a basis for the 4-eigenspace.

◦ Then $D = Q^{-1}AQ$ where $D = \begin{bmatrix} 1 & 0 \\ 0 & 4 \end{bmatrix}$ and $Q = \begin{bmatrix} -2 & -1 \\ 1 & 1 \end{bmatrix}$, and also $Q^{-1} = \begin{bmatrix} -1 & -1 \\ 1 & 2 \end{bmatrix}$.

◦ Then $D^k = \begin{bmatrix} 1 & 0 \\ 0 & 4^k \end{bmatrix}$, so $A^k = QD^kQ^{-1} = \begin{bmatrix} -2 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 4^k \end{bmatrix} \begin{bmatrix} -1 & -1 \\ 1 & 2 \end{bmatrix} = \boxed{\begin{bmatrix} 2 - 4^k & 2 - 2 \cdot 4^k \\ -1 + 4^k & -1 + 2 \cdot 4^k \end{bmatrix}}$.

◦ Remark: This formula also makes sense for values of k which are not positive integers. For example, if $k = -1$ we get the matrix $\begin{bmatrix} 7/4 & 3/2 \\ -3/4 & -1/2 \end{bmatrix}$, which is actually the inverse matrix A^{-1} . And if we set

$k = \frac{1}{2}$ we get the matrix $B = \begin{bmatrix} 0 & -2 \\ 1 & 3 \end{bmatrix}$, whose square satisfies $B^2 = \begin{bmatrix} -2 & -6 \\ 3 & 7 \end{bmatrix} = A$.

• By diagonalizing a given matrix, we can often prove theorems in a much simpler way. Here is a typical example:

• Definition: If $T : V \rightarrow V$ is a linear operator and $p(x) = a_0 + a_1x + \dots + a_nx^n$ is a polynomial, we define $p(T) = a_0I + a_1T + \dots + a_nT^n$. Similarly, if A is an $n \times n$ matrix, we define $p(A) = a_0I_n + a_1A + \dots + a_nA^n$.

◦ Since conjugation preserves sums and products, it is easy to check that $Q^{-1}p(A)Q = p(Q^{-1}AQ)$ for any invertible Q .

• Theorem (Cayley-Hamilton): If $p(x)$ is the characteristic polynomial of a matrix A , then $p(A)$ is the zero matrix $\mathbf{0}$.

◦ The same result holds for the characteristic polynomial of a linear operator $T : V \rightarrow V$.

- Example: For the matrix $A = \begin{bmatrix} 2 & 2 \\ 3 & 1 \end{bmatrix}$, we have $\det(tI - A) = \begin{vmatrix} t-2 & -2 \\ -3 & t-1 \end{vmatrix} = (t-1)(t-2) - 6 = t^2 - 3t - 4$. We can compute $A^2 = \begin{bmatrix} 10 & 6 \\ 9 & 7 \end{bmatrix}$, and then indeed we have $A^2 - 3A - 4I_2 = \begin{bmatrix} 10 & 6 \\ 9 & 7 \end{bmatrix} - \begin{bmatrix} 6 & 6 \\ 9 & 3 \end{bmatrix} - \begin{bmatrix} 4 & 0 \\ 0 & 4 \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$.
- Proof (if A is diagonalizable): If A is diagonalizable, then let $D = Q^{-1}AQ$ with D diagonal, and $p(x)$ be the characteristic polynomial of A .
- The diagonal entries of D are the eigenvalues $\lambda_1, \dots, \lambda_n$ of A , hence are roots of the characteristic polynomial of A . So $p(\lambda_1) = \dots = p(\lambda_n) = 0$.
- Then, because raising D to a power just raises all of its diagonal entries to that power, we can see that
$$p(D) = p\left(\begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}\right) = \begin{bmatrix} p(\lambda_1) & & \\ & \ddots & \\ & & p(\lambda_n) \end{bmatrix} = \begin{bmatrix} 0 & & \\ & \ddots & \\ & & 0 \end{bmatrix} = \mathbf{0}.$$
- Now by conjugating each term and adding the results, we see that $\mathbf{0} = p(D) = p(Q^{-1}AQ) = Q^{-1}[p(A)]Q$. So by conjugating back, we see that $p(A) = Q \cdot \mathbf{0} \cdot Q^{-1} = \mathbf{0}$, as claimed.
- In the case where A is not diagonalizable, the proof of the Cayley-Hamilton theorem is more difficult. One method is to use the Jordan canonical form, mentioned in the next section.

5.3 Applications of Diagonalization

- In this section we discuss a few applications of diagonalization. Our analysis is not intended to be a deep survey of all the applications of diagonalization, but rather a broad overview of a few important topics, with examples intended to convey many of the main ideas.

