

Contents

4 Eigenvalues, Diagonalization, and the Jordan Canonical Form	1
4.1 Eigenvalues, Eigenvectors, and The Characteristic Polynomial	2
4.1.1 Eigenvalues and Eigenvectors	2
4.1.2 Eigenvalues and Eigenvectors of Matrices	3
4.1.3 Eigenspaces	6
4.2 Diagonalization	10
4.2.1 Diagonalizability	10
4.2.2 Computing Diagonalizations	13
4.3 Generalized Eigenvectors and the Jordan Canonical Form	15
4.3.1 Generalized Eigenvectors	17
4.3.2 The Jordan Canonical Form	20
4.4 Applications of Diagonalization and the Jordan Canonical Form	24
4.4.1 Spectral Mapping and the Cayley-Hamilton Theorem	24
4.4.2 The Spectral Theorem for Hermitian Operators	25
4.4.3 Stochastic Matrices and Markov Chains	27
4.4.4 Systems of Linear Differential Equations	28
4.4.5 Matrix Exponentials and the Jordan Form	32

4 Eigenvalues, Diagonalization, and the Jordan Canonical Form

In this chapter, we will discuss eigenvalues and eigenvectors: these are characteristic values and 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.

Unfortunately, not all linear operators are diagonalizable, so we will then discuss a method for computing the Jordan canonical form of matrix, which is the representation that is as close to a diagonal matrix as possible. To do so requires some substantial study of the closely associated notion of generalized eigenvectors, which we pursue first; then we establish the existence and uniqueness of the Jordan canonical form.

We close with a few applications of diagonalization and the Jordan canonical form, including a proof of the Cayley-Hamilton theorem that any matrix satisfies its characteristic polynomial, a proof of the spectral theorem for Hermitian operators, modeling discrete stochastic processes and Markov chains, and various applications to solving systems of differential equations and computing matrix exponentials.

4.1 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 $1 \leq i \leq n$, meaning that the linear transformation T acts on the vector \mathbf{v}_i by scalar multiplication by λ_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 .

4.1.1 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 .
 - **Remark:** 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.)
 - Note also that (implicitly) λ must be an element of the scalar field of V , since otherwise $\lambda \mathbf{v}$ does not make sense.
 - 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 $T : \mathbb{C}^2 \rightarrow \mathbb{C}^2$ 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 T 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 $\ker(\lambda I - T) = \ker(T - \lambda I)$, 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 $\lambda I - T$. The equivalence $\ker(\lambda I - T) = \ker(T - \lambda I)$ is also immediate.
- We will remark that some linear operators may have no eigenvectors at all.
- **Example:** If $I : \mathbb{R}[x] \rightarrow \mathbb{R}[x]$ 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.
- In other cases, the existence of eigenvectors may depend on the scalar field being used.
- **Example:** Show that $T : F^2 \rightarrow F^2$ defined by $T(x, y) = \langle y, -x \rangle$ has no eigenvectors when $F = \mathbb{R}$, but does have eigenvectors when $F = \mathbb{C}$.
 - If $T(x, y) = \lambda \langle x, y \rangle$, we get $y = \lambda x$ and $-x = \lambda y$, so that $(\lambda^2 + 1)y = 0$.
 - If y were zero then $x = -\lambda y$ would also be zero, impossible. Thus $y \neq 0$ and so $\lambda^2 + 1 = 0$.
 - When $F = \mathbb{R}$ there is no such scalar λ , so there are no eigenvectors in this case.
 - However, when $F = \mathbb{C}$, we get $\lambda = \pm i$, and then the eigenvectors are $\langle x, -ix \rangle$ with eigenvalue i and $\langle x, ix \rangle$ with eigenvalue $-i$.
- 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 (Eigenvalues and Matrices):** 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 .

4.1.2 Eigenvalues and Eigenvectors of Matrices

- We 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 .

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 .

◦ 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. Because of this, we will typically assume that the underlying field of scalars is \mathbb{C} (or another algebraically closed field²) 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 could 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 $\det(\lambda I - A) = 0$.
 - 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, such as Newton's method, to find roots. (But we will arrange the examples so that the polynomials will factor nicely.)

²It is a nontrivial fact from field theory, which we take for granted, that every field can be considered as a subfield of an algebraically closed field: a field in which every polynomial of positive degree can be factored into a product of linear factors.

- 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 $\boxed{2 \text{ and } 5}$.

- 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 $\boxed{1, 3, \pi}$.

- 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 $\boxed{1, 2i, -2i}$.

4.1.3 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 reduced 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

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

- 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

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

- 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

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

- 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 $\overline{A\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$, 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} , it is a straightforward induction argument to see 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 .

4.2 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 finite-dimensional vector space V , can we find a basis β of V such that the associated matrix $[T]_{\beta}^{\beta}$ is a diagonal matrix?

4.2.1 Diagonalizability

- 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 $T : F^n \rightarrow F^n$ 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 $Q \in M_{n \times n}(F)$, 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 $A \in M_{n \times n}(F)$ is diagonalizable over F if there exists an invertible $n \times n$ matrix $Q \in M_{n \times n}(F)$ for which $Q^{-1}AQ$ is a diagonal matrix.
 - Warning: We will often leave the field F implicit in our discussion. Whether a particular matrix is diagonalizable does partly depend on the field F we are working in.
 - 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 $B = Q^{-1}AQ$.
 - Thus, a matrix is diagonalizable precisely when it is similar to a diagonal matrix.
- 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 transformations, 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 T 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.
- We also must formalize the notion of what it means to have all of the necessary eigenvalues in F :
- **Definition:** If $p(x) \in F[x]$, we say that $p(x)$ splits completely over F if $p(x)$ can be written as a product of linear factors in $F[x]$: i.e., as $p(x) = a(x - r_1)(x - r_2) \cdots (x - r_d)$ for some $a, r_1, r_2, \dots, r_d \in F$.
 - Informally, a polynomial splits completely over F when all of its roots are actually elements of F , rather than in some larger field.
 - **Example:** The polynomial $x^2 - 1$ splits completely over \mathbb{Q} , since we can write $x^2 - 1 = (x - 1)(x + 1)$ in $\mathbb{Q}[x]$.
 - **Example:** The polynomial $x^2 - 2$ does not split completely over \mathbb{Q} , but it does split completely over \mathbb{R} since we can write $x^2 - 2 = (x - \sqrt{2})(x + \sqrt{2})$ in $\mathbb{R}[x]$. Notice here that the roots $\sqrt{2}$ and $-\sqrt{2}$ of this polynomial are not elements of \mathbb{Q} but are elements of \mathbb{R} .
 - If A is an $n \times n$ matrix, we say that all of the eigenvalues of A lie in F when the characteristic polynomial of A splits completely over F .
- Now we can establish our diagonalizability criterion for matrices:
- **Theorem (Diagonalizability Criterion):** A matrix $A \in M_{n \times n}(F)$ is diagonalizable (over F) if and only if all of its eigenvalues lie in F and, 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 the diagonal entries of its diagonalization are the eigenvalues of A , so they must all lie in the scalar field F .
 - Furthermore, 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 F 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 $A \in M_{n \times n}(F)$ has n distinct eigenvalues in F , then A is diagonalizable over F .
 - **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.

4.2.2 Computing Diagonalizations

- 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 $\beta = \{\langle 1, -1 \rangle, \langle -2, 3 \rangle\}$, 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 $D =$

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{bmatrix} \text{ with } Q = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 1 \\ 0 & 1 & 2 \end{bmatrix}.$$

- 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 $A^k = QD^kQ^{-1} = \begin{bmatrix} -2 & -1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 4^k \end{bmatrix} \begin{bmatrix} -1 & -1 \\ 1 & 2 \end{bmatrix} = \begin{bmatrix} 2 - 4^k & 2 - 2 \cdot 4^k \\ -1 + 4^k & -1 + 2 \cdot 4^k \end{bmatrix}$.

- 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 let $p(x)$ 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 substantially more difficult. In the next section, we will treat this case using the Jordan canonical form.

4.3 Generalized Eigenvectors and the Jordan Canonical Form

- As we saw in the previous section, there exist matrices which are not conjugate to any diagonal matrix. For computational purposes, however, we might still like to know what the simplest form to which a non-diagonalizable matrix is similar. 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 $J_2 = [0]$.

- Our goal is to prove that every matrix is similar to a Jordan canonical form and to give a procedure for computing the Jordan canonical form of a matrix.
 - 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.
 - In order to describe the procedure, however, we require some preliminary results.
- We will begin by proving that any linear transformation can be represented by an upper-triangular matrix with respect to some basis.

• Theorem (Upper-Triangular Associated Matrix): Suppose $T : V \rightarrow V$ is a linear operator on a finite-dimensional complex vector space. Then there exists a basis β of V such that the associated matrix $[T]_\beta^\beta$ is upper-triangular.

- Proof: We induct on $n = \dim(V)$.
- For the base case $n = 1$, the result holds trivially, since any basis will yield an upper-triangular matrix.
- For the inductive step, now assume $n \geq 2$, and let λ be any eigenvalue of T . (From our earlier results, T necessarily has at least one eigenvalue.)
- Define $W = \text{im}(T - \lambda I)$: since λ is an eigenvalue of T , $\ker(T - \lambda I)$ has positive dimension, so $\dim(W) < \dim(V)$.
- We claim that the map $S : W \rightarrow V$ given by $S(\mathbf{w}) = T(\mathbf{w})$ has $\text{im}(S)$ contained in W , so that S will be a linear operator on W (to which we can then apply the inductive hypothesis).
- To see this, let \mathbf{w} be any vector in W . Then $S(\mathbf{w}) = (T - \lambda I)\mathbf{w} + \lambda\mathbf{w}$, and both $(T - \lambda I)\mathbf{w}$ and $\lambda\mathbf{w}$ are in W : since W is a subspace, we conclude that $S(\mathbf{w})$ also lies in W .
- Now since S is a linear operator on W , by hypothesis there exists a basis $\gamma = \{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ for W such that the matrix $[S]_\gamma^\gamma$ is upper-triangular.
- Extend γ to a basis $\beta = \{\mathbf{w}_1, \dots, \mathbf{w}_k, \mathbf{v}_{k+1}, \dots, \mathbf{v}_n\}$ of V . We claim that $[T]_\beta^\beta$ is upper-triangular.
- The upper $k \times k$ portion of $[T]_\beta^\beta$ is the matrix $[S]_\gamma^\gamma$ which is upper-triangular by hypothesis. Furthermore, for each \mathbf{v}_i we can write $T(\mathbf{v}_i) = (T - \lambda I)\mathbf{v}_i + \lambda\mathbf{v}_i$, and $(T - \lambda I)\mathbf{v}_i$ is in W , hence is a linear combination of $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$.
- Thus, $[T]_\beta^\beta$ is upper-triangular, as claimed.

• We will now build on this result by showing that we can improve our choice of basis to yield a matrix in Jordan canonical form. We will in particular need the following refinement:

• Corollary: Suppose $T : V \rightarrow V$ is a linear operator on a finite-dimensional vector space such that the scalar field of V contains all eigenvalues of T . If λ is an eigenvalue of T having multiplicity d , then there exists a basis β of V such that the associated matrix $[T]_\beta^\beta$ is upper-triangular and where the last d entries on the diagonal of this matrix are equal to λ .

- Proof: Apply the same inductive construction as the proof above, using the eigenvalue λ at each stage of the construction where it remains an eigenvalue of the subspace W .
- We observe that the diagonal entries of $[T]_\beta^\beta$ are the eigenvalues of T (counted with multiplicity).
- Also observe that $\det(tI - T) = \det(tI - S) \cdot (t - \lambda)^{\dim(E_\lambda)}$, where E_λ is the λ -eigenspace of T . Thus, all eigenvalues of S will also lie in the scalar field of V .
- Thus, if at any stage of the construction we have not yet reached d diagonal entries equal to λ , the operator S will still have λ as an eigenvalue, and we will generate at least one additional entry of λ on the diagonal in the next step of the construction.