5.3.1 Transition Matrices and Incidence Matrices

- In many applications, we can use linear algebra to model the behavior of an iterated system. Such models are quite common in applied mathematics, the social sciences (particularly economics), and the life sciences.
 - For example, consider a state with two cities A and B whose populations flow back and forth over time: after one year passes a resident of city A has a 10% chance of moving to city B and a 90% chance of staying in city A , while a resident of city B has a 30% chance of moving to A and a 70% chance of staying in B .
 - We would like to know what will happen to the relative populations of cities A and B over a long period of time.
 - If city A has a population of A_{old} and city B has a population of B_{old} , then one year later, we can see that city A 's population will be $A_{\text{new}} = 0.9A_{\text{old}} + 0.3B_{\text{old}}$, while B 's population will be $B_{\text{new}} = 0.1A_{\text{old}} + 0.7B_{\text{old}}$.
 - By iterating this calculation, we can in principle compute the cities' populations as far into the future as desired, but the computations rapidly become quite messy to do exactly.
 - For example, with the starting populations $(A, B) = (1000, 3000)$, here is a table of the populations (to the nearest whole person) after n years:

n	0	1	2	3	4	5	6	7	8	9	10	15	20	30
A	1000	1800	2280	2568	2741	2844	2907	2944	2966	2980	2988	2999	3000	3000
B	3000	2200	1720	1432	1259	1156	1093	1056	1034	1020	1012	1001	1000	1000

- We can see that the populations seem to approach (rather rapidly) having 3000 people in city A and 1000 in city B .
- We can do the computations above much more efficiently by writing the iteration in matrix form:
$$\begin{bmatrix} A_{\text{new}} \\ B_{\text{new}} \end{bmatrix} = \begin{bmatrix} 0.9 & 0.3 \\ 0.1 & 0.7 \end{bmatrix} \begin{bmatrix} A_{\text{old}} \\ B_{\text{old}} \end{bmatrix}.$$

- Since the population one year into the future is obtained by left-multiplying the population vector by $M = \begin{bmatrix} 0.9 & 0.3 \\ 0.1 & 0.7 \end{bmatrix}$, the population k years into the future can then be obtained by left-multiplying the population vector by M^k .
- By diagonalizing this matrix, we can easily compute M^k , and thus analyze the behavior of the population as time extends forward.
- In this case, M is diagonalizable: $M = QDQ^{-1}$ with $D = \begin{bmatrix} 1 & 0 \\ 0 & 3/5 \end{bmatrix}$ and $Q = \begin{bmatrix} 3 & -1 \\ 1 & 1 \end{bmatrix}$.
- Then $M^k = QD^kQ^{-1}$, and as $k \rightarrow \infty$, we see that $D^k \rightarrow \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$, so M^k will approach $Q \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} Q^{-1} = \begin{bmatrix} 3/4 & 3/4 \\ 1/4 & 1/4 \end{bmatrix}$.
- From this calculation, we can see that as time extends on, the cities' populations will approach the situation where 3/4 of the residents live in city A and 1/4 of the residents live in city B .
- Notice that this “steady-state” solution where the cities' populations both remain constant represents an eigenvector of the original matrix with eigenvalue $\lambda = 1$.
- The system above, in which members of a set (in this case, residents of the cities) are identified as belonging to one of several states that can change over time, is known as a stochastic process.
 - If, as in our example, the probabilities of changing from one state to another are independent of time, the system is called a Markov chain.
 - Markov chains and their continuous analogues (known as Markov processes) arise (for example) in probability problems involving repeated wagers or random walks, in economics modeling the flow of goods among industries and nations, in biology modeling the gene frequencies in populations, and in civil engineering modeling the arrival of people to buildings.
 - A Markov chain model was also used for one of the original versions of the PageRank algorithm used by Google to rank internet search results.
- Definition: A square matrix whose entries are nonnegative and whose columns sum to 1 is called a transition matrix (or a stochastic matrix).
 - Equivalently, a square matrix M is a transition matrix precisely when $M^T \mathbf{v} = \mathbf{v}$, where \mathbf{v} is the column vector of all 1s.
 - From this description, we can see that \mathbf{v} is an eigenvector of M^T of eigenvalue 1, and since M^T and M have the same characteristic polynomial, we conclude that M has 1 as an eigenvalue.
 - If it were true that M were diagonalizable and every eigenvalue of M had absolute value less than 1 (except for the eigenvalue 1), then we could apply the same argument as we did in the example to conclude that the powers of M approached a limit.
 - Unfortunately, this is not true in general: the transition matrix $M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ has M^2 equal to the identity matrix, so odd powers of M are equal to M while even powers are equal to the identity. (In this case, the eigenvalues of M are 1 and -1 .)
 - Fortunately, the argument does apply to a large class of transition matrices:
- Theorem (Markov Chains): If M is a transition matrix, then every eigenvalue λ of M has $|\lambda| \leq 1$. Furthermore, if some power of M has all entries positive, then the only eigenvalue of M of absolute value 1 is $\lambda = 1$, and the 1-eigenspace has dimension 1. In such a case, the “matrix limit” $\lim_{k \rightarrow \infty} M^k$ exists and has all columns equal to a “steady-state” solution of the Markov chain whose transition matrix is M .
 - We will not prove this theorem, although most of the arguments (when M is diagonalizable) are similar to the computations we did in the example above.
- Another situation, in a somewhat different direction, concerns the analysis of groups in networks.