4.3.1 Generalized Eigenvectors

- 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, we can fill in these “missing” basis entries.
- 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 $k = 2$.
- Although it may seem that we have also generalized the idea of an eigenvalue, in fact generalized eigenvectors can only have their associated constant λ be an eigenvalue of T :
- Proposition (Eigenvalues for Generalized Eigenvectors): If $T : V \rightarrow V$ is a linear operator and \mathbf{v} is a nonzero vector satisfying $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$ for some positive integer k and some scalar λ , then λ is an eigenvalue of T . Furthermore, the eigenvalue associated to a generalized eigenvector is unique.
 - Proof: Let k be the smallest positive integer for which $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$. Then by assumption, $\mathbf{w} = (T - \lambda I)^{k-1} \mathbf{v}$ is not the zero vector, but $(T - \lambda I)\mathbf{w} = \mathbf{0}$. Thus, \mathbf{w} is an eigenvector of T with corresponding eigenvalue λ .
 - For uniqueness, we claim that $T - \mu I$ is one-to-one on the generalized λ -eigenspace for any $\mu \neq \lambda$. Then by a trivial induction, $(T - \mu I)^n$ will also be one-to-one on the generalized λ -eigenspace for each n , so no nonzero vector can be in the kernel.
 - So suppose that \mathbf{v} is a nonzero vector in the generalized λ -eigenspace and that $(T - \mu I)\mathbf{v} = \mathbf{0}$. Let k be the smallest positive integer such that $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$: then $\mathbf{w} = (T - \lambda I)^{k-1} \mathbf{v}$ is nonzero and $(T - \lambda I)\mathbf{w} = \mathbf{0}$.
 - Also, we see that $(T - \mu I)\mathbf{w} = (T - \mu I)(T - \lambda I)^{k-1} \mathbf{v} = (T - \lambda I)^{k-1} (T - \mu I)\mathbf{v} = (T - \lambda I)^{k-1} \mathbf{0} = \mathbf{0}$.
 - Then \mathbf{w} would be a nonzero vector in both the λ -eigenspace and the μ -eigenspace, which is impossible.
- Like the (regular) eigenvectors, the generalized λ -eigenvectors (together with the zero vector) also form a subspace, called the generalized λ -eigenspace:
- Proposition (Generalized Eigenspaces): For a linear operator $T : V \rightarrow V$, the set of vectors \mathbf{v} satisfying $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$ for some positive integer k is a subspace of V .
 - Proof: We verify the subspace criterion.
 - [S1]: Clearly, the zero vector satisfies the condition.
 - [S2]: If \mathbf{v}_1 and \mathbf{v}_2 have $(T - \lambda I)^{k_1} \mathbf{v}_1 = \mathbf{0}$ and $(T - \lambda I)^{k_2} \mathbf{v}_2 = \mathbf{0}$, then $(T - \lambda I)^{\max(k_1, k_2)} (\mathbf{v}_1 + \mathbf{v}_2) = \mathbf{0}$.
 - [S3]: If $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$, then $(T - \lambda I)^k (c\mathbf{v}) = \mathbf{0}$ as well.
- From the definition of generalized eigenvector alone, it may seem from the definition that the value k with $(\lambda I - T)^k \mathbf{v} = \mathbf{0}$ may be arbitrarily large. But in fact, it is always the case that we can choose $k \leq \dim(V)$ when V is finite-dimensional:

- **Theorem** (Computing Generalized Eigenspaces): If $T : V \rightarrow V$ is a linear operator and V is finite-dimensional, then the generalized λ -eigenspace of T is equal to $\ker(T - \lambda I)^{\dim(V)}$. In other words, if $(T - \lambda I)^k \mathbf{v} = \mathbf{0}$ for some positive integer k , then in fact $(T - \lambda I)^{\dim(V)} \mathbf{v} = \mathbf{0}$.
 - **Proof:** Let $S = T - \lambda I$ and define $W_i = \ker(S^i)$ for each $i \geq 1$.
 - Observe that $W_1 \subseteq W_2 \subseteq W_3 \subseteq \dots$, since if $S^i \mathbf{v} = \mathbf{0}$ then $S^{i+k} \mathbf{v} = \mathbf{0}$ for each $k \geq 1$.
 - We claim that if $W_i = W_{i+1}$, then all W_{i+k} are also equal to W_i for all $k \geq 1$: in other words, that if two consecutive terms in the sequence are equal, then all subsequent terms are equal.
 - So suppose that $W_i = W_{i+1}$, and let \mathbf{v} be any vector in W_{i+2} . Then $\mathbf{0} = S^{i+2} \mathbf{v} = S^{i+1}(S\mathbf{v})$, meaning that $S\mathbf{v}$ is in $\ker(S^{i+1}) = W_{i+1} = W_i = \ker(S^i)$. Therefore, $S^i(S\mathbf{v}) = \mathbf{0}$, so that \mathbf{v} is actually in W_{i+1} .
 - Therefore, $W_{i+2} = W_{i+1}$. By iterating this argument we conclude that $W_i = W_{i+1} = W_{i+2} = \dots$ as claimed.
 - Returning to the original argument, observe that $\dim(W_1) \leq \dim(W_2) \leq \dots \leq \dim(W_k) \leq \dim(V)$ for each $k \geq 1$.
 - Thus, since the dimensions are all nonnegative integers, we must have $\dim(W_k) = \dim(W_{k+1})$ for some $k \leq \dim(V)$, as otherwise we would have $1 \leq \dim(W_1) < \dim(W_2) < \dots < \dim(W_k)$, but this is not possible since $\dim(W_k)$ would then exceed $\dim(V)$.
 - Then $W_k = W_{k+1} = W_{k+2} = \dots = W_{\dim(V)} = W_{\dim(V)+1} = \dots$.
 - Finally, if \mathbf{v} is a generalized eigenvector, then it lies in some W_i , but since the sequence of subspaces W_i stabilizes at $W_{\dim(V)}$, we conclude that \mathbf{v} is contained in $W_{\dim(V)} = \ker(S^{\dim(V)}) = \ker(T - \lambda I)^{\dim(V)}$, as claimed.
- The theorem above gives us a completely explicit way to find the vectors in a generalized eigenspace, since we need only find all possible eigenvalues λ for T , and then compute the kernel of $(T - \lambda I)^{\dim(V)}$ for each λ .
 - We will show later that it is not generally necessary to raise $T - \lambda I$ to the full power $\dim(V)$: in fact, it is sufficient to compute the kernel of $(T - \lambda I)^{d_i}$, where d_i is the multiplicity of λ as a root of the characteristic polynomial.
 - The advantage of taking the power as $\dim(V)$, however, is that it does not depend on T or λ in any way.
- **Example:** Find a basis for each generalized eigenspace of $A = \begin{bmatrix} 2 & 0 & 0 \\ -1 & 2 & 1 \\ 1 & -1 & 0 \end{bmatrix}$.
 - 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 must compute the nullspace of $(A - I)^3 = \begin{bmatrix} 1 & 0 & 0 \\ -1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}$. Upon row-reducing, we see that the generalized 1-eigenspace has dimension 2 and is spanned by the vectors $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$.
 - For the generalized 2-eigenspace, we must compute the nullspace of $(A - 2I)^3 = \begin{bmatrix} 0 & 0 & 0 \\ -1 & 2 & 3 \\ 1 & -3 & -4 \end{bmatrix}$. Upon row-reducing, we see that the generalized 2-eigenspace has dimension 1 and is spanned by the vector $\begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$.
- In the example above, note that neither of the generalized 1-eigenvectors is a 1-eigenvector, so the 1-eigenspace of A is only 1-dimensional. Thus, A is not diagonalizable, and V does not possess a basis of eigenvectors of A .