- For example, suppose we have a network of people, each of whom can send direct messages to certain other people. (In combinatorics, this object is known as a directed graph.)
- We would like to study the question of who can send messages (possibly using other people as intermediaries) to whom, and in how many different possible ways.
- Concretely, suppose that we have five people 1, 2, 3, 4, and 5 where 1 can send to 2 or 4, 2 can send to 3 or 5, 3 can send to 1 or 4, 4 can send to 5, and 5 can send to 3.
- We can summarize all of this information using an incidence matrix M whose (i, j) entry is 1 if person i

can send a message to person j , and 0 otherwise: in this case, we have $M = \begin{bmatrix} 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$.

- The entries of M^2 will give us information about messages that can be sent through one intermediate person.
- For example, $(M^2)_{3,5} = M_{3,1}M_{1,5} + M_{3,2}M_{2,5} + M_{3,3}M_{3,5} + M_{3,4}M_{4,5} + M_{3,5}M_{5,5}$: a term $M_{3,k}M_{k,5}$ is equal to 1 precisely when $M_{3,k} = M_{k,5} = 1$, which is to say that 3 can send a message to 5 via person k .
- By summing, we see that the entry $(M^2)_{i,j}$ represents the total number of ways that person i can send a message to person j via one other person.

- So, since $M^2 = \begin{bmatrix} 0 & 0 & 1 & 0 & 2 \\ 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \end{bmatrix}$, we see that (for example) there are two ways 1 can send a message to 5 via one other person.

- In a similar way, $(M^d)_{i,j}$ represents the number of ways person i can send a message to person j using $d - 1$ other people in the middle.
- By summing, we see that $(M + M^2 + M^3 + \dots + M^d)_{i,j}$ gives the number of ways i can send a message to j with at most $d - 1$ other people in the middle.

- In particular, for the matrix we gave above, since $M + M^2 + M^3 + M^4 = \begin{bmatrix} 3 & 2 & 3 & 5 & 3 \\ 2 & 2 & 4 & 4 & 5 \\ 3 & 1 & 4 & 4 & 3 \\ 1 & 1 & 1 & 2 & 2 \\ 1 & 1 & 3 & 2 & 3 \end{bmatrix}$, we see that any person can send any other person a message with at most 3 intermediaries.

- Ultimately, to analyze this type of network, we want to study the behavior of powers of M , which (in the event that M is diagonalizable) we can easily do by diagonalizing M : if $M = Q^{-1}DQ$, then $(M + M^2 + \dots + M^d) = Q^{-1}(D + D^2 + \dots + D^d)Q$.

5.3.2 Systems of Linear Differential Equations

- Consider the problem of solving a system of linear differential equations.
 - First, observe that we can reduce any system of linear differential equations to a system of *first-order* linear differential equations (in more variables): if we define new variables equal to the higher-order derivatives of our old variables, then we can rewrite the old system as a system of first-order equations.
 - For example, to convert $y''' + y' = 0$ into a system of 1st-order equations, we can define new variables $z = y'$ and $w = y'' = z'$: then the single 3rd-order equation $y''' + y' = 0$ is equivalent to the 1st-order system $y' = z$, $z' = w$, $w' = -z$.
- By rearranging the equations and defining new variables appropriately, we can put any system of linear differential equations into the form

$$\begin{aligned} y_1' &= a_{1,1}(x) \cdot y_1 + a_{1,2}(x) \cdot y_2 + \dots + a_{1,n}(x) \cdot y_n + q_1(x) \\ &\vdots \\ y_n' &= a_{n,1}(x) \cdot y_1 + a_{n,2}(x) \cdot y_2 + \dots + a_{n,n}(x) \cdot y_n + q_n(x) \end{aligned}$$

for some functions $a_{i,j}(x)$ and $q_i(x)$ for $1 \leq i, j \leq n$.

◦ We can write this system more compactly using matrices: if $A = \begin{bmatrix} a_{1,1}(x) & \cdots & a_{1,n}(x) \\ \vdots & \ddots & \vdots \\ a_{n,1}(x) & \cdots & a_{n,n}(x) \end{bmatrix}$, $\mathbf{q} =$

$\begin{bmatrix} q_1(x) \\ \vdots \\ q_n(x) \end{bmatrix}$, and $\mathbf{y} = \begin{bmatrix} y_1(x) \\ \vdots \\ y_n(x) \end{bmatrix}$ so that $\mathbf{y}' = \begin{bmatrix} y_1'(x) \\ \vdots \\ y_n'(x) \end{bmatrix}$, we can write the system more compactly as $\mathbf{y}' = \mathbf{A}\mathbf{y} + \mathbf{q}$.

◦ We say that the system is homogeneous if $\mathbf{q} = \mathbf{0}$, and it is nonhomogeneous otherwise.

- We also have a version of the Wronskian in this setting for checking whether function vectors are linearly independent:

• Definition: Given n vectors $\mathbf{v}_1 = \begin{bmatrix} y_{1,1}(x) \\ \vdots \\ y_{1,n}(x) \end{bmatrix}$, \dots , $\mathbf{v}_n = \begin{bmatrix} y_{n,1}(x) \\ \vdots \\ y_{n,n}(x) \end{bmatrix}$ of length n with functions as entries,

their Wronskian is defined as the determinant $W = \begin{vmatrix} y_{1,1} & \cdots & y_{1,n} \\ \vdots & \ddots & \vdots \\ y_{n,1} & \cdots & y_{n,n} \end{vmatrix}$.

◦ By our results on row operations and determinants, we immediately see that n function vectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ of length n are linearly independent if their Wronskian is not the zero function.

- Our goal is only to outline some of the applications of linear algebra to the study of differential equations, so we will now assume that all of the entries in the matrix A are constants and that the system is homogeneous. In this case, we have the following fundamental theorem:

• Theorem (Homogeneous Systems): If the $n \times n$ coefficient matrix A is constant and I is any interval, then the set of solutions \mathbf{y} to the homogeneous system $\mathbf{y}' = \mathbf{A}\mathbf{y}$ on I is an n -dimensional vector space.

◦ This theorem guarantees the existence of solutions to the system $\mathbf{y}' = \mathbf{A}\mathbf{y}$, and gives us some information about the nature of the solution space (namely, that it is n -dimensional).

◦ We, of course, would actually like to write down the solutions explicitly.

• Our key observation is: if $\mathbf{v} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix}$ is an eigenvector of A with eigenvalue λ , then $\mathbf{y} = \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{bmatrix} e^{\lambda x}$ is a

solution to $\mathbf{y}' = \mathbf{A}\mathbf{y}$.

◦ This follows simply by differentiating $\mathbf{y} = e^{\lambda x}\mathbf{v}$ with respect to x : we see $\mathbf{y}' = \lambda e^{\lambda x}\mathbf{v} = \lambda\mathbf{y} = \mathbf{A}\mathbf{y}$.

◦ In the event that A has n linearly independent eigenvectors (which is to say, if A is diagonalizable), we will therefore obtain n solutions to the differential equation.

◦ If these solutions are linearly independent, then since we know the solution space is n -dimensional, we would be able to conclude that our solutions are a basis for the solution space. This turns out to be true:

• Theorem (Eigenvalue Method): If A has n linearly independent eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ with associated eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$, then the general solution to the matrix differential system $\mathbf{y}' = \mathbf{A}\mathbf{y}$ is given by $\mathbf{y} = C_1 e^{\lambda_1 x} \mathbf{v}_1 + C_2 e^{\lambda_2 x} \mathbf{v}_2 + \cdots + C_n e^{\lambda_n x} \mathbf{v}_n$, where C_1, \dots, C_n are arbitrary constants.

◦ Recall that the matrix A will have n linearly independent eigenvectors precisely when it is diagonalizable, which is equivalent to saying that the dimension of each eigenspace is equal to the multiplicity of the corresponding eigenvalue as a root of the characteristic polynomial of A .