- On the other hand, we can also easily see from our description that V does possess a basis of *generalized* eigenvectors of A .
- Our goal is now to prove that there always exists a basis of generalized eigenvectors for V . Like in our argument for (regular) eigenvectors, we first prove that generalized eigenvectors associated to different eigenvalues are linearly independent.
- **Theorem** (Independent Generalized Eigenvectors): If $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ are generalized 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 a generalized 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 generalized eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_n$ having distinct eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$.
 - Suppose that $(T - \lambda_1 I)^k \mathbf{v}_1 = \mathbf{0}$. Then applying $(T - \lambda_1 I)^k$ to both sides yields $\mathbf{0} = T(\mathbf{0}) = a_1(T - \lambda_1 I)^k \mathbf{v}_1 + \dots + a_n(T - \lambda_1 I)^k \mathbf{v}_n = a_2(T - \lambda_1 I)^k \mathbf{v}_2 + \dots + a_n(T - \lambda_1 I)^k \mathbf{v}_n$.
 - Now observe that $(T - \lambda_1 I)^k \mathbf{v}_j$ lies in the generalized λ_j -eigenspace, for each j , because if $(T - \lambda_j I)^a \mathbf{v}_j = \mathbf{0}$, then $(T - \lambda_j I)^a [(T - \lambda_1 I)^k \mathbf{v}_j] = (T - \lambda_1 I)^k [(T - \lambda_j I)^a \mathbf{v}_j] = (T - \lambda_1 I)^k \mathbf{0} = \mathbf{0}$.
 - By the inductive hypothesis, each of these vectors $a_j(T - \lambda_1 I)^k \mathbf{v}_j$ must be zero. If $a_j \neq 0$, then this would imply that \mathbf{v}_j is a nonzero vector in both the generalized λ_j -eigenspace and the generalized λ_1 -eigenspace, which is impossible. Therefore, $a_j = 0$ for all $j \geq 2$. We then have $a_1\mathbf{v}_1 = \mathbf{0}$ so $a_1 = 0$ as well, meaning that the \mathbf{v}_i are linearly independent.

- Next, we compute the dimension of a generalized eigenspace.

- **Theorem** (Dimension of Generalized Eigenspace): If V is finite-dimensional, $T : V \rightarrow V$ is linear, and λ is a scalar, then the dimension of the generalized λ -eigenspace is equal to the multiplicity d of λ as a root of the characteristic polynomial of T , and in fact the generalized λ -eigenspace is the kernel of $(T - \lambda I)^d$.

- **Proof:** Suppose the multiplicity of λ as a root of the characteristic polynomial of T is d .
- As we proved earlier, there exists a basis β of V for which the associated matrix $A = [T]_\beta^\beta$ is upper-triangular and has the last d diagonal entries equal to λ . (The remaining diagonal entries are the other eigenvalues of T , which by hypothesis are not equal to λ .)
- Then, for $B = A - \lambda I$, we see that $B = \begin{bmatrix} D & * \\ 0 & U \end{bmatrix}$, where D is upper-triangular with nonzero entries on the diagonal and U is a $d \times d$ upper-triangular matrix with zeroes on the diagonal.
- Observe that $B^{\dim(V)} = \begin{bmatrix} D^{\dim(V)} & * \\ 0 & U^{\dim(V)} \end{bmatrix}$, and also, by a straightforward induction argument, U^d is the zero matrix, so $U^{\dim(V)}$ is also the zero matrix, since $d \leq \dim(V)$.
- The generalized λ -eigenspace then has dimension equal to the nullity of $(A - \lambda I)^{\dim(V)} = B^{\dim(V)}$, but since $D^{\dim(V)}$ is upper-triangular with nonzero entries on the diagonal, we see that the nullity of $B^{\dim(V)}$ is exactly d .
- The last statement follows from the observation that U^d is the zero matrix.

- **Example:** Find the dimensions of the generalized eigenspaces of $A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 2 & -3 & 1 \\ 0 & 1 & -2 & 1 \\ 0 & 0 & -1 & 1 \end{bmatrix}$, and then verify the result by finding a basis for each generalized eigenspace.

- Some computation produces $\det(tI - A) = t^3(t - 1)$. Thus, the eigenvalues of A are $\lambda = 0, 0, 0, 1$.
- So by the theorem above, the dimension of the generalized 0-eigenspace is 3 and the dimension of the generalized 1-eigenspace is 1.

- For the generalized 0-eigenspace, the nullspace of $A^4 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 \end{bmatrix}$ has basis $\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}$, $\begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}$, $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$.

- Since 1 is a root of multiplicity 1, the generalized 1-eigenspace is simply the 1-eigenspace, and row-reducing $I - A = \begin{bmatrix} 1 & 0 & -1 & 0 \\ 0 & -1 & 3 & -1 \\ 0 & -1 & 3 & -1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$ yields a basis vector $\begin{bmatrix} 0 \\ 1 \\ 0 \\ -1 \end{bmatrix}$.

- At last, we can show that any finite-dimensional (complex) vector space has a basis of generalized eigenvectors:
- **Theorem** (Spectral Decomposition): If V is finite-dimensional, $T : V \rightarrow V$ is linear, and all eigenvalues of T lie in the scalar field of V , then V has a basis of generalized eigenvectors of T .
 - **Proof:** Suppose the eigenvalues of T are λ_i with respective multiplicities d_i as roots of the characteristic polynomial, and let $\beta_i = \{\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,d_i}\}$ be a basis for the generalized λ_i -eigenspace for each $1 \leq i \leq k$.
 - We claim that $\beta = \beta_1 \cup \dots \cup \beta_k$ is a basis for V .
 - By the previous theorem, the number of elements in β_i is d_i : then β contains $\sum_i d_i = \dim(V)$ vectors, so to show β is a basis it suffices to prove linear independence.
 - So suppose we have a dependence $a_{1,1}\mathbf{v}_{1,1} + \dots + a_{k,j}\mathbf{v}_{k,j} = \mathbf{0}$. Let $\mathbf{w}_i = \sum_j a_{i,j}\mathbf{v}_{i,j}$: observe that \mathbf{w}_i lies in the generalized λ_i -eigenspace 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 generalized eigenvectors of T having distinct eigenvalues, which is impossible.
 - Therefore, each $\mathbf{w}_i = \mathbf{0}$, meaning that $a_{i,1}\mathbf{v}_{i,1} + \dots + a_{i,d_i}\mathbf{v}_{i,d_i} = \mathbf{0}$. 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 .