- Proof: By the observation above, each of $e^{\lambda_1 x} \mathbf{v}_1, e^{\lambda_2 x} \mathbf{v}_2, \dots, e^{\lambda_n x} \mathbf{v}_n$ is a solution to $\mathbf{y}' = \mathbf{A}\mathbf{y}$. We claim that they are a basis for the solution space.
 - To show this, we know by our earlier results that the solution space of the system $\mathbf{y}' = \mathbf{A}\mathbf{y}$ is n -dimensional: thus, if we can show that these solutions are linearly independent, we would be able to conclude that our solutions are a basis for the solution space.
 - We can compute the Wronskian of these solutions: after factoring out the exponentials from each column, we obtain $W = e^{(\lambda_1 + \dots + \lambda_n)x} \det(M)$, where $M = \begin{bmatrix} | & | & | \\ \mathbf{v}_1 & \cdots & \mathbf{v}_n \\ | & | & | \end{bmatrix}$.
 - The exponential is always nonzero and the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ are (by hypothesis) linearly independent, meaning that $\det(M)$ is also nonzero. Thus, W is nonzero, so $e^{\lambda_1 x} \mathbf{v}_1, e^{\lambda_2 x} \mathbf{v}_2, \dots, e^{\lambda_n x} \mathbf{v}_n$ are linearly independent.
 - Since these solutions are therefore a basis for the solution space, we immediately conclude that the general solution to $\mathbf{y}' = \mathbf{A}\mathbf{y}$ has the form $\mathbf{y} = C_1 e^{\lambda_1 x} \mathbf{v}_1 + C_2 e^{\lambda_2 x} \mathbf{v}_2 + \dots + C_n e^{\lambda_n x} \mathbf{v}_n$, for arbitrary constants C_1, \dots, C_n .
- Example: Find all functions y_1 and y_2 such that
$$\begin{aligned} y_1' &= y_1 - 3y_2 \\ y_2' &= y_1 + 5y_2 \end{aligned}$$
 - The coefficient matrix is $A = \begin{bmatrix} 1 & -3 \\ 1 & 5 \end{bmatrix}$, whose characteristic polynomial is $\det(tI - A) = \begin{vmatrix} t-1 & 3 \\ -1 & t-5 \end{vmatrix} = (t-1)(t-5) + 3 = t^2 - 6t + 8 = (t-2)(t-4)$, so the eigenvalues of A are $\lambda = 2, 4$.
 - Since the eigenvalues are distinct, A is diagonalizable, and some calculation will produce the eigenvectors $\begin{bmatrix} -3 \\ 1 \end{bmatrix}$ for $\lambda = 2$ and $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$ for $\lambda = 4$.
 - Thus, the general solution to the system is
$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \boxed{C_1 \begin{bmatrix} -3 \\ 1 \end{bmatrix} e^{2x} + C_2 \begin{bmatrix} -1 \\ 1 \end{bmatrix} e^{4x}}$$
 - We also remark that in the event that the coefficient matrix has nonreal eigenvalues, by taking an appropriate linear combination we can produce real-valued solution vectors.
 - There is also another, quite different, method for using diagonalization to solve a homogeneous system of linear differential equations with constant coefficients.
 - As motivation, if we consider the differential equation $y' = ky$ with the initial condition $y(0) = C$, it is not hard to verify that the general solution is $y(x) = e^{kx}C$.
 - We would like to find some way to extend this result to an $n \times n$ system $\mathbf{y}' = \mathbf{A}\mathbf{y}$ with initial condition $\mathbf{y}(0) = \mathbf{c}$.
 - The natural way would be to try to define the “exponential of a matrix” e^A in such a way that e^{At} has the property that $\frac{d}{dt}[e^{At}] = Ae^{At}$: then $\mathbf{y}(t) = e^{At}\mathbf{c}$ will have $\mathbf{y}'(t) = Ae^{At}\mathbf{c} = \mathbf{A}\mathbf{y}$.
 - Definition: If A is an $n \times n$ matrix, the exponential of A is defined to be the infinite sum $e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!}$.
 - The definition is motivated by the Taylor series for the exponential of a real or complex number z ; namely, $e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!}$.
 - Like with the Taylor series, it can be shown that for any matrix A , the infinite series $\sum_{n=0}^{\infty} \frac{A^n}{n!}$ converges absolutely, in the sense that the series in each of the entries of the matrix converges absolutely.
 - Theorem (Exponential Solutions): For any $n \times n$ matrix A , the general solution to the matrix system $\mathbf{y}' = \mathbf{A}\mathbf{y}$ on an interval I is given by $\mathbf{y}(t) = e^{At}\mathbf{c}$, for an arbitrary constant vector \mathbf{c} .

- Proof: Since the infinite series for e^{At} converges absolutely, we can differentiate it term-by-term: with $e^{At} = \sum_{n=0}^{\infty} \frac{A^n}{n!} t^n$, we compute $\frac{d}{dt}[e^{At}] = \sum_{n=1}^{\infty} \frac{A^n}{n!} (nt^{n-1}) = A \sum_{n=1}^{\infty} \frac{A^{n-1}}{(n-1)!} t^{n-1} = A e^{At}$.
- Thus, $\frac{d}{dt}[\mathbf{y}(t)] = [A e^{At} \mathbf{c}] = A \mathbf{y}(t)$, as required.
- Furthermore, since \mathbf{c} is a vector of length n , the vectors of the form $e^{At} \mathbf{c}$ form an n -dimensional space, hence are all the solutions.
- All that remains is actually to *compute* the exponential of a matrix, which we have not yet explained.
 - When the matrix is diagonalizable, we can do this comparatively easily: explicitly, if $A = Q^{-1} D Q$, then
$$e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!} = \sum_{n=0}^{\infty} \frac{(Q^{-1} D Q)^n}{n!} = \sum_{n=0}^{\infty} \frac{Q^{-1} D^n Q}{n!} = Q^{-1} \left[\sum_{n=0}^{\infty} \frac{D^n}{n!} \right] Q = Q^{-1} e^D Q.$$
 - Furthermore, again from the power series definition, if $D = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{bmatrix}$, then $e^D = \begin{bmatrix} e^{\lambda_1} & & \\ & \ddots & \\ & & e^{\lambda_n} \end{bmatrix}$.
 - Thus, by using the diagonalization, we can compute the exponential of the original matrix A , and thereby use it to solve the differential equation $\mathbf{y}' = A \mathbf{y}$.
- Example: Find all functions y_1 and y_2 such that
$$\begin{aligned} y_1' &= 2y_1 - y_2 \\ y_2' &= -2y_1 + 3y_2 \end{aligned}$$
 - The coefficient matrix is $A = \begin{bmatrix} 2 & -1 \\ -2 & 3 \end{bmatrix}$, with eigenvalues $\lambda = 1, 4$. Since the eigenvalues are distinct, A is diagonalizable, and we can find eigenvectors $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ for $\lambda = 1$ and $\begin{bmatrix} 1 \\ -2 \end{bmatrix}$ for $\lambda = 4$.
 - Then with $Q = \begin{bmatrix} 1 & 1 \\ 1 & -2 \end{bmatrix}$, with $Q^{-1} = \frac{1}{3} \begin{bmatrix} 2 & 1 \\ 1 & -1 \end{bmatrix}$, we have $Q^{-1} A Q = D = \begin{bmatrix} 1 & 0 \\ 0 & 4 \end{bmatrix}$.
 - Thus, $e^{At} = Q e^{Dt} Q^{-1} = Q \begin{bmatrix} e^t & 0 \\ 0 & e^{4t} \end{bmatrix} Q^{-1} = \frac{1}{3} \begin{bmatrix} 2e^t + e^{4t} & e^t - e^{4t} \\ 2e^t - 2e^{4t} & e^t + 2e^{4t} \end{bmatrix}$.
 - Then $\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2e^t + e^{4t} & e^t - e^{4t} \\ 2e^t - 2e^{4t} & e^t + 2e^{4t} \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \end{bmatrix}$ for arbitrary constants C_1 and C_2 .

5.3.3 Non-Diagonalizable Matrices and the Jordan Canonical Form

- As we saw in the previous section, there exist matrices which are not diagonalizable. For computational purposes, however, we might still like to know what the simplest form such a non-diagonalizable matrix is similar to. The answer is given by what is called the Jordan canonical form:
- Definition: The $n \times n$ Jordan block with eigenvalue λ is the $n \times n$ matrix J having λ s on the diagonal, 1s directly above the diagonal, and zeroes elsewhere.

- Here are the general Jordan block matrices of sizes 2, 3, and 4: $\begin{bmatrix} \lambda & 1 \\ 0 & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & 1 & 0 \\ 0 & \lambda & 1 \\ 0 & 0 & \lambda \end{bmatrix}, \begin{bmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 1 & 0 \\ 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & \lambda \end{bmatrix}$.

- Definition: A matrix is in Jordan canonical form if it is a “block-diagonal matrix” of the form
$$\begin{bmatrix} J_1 & & & \\ & J_2 & & \\ & & \ddots & \\ & & & J_k \end{bmatrix},$$
 where each J_1, \dots, J_k is a square Jordan block matrix (possibly with different eigenvalues and different sizes).

- Example: The matrix $\begin{bmatrix} 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 4 \end{bmatrix}$ is in Jordan canonical form, with $J_1 = [2]$, $J_2 = [3]$, $J_3 = [4]$.
- Example: The matrix $\begin{bmatrix} 2 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{bmatrix}$ is in Jordan canonical form, with $J_1 = \begin{bmatrix} 2 & 1 \\ 0 & 2 \end{bmatrix}$ and $J_2 = [3]$.
- Example: The matrix $\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$ is in Jordan canonical form, with $J_1 = [1]$, $J_2 = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$, $J_3 = [1]$.
- Example: The matrix $\begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$ is in Jordan canonical form, with $J_1 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$ and $J_2 = [0]$.