4.3.2 The Jordan Canonical Form

- Now that we have established the existence of a basis of generalized eigenvectors (under the assumption that V is finite-dimensional and that its scalar field contains all eigenvalues of T), our goal is to find as simple a basis as possible for each generalized eigenspace.
- To motivate our discussion, 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 a trivial induction, 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.

- By running the calculation above in reverse (assuming for now that the \mathbf{v}_i are linearly independent), 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 has the form $\begin{bmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 1 & 0 \\ 0 & 0 & \ddots & 1 \\ 0 & 0 & 0 & \lambda \end{bmatrix}$: in other words, a Jordan-block matrix.

- Our goal is to prove that there exists a basis for the generalized λ -eigenspace consisting of chains of generalized eigenvectors: by applying this to each generalized eigenspace, we obtain a Jordan canonical form for T .
- A simple way to construct chains of generalized eigenvectors is simply to find a generalized eigenvector and then repeatedly apply $T - \lambda I$ to it.

- Example: If $A = \begin{bmatrix} -1 & 2 & -2 & 1 \\ -1 & 2 & -1 & 1 \\ 0 & 0 & 1 & 0 \\ -1 & 1 & -2 & 1 \end{bmatrix}$, find a chain of generalized 1-eigenvectors for A having length 3.

- We compute $\det(tI - A) = t(t - 1)^3$. Thus, the eigenvalues of A are $\lambda = 0, 1, 1, 1$.
- By our theorems, the 1-eigenspace is 3-dimensional and equal to the nullspace of the matrix

$$(A - I)^3 = \begin{bmatrix} -1 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -1 & 1 & -1 & 0 \end{bmatrix}, \text{ hence has a basis } \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}.$$

- The first vector is an eigenvector of A (so it only produces a chain of length 0), but if we instead take

$$\mathbf{v} = \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \text{ we get } (A - I)\mathbf{v} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ -1 \end{bmatrix} \text{ and } (A - I)^2\mathbf{v} = \begin{bmatrix} -1 \\ -1 \\ 0 \\ 0 \end{bmatrix}, \text{ which has length 3.}$$

- However, this procedure of constructing a chain starting from an arbitrary generalized eigenvector is rather haphazard.
 - If we are looking to construct a chain of generalized eigenvectors in a more careful manner, we could instead run the construction in the opposite direction, by starting with a collection of eigenvectors and trying to find generalized eigenvectors that are mapped to them by $T - \lambda I$.
 - By refining this idea appropriately, we can give a method for constructing a basis for V consisting of chains of generalized eigenvectors.
- Theorem (Existence of Jordan Basis): If V is finite-dimensional, $T : V \rightarrow V$ is linear, and all eigenvalues of T lie in the scalar field of V , then V has a basis consisting of chains of generalized eigenvectors of T .

- Proof: It suffices to show that each eigenspace has a basis consisting of chains of generalized eigenvectors, since (as we already showed) the union of bases for the generalized eigenspaces will be a basis for V .
- So suppose λ is an eigenvalue of T , let W be the generalized λ -eigenspace of V , with $\dim(W) = d$.
- Also, take $S : W \rightarrow W$ to be the map $S = T - \lambda I$, and note (as we showed) that S^d is the zero transformation on W .
- We must then prove that there exist vectors $\mathbf{w}_1, \dots, \mathbf{w}_k$ and integers a_1, \dots, a_k such that $S^{a_i}(\mathbf{w}_i) = \mathbf{0}$ and the set $\{\mathbf{w}_1, S\mathbf{w}_1, \dots, S^{a_1-1}\mathbf{w}_1, \mathbf{w}_2, S\mathbf{w}_2, \dots, S^{a_2-1}\mathbf{w}_2, \dots, \mathbf{w}_k, \dots, S^{a_k-1}\mathbf{w}_k\}$ is a basis of W .
- We will show this result by (strong) induction on d . If $d = 1$ then the result is trivial, since then S is the zero transformation so we can take $a_1 = 1$ and \mathbf{w}_1 to be any nonzero vector in W .
- Now assume $d > 2$ and that the result holds for all spaces of dimension less than d .
- Since $S : W \rightarrow W$ is not one-to-one (else it would be an isomorphism, but then S^d could not be zero) $W' = \text{im}(S)$ has dimension strictly less than $d = \dim(W)$.
- If W' is the zero space, then we can take $a_1 = \dots = a_k = 1$ and $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$ to be any basis of W .
- Otherwise, if W' is not zero, then by the inductive hypothesis, there exist vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ and integers a_1, \dots, a_k such that $S^{a_i}(\mathbf{v}_i) = \mathbf{0}$ and the set $\beta' = \{\mathbf{v}_1, \dots, S^{a_1-1}\mathbf{v}_1, \dots, \mathbf{v}_k, \dots, S^{a_k-1}\mathbf{v}_k\}$ is a basis of W' .
- Now, since each \mathbf{v}_i is in $W' = \text{im}(S)$, by definition there exists a vector \mathbf{w}_i in W with $S\mathbf{w}_i = \mathbf{v}_i$. (In other words, can “extend” each of the chains for W' to obtain chains for W .)

- Furthermore, note that $\{S^{a_1-1}\mathbf{v}_1, \dots, S^{a_k-1}\mathbf{v}_k\}$ are linearly independent vectors in $\ker(S)$, so we can extend that set to obtain a basis $\gamma = \{S^{a_1-1}\mathbf{v}_1, \dots, S^{a_k-1}\mathbf{v}_k, \mathbf{z}_1, \dots, \mathbf{z}_s\}$ of $\ker(S)$.
 - We claim that the set $\beta = \{\mathbf{w}_1, \dots, S^{a_1}\mathbf{w}_1, \dots, \mathbf{v}_k, \dots, S^{a_k}\mathbf{w}_k, \mathbf{z}_1, \dots, \mathbf{z}_s\}$ is the desired basis for W . It clearly has the proper form, since $S\mathbf{z}_i = \mathbf{0}$ for each i , and the total number of vectors is $a_1 + \dots + a_k + s + k$.
 - Furthermore, since $\{\mathbf{v}_1, \dots, S^{a_1-1}\mathbf{v}_1, \dots, \mathbf{v}_k, \dots, S^{a_k-1}\mathbf{v}_k\}$ is a basis of W' , $\dim(\text{im}T) = a_1 + \dots + a_k$, and since $\{S^{a_1-1}\mathbf{v}_1, \dots, S^{a_k-1}\mathbf{v}_k, \mathbf{z}_1, \dots, \mathbf{z}_s\}$ is a basis of $\ker(T)$, we see $\dim(\ker T) = s + k$.
 - Then $\dim(W) = \dim(\ker T) + \dim(\text{im}T) = a_1 + \dots + a_k + s + k$, and so we see that the set β contains the proper number of vectors.
 - It remains to verify that β is linearly independent. So suppose that $c_{1,1}\mathbf{w}_1 + \dots + c_{k,a_k}S^{a_k-1}\mathbf{w}_k + b_1\mathbf{z}_1 + \dots + b_s\mathbf{z}_s = \mathbf{0}$.
 - Since $S^m\mathbf{w}_i = S^{m-1}\mathbf{v}_i$, applying S to both sides yields $c_{1,1}\mathbf{v}_1 + \dots + c_{k,a_k-1}S^{a_k-1}\mathbf{v}_k = \mathbf{0}$, so since β' is linearly independent, all coefficients must be zero.
 - The original dependence then reduces to $c_{1,a_1}S^{a_1}\mathbf{w}_1 + \dots + c_{k,a_k}\mathbf{w}_k + b_1\mathbf{z}_1 + \dots + b_s\mathbf{z}_s = \mathbf{0}$, but since γ is linearly independent, all coefficients must be zero. Thus, β is linearly independent and therefore a basis for W .
- Using the theorem above, we can establish the existence of the Jordan form, which also turns out to be essentially unique:
 - **Theorem** (Jordan Canonical Form): If V is finite-dimensional, $T : V \rightarrow V$ is linear, and all eigenvalues of T lie in the scalar field of V , then there exists a basis β of V such that $[T]_\beta^\beta$ is a matrix in Jordan canonical form. Furthermore, the Jordan canonical form is unique up to rearrangement of the Jordan blocks.
 - **Proof:** By the theorem above, each eigenspace of T has a basis consisting of chains of generalized eigenvectors. If $\{\mathbf{v}, S\mathbf{v}, \dots, S^{a-1}\mathbf{v}\}$ is such a chain, where $S = T - \lambda I$ and $S^a\mathbf{v} = \mathbf{0}$, then we can easily see that $T(S^b\mathbf{v}) = (S + \lambda)S^b\mathbf{v} = S^{b+1}\mathbf{v} + \lambda(S^b\mathbf{v})$, and so the associated matrix for this portion of the basis is a Jordan-block matrix of size a and eigenvalue λ .
 - Therefore, if we take β to be the union of chains of generalized eigenvectors for each eigenspace, then $[T]_\beta^\beta$ is a matrix in Jordan canonical form.
 - For the uniqueness, we claim that the number of Jordan blocks of eigenvalue λ having size at least d is equal to $\dim(\ker(T - \lambda I)^{d-1}) - \dim(\ker(T - \lambda I)^d)$. Since this quantity depends only on T (and not on the particular choice of basis) and completely determines the exact number of each type of Jordan block, the number of Jordan blocks of each size and eigenvalue must be the same in any Jordan canonical form.
 - To see this, let $S = T - \lambda I$ and take $\{\mathbf{w}_1, S\mathbf{w}_1, \dots, S^{a_1-1}\mathbf{w}_1, \mathbf{w}_2, S\mathbf{w}_2, \dots, S^{a_2-1}\mathbf{w}_2, \dots, \mathbf{w}_k, \dots, S^{a_k-1}\mathbf{w}_k\}$ to be a Jordan basis for the generalized λ -eigenspace: the sizes of the Jordan blocks are then $a_1 \leq a_2 \leq \dots \leq a_k$.
 - Then a basis for the kernel of S^d is given by $\{S^{a_i-d}\mathbf{w}_i, \dots, S^{a_i-1}\mathbf{w}_i, \dots, S^{a_i-d}\mathbf{w}_k, \dots, S^{a_k-1}\mathbf{w}_k\}$, where i is the smallest value such that $d \leq a_i$.
 - We can see that in extending the basis of $\ker(S^{d-1})$ to a basis of $\ker(S^d)$, we adjoin the additional vectors $\{S^{a_i-d}\mathbf{w}_i, S^{a_{i+1}-d}\mathbf{w}_{i+1}, \dots, S^{a_k-d}\mathbf{w}_k\}$, and the number of such vectors is precisely the number of a_i that are at least d .
 - Thus, $\dim(\ker S^{d-1}) - \dim(\ker S^d)$ is the number of Jordan blocks of size at least d , as claimed.
 - In addition to proving the existence of the Jordan canonical form, the theorem above also gives us a method for computing it explicitly: all we need to do is find the dimensions of $\ker(T - \lambda I)$, $\ker(T - \lambda I)^2$, \dots , $\ker(T - \lambda I)^d$ where d is the multiplicity of the eigenvalue λ , and then use the results to find the number of each type of Jordan block.
 - From the analysis above, the number of $d \times d$ Jordan blocks with eigenvalue λ is equal to $-\dim(\ker(T - \lambda I)^{d+1}) + 2\dim(\ker(T - \lambda I)^d) - \dim(\ker(T - \lambda I)^{d-1})$, which, by the nullity-rank theorem, is also equal to $\text{rank}((T - \lambda I)^{d+1}) - 2\text{rank}((T - \lambda I)^d) + \text{rank}((T - \lambda I)^{d-1})$.
 - When actually working with the Jordan form J of a particular matrix A , one also wants to know the conjugating matrix Q with $A = Q^{-1}JQ$.

- By our theorems, we can take the columns of Q to be chains of generalized eigenvectors, but actually computing these chains is more difficult. A procedure for doing these calculations can be extracted from our proof of the theorem above, but we will not describe it explicitly.

- Example: Find the Jordan canonical form of $A = \begin{bmatrix} 0 & 1 & 0 & 1 \\ -4 & 3 & 1 & 3 \\ -5 & 3 & 2 & 4 \\ 3 & -1 & -1 & -1 \end{bmatrix}$.