- The main result is that, over the complex numbers, every matrix is similar to a matrix in Jordan canonical form, and (furthermore) the Jordan canonical form is unique up to rearrangement of the Jordan blocks:
- Theorem (Jordan Canonical Form): If A is any $n \times n$ matrix, then A is similar to a matrix in Jordan canonical form. Furthermore, the Jordan canonical form is unique up to rearrangement of the Jordan blocks.
 - The Jordan canonical form therefore serves as an “approximate diagonalization” for non-diagonalizable matrices, since the Jordan blocks are very close to being diagonal matrices.
 - The idea behind the Jordan canonical form is that ultimately, a non-diagonalizable linear transformation (or matrix) fails to have enough eigenvectors for us to construct a diagonal basis. By generalizing the definition of eigenvector and eigenspace, we can fill in these “missing” basis entries, and then by constructing a suitable basis of “generalized eigenvectors”, we can establish the existence of the Jordan canonical form.
- Definition: For a linear operator $T : V \rightarrow V$, a nonzero vector \mathbf{v} satisfying $(A - \lambda I)^k \mathbf{v} = \mathbf{0}$ for some positive integer k and some scalar λ is called a generalized eigenvector of T .
 - We take the analogous definition for matrices: a generalized eigenvector for A is a nonzero vector \mathbf{v} with $(A - \lambda I)^k \mathbf{v} = \mathbf{0}$ for some positive integer k and some scalar λ .
 - Observe that (regular) eigenvectors correspond to $k = 1$, and so every eigenvector is a generalized eigenvector. The converse, however, is not true:
- Example: Show that $\mathbf{v} = \begin{bmatrix} 4 \\ 1 \end{bmatrix}$ is a generalized 2-eigenvector for $A = \begin{bmatrix} 1 & -1 \\ 1 & 3 \end{bmatrix}$ that is not a (regular) 2-eigenvector.
 - We compute $(A - 2I)\mathbf{v} = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} 4 \\ 1 \end{bmatrix} = \begin{bmatrix} 5 \\ -5 \end{bmatrix}$, and since this is not zero, \mathbf{v} is not a 2-eigenvector.
 - However, $(A - 2I)^2 \mathbf{v} = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} 5 \\ -5 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, and so \mathbf{v} is a generalized 2-eigenvector, with $k = 2$.
- Here are some of the fundamental properties of generalized eigenvectors for a linear transformation $T : V \rightarrow V$ (the proofs of many of these properties are difficult, and we will take them for granted):
 - The associated constant λ for a generalized eigenvector is always an eigenvalue of T : that is, if \mathbf{v} is a nonzero vector with $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$ for some k and some λ , then λ is an eigenvalue of T .
 - Generalized eigenvectors with different associated eigenvalues are linearly independent.
 - Like (regular) eigenvectors, the generalized λ -eigenvectors (together with the zero vector) also form a subspace, called the generalized λ -eigenspace.

- The generalized λ -eigenspace of T is equal to the kernel of $(T - \lambda I)^d$, where d is the multiplicity of λ as a root of the characteristic polynomial of T . Furthermore, the dimension of the generalized λ -eigenspace is equal to d .
 - There exists a basis of V consisting of generalized eigenvectors of T .
- **Example:** Find a basis for each generalized eigenspace of $A = \begin{bmatrix} 2 & 0 & 0 \\ -1 & 2 & 1 \\ 1 & -1 & 0 \end{bmatrix}$, and verify that \mathbb{R}^3 has a basis of generalized eigenvectors for A .
 - By expanding along the top row, we see $\det(tI - A) = (t - 1)^2(t - 2)$. Thus, the eigenvalues of A are $\lambda = 1, 1, 2$.
 - For the generalized 1-eigenspace, we compute the nullspace of $(A - I)^2 = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$. Upon row-reducing, we see that the generalized 1-eigenspace has dimension 2 and is spanned by $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$.
 - For the generalized 2-eigenspace, we compute the nullspace of $A - 2I = \begin{bmatrix} 0 & 0 & 0 \\ -1 & 0 & 1 \\ 1 & -1 & -2 \end{bmatrix}$. Upon row-reducing, we see that the generalized 2-eigenspace has dimension 1 and is spanned by $\begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$.
 - Furthermore, the union of the bases for the eigenspaces, $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$, $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$, $\begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$, clearly forms a basis for \mathbb{R}^3 , as claimed. We also see that the dimension of each generalized eigenspace is indeed equal to the multiplicity of the associated eigenvalue as a root of the characteristic polynomial.
- Now that we know that V has a basis of generalized eigenvectors of T , our goal is to find as simple a basis as possible for each generalized eigenspace.
 - As motivation, suppose that there is a basis $\beta = \{\mathbf{v}_{k-1}, \mathbf{v}_{k-2}, \dots, \mathbf{v}_1, \mathbf{v}_0\}$ of V such that $T : V \rightarrow V$ has associated matrix $[T]_{\beta}^{\beta} = \begin{bmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 1 & 0 \\ 0 & 0 & \ddots & 1 \\ 0 & 0 & 0 & \lambda \end{bmatrix}$, a Jordan block matrix.
 - Then $T\mathbf{v}_{k-1} = \lambda\mathbf{v}_{k-1}$ and $T(\mathbf{v}_i) = \lambda\mathbf{v}_i + \mathbf{v}_{i+1}$ for each $0 \leq i \leq k - 2$.
 - Rearranging, we see that $(T - \lambda I)\mathbf{v}_{k-1} = \mathbf{0}$ and $(T - \lambda I)\mathbf{v}_i = \mathbf{v}_{i+1}$ for each $0 \leq i \leq k - 2$.
 - Thus, by plugging these relations into one another, we see that \mathbf{v}_0 is a generalized λ -eigenvector of T and that $\mathbf{v}_i = (T - \lambda I)^i \mathbf{v}_0$ for each $0 \leq i \leq k - 1$.
 - In other words, the basis β is composed of a “chain” of generalized eigenvectors obtained by successively applying the operator $T - \lambda I$ to a particular generalized eigenvector \mathbf{v}_0 .
- **Definition:** Suppose $T : V \rightarrow V$ is linear and \mathbf{v} is a generalized λ -eigenvector of T such that $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$ and k is minimal. The list $\{\mathbf{v}_{k-1}, \mathbf{v}_{k-2}, \dots, \mathbf{v}_1, \mathbf{v}_0\}$, where $\mathbf{v}_i = (T - \lambda I)^i \mathbf{v}$ for each $0 \leq i \leq k - 1$, is called a chain of generalized eigenvectors.
 - It can be verified that any chain of generalized eigenvectors is linearly independent. By running the calculation above in reverse, if we take $\beta = \{\mathbf{v}_{k-1}, \dots, \mathbf{v}_1, \mathbf{v}_0\}$ as an ordered basis of $W = \text{span}(\beta)$, then the matrix associated to T on W is a Jordan-block matrix.
- The final necessary piece for constructing the Jordan canonical form is then the following:

- **Theorem** (Existence of Jordan Basis): If V is finite-dimensional and $T : V \rightarrow V$ is linear, then V has an ordered basis β consisting of chains of generalized eigenvectors of T , and thus the associated matrix $[T]_{\beta}^{\beta}$ is in Jordan canonical form. Furthermore, the number of $k \times k$ Jordan blocks with eigenvalue λ in any Jordan canonical form of T is equal to $\text{rank}((T - \lambda I)^{k+1}) - 2\text{rank}((T - \lambda I)^k) + \text{rank}((T - \lambda I)^{k-1})$.

- The first part of this theorem says that the Jordan canonical form exists, while the second part guarantees that the Jordan form is unique up to rearranging the blocks, since the number of blocks of any possible size and eigenvalue depend only on the transformation T itself and not on the choice of basis.

- **Example:** Find the Jordan canonical form of $A = \begin{bmatrix} 0 & -1 & 3 & 2 \\ 1 & 0 & -2 & 0 \\ -1 & 0 & 3 & 1 \\ 2 & -1 & -3 & 0 \end{bmatrix}$.

- We compute $\det(tI - A) = t(t - 1)^3$, so the eigenvalues of A are $\lambda = 0, 1, 1, 1$. Since 0 is a non-repeated eigenvalue, there can only be a Jordan block of size 1 associated to it.

- To find the Jordan blocks with $\lambda = 1$, we have $A - I = \begin{bmatrix} -1 & -1 & 3 & 2 \\ 1 & -1 & -2 & 0 \\ -1 & 0 & 2 & 1 \\ 2 & -1 & -3 & -1 \end{bmatrix}$, with $\text{rank}(A - I) = 3$.

- Next, we compute $(A - I)^2 = \begin{bmatrix} 1 & 0 & -1 & -1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & -2 & -1 \\ -2 & 0 & 5 & 2 \end{bmatrix}$, with $\text{rank}(A - I)^2 = 2$.

- Finally, $(A - I)^3 = \begin{bmatrix} -2 & 0 & 4 & 2 \\ -1 & 0 & 2 & 1 \\ -1 & 0 & 2 & 1 \\ 1 & 0 & -2 & -1 \end{bmatrix}$ so $\text{rank}(A - I)^3 = 1$.

- Therefore, for $\lambda = 1$, we see that there are $2 - 2 \cdot 3 + 4 = 0$ blocks of size 1, $1 - 2 \cdot 2 + 3 = 0$ blocks of size 2, and $1 - 2 \cdot 1 + 2 = 1$ block of size 3.

- This means there is a Jordan 1-block of size 3 (along with the Jordan 0-block of size 1), and so the

Jordan canonical form is $\begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$.

- The Jordan canonical form has a variety of applications: its primary utility is as a general version of diagonalization, since the Jordan canonical form exists for all matrices, even non-diagonalizable ones.

- For example, we can use the Jordan canonical form (in place of the diagonalization) to prove the Cayley-Hamilton theorem for non-diagonalizable matrices.

- In most practical applications, diagonalization is sufficient, so the Jordan canonical form tends to be more useful as a theoretical tool, although it does also have some important practical applications to performing computations with matrices as well.

- For example, we can give a formula for the exponential of a matrix in Jordan canonical form, which we can use to solve systems of homogeneous first-order linear differential equations where the coefficient matrix is not diagonalizable.

Well, you're at the end of my handout. Hope it was helpful.

Copyright notice: This material is copyright Evan Dummit, 2012-2017. You may not reproduce or distribute this material without my express permission.