- We compute $\det(tI - A) = (t - 1)^4$, so the eigenvalues of A are $\lambda = 1, 1, 1, 1$, meaning that all of the Jordan blocks have eigenvalue 1.

- To find the sizes, we have $A - I = \begin{bmatrix} -1 & 1 & 0 & 1 \\ -4 & 2 & 1 & 3 \\ -5 & 3 & 1 & 4 \\ 3 & -1 & -1 & -2 \end{bmatrix}$. Row-reducing $A - I$ yields $\begin{bmatrix} 1 & -1 & 0 & -1 \\ 0 & 2 & -1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$,

so $\text{rank}(A - I) = 2$. Furthermore, we can compute that $(A - I)^2$ is the zero matrix, so $\text{rank}(A - I)^2 = 0$.

- Thus, the number of 1×1 Jordan blocks is $\text{rank}(A - I)^2 - 2\text{rank}(A - I)^1 + \text{rank}(A - I)^0 = 0 - 2 \cdot 2 + 4 = 0$, and the number of 2×2 Jordan blocks is $\text{rank}(A - I)^3 - 2\text{rank}(A - I)^2 + \text{rank}(A - I)^1 = 0 - 2 \cdot 0 + 2 = 2$.
- Thus, there are 2 blocks of size 2 with eigenvalue 1 (and no blocks of other sizes or other eigenvalues),

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

- 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 eigenvalue 1, we have $A - I = \begin{bmatrix} -1 & -1 & 3 & 2 \\ 1 & -1 & -2 & 0 \\ -1 & 0 & 2 & 1 \\ 2 & -1 & -3 & -1 \end{bmatrix}$. Row-reducing

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

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

- 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 $\begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$.

4.4 Applications of Diagonalization and the Jordan Canonical Form

- The Jordan canonical form, and also the special case of diagonalization, have a wide variety of applications. The Jordan form is primarily useful as a theoretical tool, although it does also have some important practical applications to performing computations with matrices as well.

4.4.1 Spectral Mapping and the Cayley-Hamilton Theorem

- First, we establish the Cayley-Hamilton theorem for arbitrary matrices:
- **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 $T : V \rightarrow V$ on a finite-dimensional vector space.
 - **Proof:** Since the characteristic polynomial of a matrix does not depend on the underlying field of coefficients, we may assume that the characteristic polynomial factors completely over the field (i.e., that all of the eigenvalues of A lie in the field) by replacing the field with its algebraic closure.
 - Then by our results, A has a Jordan canonical form J such that $J = Q^{-1}AQ$ for some invertible Q . Also let $p(x) = (x - \lambda_1)^{d_1} \cdots (x - \lambda_k)^{d_k}$ be the characteristic polynomial of A .
 - We first claim that for a $d \times d$ Jordan block matrix J_i with associated eigenvalue λ_i , we have $(J_i - \lambda_i I)^d = \mathbf{0}$.
 - To see this, let $T : V \rightarrow V$ be a linear transformation on a d -dimensional vector space with ordered basis $\{\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{d-1}\}$ having associated matrix J_i and let $S = T - \lambda_i I$.
 - Then by construction, $\mathbf{v}_{i+1} = S\mathbf{v}_i$ for each $0 \leq i \leq d-2$, and $S\mathbf{v}_{d-1} = \mathbf{0}$: we then see $S^d\mathbf{v}_i = S^{i+d}\mathbf{v}_0 = S^i\mathbf{v}_{d-1} = \mathbf{0}$, so S^d is the zero transformation on V , as required.
 - Now, if J_i is any $d \times d$ Jordan block in J of eigenvalue λ_i , the characteristic polynomial of A is divisible by $(t - \lambda_i)^d$, since λ_i occurs as an eigenvalue with multiplicity at least d . Therefore, $p(J_i) = (J_i - \lambda_i I)^{d_1} \cdots (J_i - \lambda_i I)^{d_i} \cdots (J_i - \lambda_k I)^{d_k}$, and by the calculation above, $(J_i - \lambda_i I)^{d_i} = \mathbf{0}$, so $p(J_i) = \mathbf{0}$.
 - We then see $p(J) = \begin{bmatrix} p(J_1) & & \\ & \ddots & \\ & & p(J_n) \end{bmatrix} = \mathbf{0}$, and then finally, $p(A) = Q[p(J)]Q^{-1} = Q(\mathbf{0})Q^{-1} = \mathbf{0}$, as required.
- Using the same ideas, we can also establish the spectral mapping theorem:
- **Theorem** (Spectral Mapping): If $T : V \rightarrow V$ is a linear operator on an n -dimensional vector space having eigenvalues $\lambda_1, \dots, \lambda_n$ (counted with multiplicity), then for any polynomial $q(x)$, the eigenvalues of $q(T)$ are $q(\lambda_1), \dots, q(\lambda_n)$.
 - In fact, this result holds if q is replaced by any function that can be written as a convergent power series (for example, the exponential function).
 - **Proof:** Let β be a basis for V such that $[T]_\beta^\beta = J$ is in Jordan canonical form. Then $[q(T)]_\beta^\beta = q(J)$, so it suffices to find the eigenvalues of $q(J)$.
 - Now observe that if B is any upper-triangular matrix with diagonal entries $b_{1,1}, \dots, b_{n,n}$, then $q(B)$ is also upper-triangular and has diagonal entries $q(b_{1,1}), \dots, q(b_{n,n})$.
 - Applying this to the Jordan canonical form J , we see that the diagonal entries of $q(J)$ are $q(\lambda_1), \dots, q(\lambda_n)$, and the diagonal entries of any upper-triangular matrix are its eigenvalues (counted with multiplicity).

4.4.2 The Spectral Theorem for Hermitian Operators

- We now use our results on generalized eigenvectors and the Jordan canonical form to establish a fundamental result about the diagonalizability of self-adjoint operators known as the spectral theorem:
- **Definition:** If $T : V \rightarrow V$ is a linear transformation and T^* exists, we say T is Hermitian (or self-adjoint) if $T^* = T$, and that T is skew-Hermitian if $T^* = -T$.

- We extend this definition to matrices in the natural way: we say a matrix A is (skew)-Hermitian if $A = [T]_{\beta}^{\beta}$ for some basis β of V and some (skew)-Hermitian linear transformation T .
- As we showed above, the matrix associated to T^* is A^* , the conjugate-transpose of A , so A is Hermitian precisely when $A^* = A$ and A is skew-Hermitian precisely when $A^* = -A$.
- If A is a matrix with real entries, then A is Hermitian if and only if $A^T = A$ (i.e., A is a symmetric matrix), and A is skew-Hermitian if and only if $A^T = -A$ (i.e., A is a skew-symmetric matrix).

- Hermitian linear operators (and Hermitian matrices) have a variety of very nice properties. Among the most fundamental of these properties is that all of their eigenvalues are real, and that they are diagonalizable:
- **Theorem** (Properties of Hermitian Operators): Suppose V is a finite-dimensional inner product space and $T : V \rightarrow V$ is a Hermitian linear transformation. Then the following hold:

1. For any $\mathbf{v} \in V$, $\langle T(\mathbf{v}), \mathbf{v} \rangle$ is a real number.

- **Proof:** We have $\langle T(\mathbf{v}), \mathbf{v} \rangle = \langle \mathbf{v}, T^*(\mathbf{v}) \rangle = \langle \mathbf{v}, T(\mathbf{v}) \rangle = \overline{\langle T(\mathbf{v}), \mathbf{v} \rangle}$, so $\langle T(\mathbf{v}), \mathbf{v} \rangle$ is equal to its complex conjugate, hence is real.

2. All eigenvalues of T are real numbers.

- **Proof:** Suppose λ is an eigenvalue of T with eigenvector $\mathbf{v} \neq \mathbf{0}$.
- Then $\langle T(\mathbf{v}), \mathbf{v} \rangle = \langle \lambda \mathbf{v}, \mathbf{v} \rangle = \lambda \langle \mathbf{v}, \mathbf{v} \rangle$ is real. Since \mathbf{v} is not the zero vector we conclude that $\langle \mathbf{v}, \mathbf{v} \rangle$ is a nonzero real number, so λ is also real.

3. Eigenvectors of T with different eigenvalues are orthogonal.

- **Proof:** Suppose that $T\mathbf{v}_1 = \lambda_1\mathbf{v}_1$ and $T\mathbf{v}_2 = \lambda_2\mathbf{v}_2$.
- Then $\lambda_1 \langle \mathbf{v}_1, \mathbf{v}_2 \rangle = \langle T\mathbf{v}_1, \mathbf{v}_2 \rangle = \langle \mathbf{v}_1, T^*\mathbf{v}_2 \rangle = \langle \mathbf{v}_1, \lambda_2\mathbf{v}_2 \rangle = \lambda_2 \langle \mathbf{v}_1, \mathbf{v}_2 \rangle$ since λ_2 is real. But since $\lambda_1 \neq \lambda_2$, this means $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle = 0$.

4. Every generalized eigenvector of T is an eigenvector of T .

- **Proof:** We show by induction that if $(T - \lambda I)^k \mathbf{w} = \mathbf{0}$ then in fact $(T - \lambda I)\mathbf{w} = \mathbf{0}$.
- For the base case we take $k = 2$, so that $(\lambda I - T)^2 \mathbf{w} = \mathbf{0}$. Then since λ is an eigenvalue of T and therefore real, we have

$$\begin{aligned} \mathbf{0} = \langle (T - \lambda I)^2 \mathbf{w}, \mathbf{w} \rangle &= \langle (T - \lambda I)\mathbf{w}, (T - \lambda I)^* \mathbf{w} \rangle \\ &= \langle (T - \lambda I)\mathbf{w}, (T^* - \bar{\lambda} I)\mathbf{w} \rangle \\ &= \langle (T - \lambda I)\mathbf{w}, (T - \lambda I)\mathbf{w} \rangle \end{aligned}$$

and thus the inner product of $(T - \lambda I)\mathbf{w}$ with itself is zero, so $(T - \lambda I)\mathbf{w}$ must be zero.

- For the inductive step, observe that $(T - \lambda I)^{k+1} \mathbf{w} = \mathbf{0}$ implies $(T - \lambda I)^k [(T - \lambda I)\mathbf{w}] = \mathbf{0}$, and therefore by the inductive hypothesis this means $(T - \lambda I)[(T - \lambda I)\mathbf{w}] = \mathbf{0}$, or equivalently, $(T - \lambda I)^2 \mathbf{w} = \mathbf{0}$. Applying the result for $k = 2$ from above yields $(T - \lambda I)\mathbf{w} = \mathbf{0}$, as required.

- Using these basic properties, we can prove that Hermitian operators are diagonalizable, and in fact that they are diagonalizable in a particularly nice way:

- **Theorem** (Spectral Theorem): Suppose V is a finite-dimensional inner product space over \mathbb{R} or \mathbb{C} and $T : V \rightarrow V$ is a Hermitian linear transformation. Then V has an orthonormal basis β of eigenvectors of T , so in particular, T is diagonalizable.

- The set of scalars λ for which $T - \lambda I$ is not invertible is called the spectrum of T , which when V is finite-dimensional is simply the set of eigenvalues of T . The spectral theorem shows that V is the direct sum of the eigenspaces of T , meaning that the action of T on V can be decomposed into simple pieces (acting as scalar multiplication), with one piece coming from each element of the spectrum. (This is the reason for the name of the theorem.)
 - Proof: By the theorem above, every eigenvalue of T is real hence lies in the scalar field.
 - Then every generalized eigenvector of T is an eigenvector of T , and so since V has a basis of generalized eigenvectors, it has a basis of eigenvectors and is therefore diagonalizable.
 - To finish the proof, start with a basis for each eigenspace, and then apply Gram-Schmidt, yielding an orthonormal basis for each eigenspace.
 - Since T is diagonalizable, the union of these bases is a basis for V : furthermore, each of the vectors has norm 1, and they are all orthogonal by the orthogonal result above.
 - By construction, each vector is orthogonal to the others in its eigenspace, and by the observation above it is also orthogonal to the vectors in the other eigenspaces, so we obtain an orthonormal basis β of eigenvectors of T .
- A very useful special case is the situation where $V = \mathbb{R}^n$ or \mathbb{C}^n (i.e., the matrix version), in which case the theorem says that every Hermitian matrix A can be written as $A = UDU^{-1}$ where D is a real diagonal matrix and U is a unitary matrix (i.e., satisfying $U^* = U^{-1}$).

- Example: The real symmetric matrix $A = \begin{bmatrix} 3 & 6 \\ 6 & 8 \end{bmatrix}$ has eigenvalues $\lambda = -1, 12$ and has $A = UDU^{-1}$ where $D = \begin{bmatrix} -1 & 0 \\ 0 & 12 \end{bmatrix}$ and $U = \frac{1}{\sqrt{13}} \begin{bmatrix} -3 & 2 \\ 2 & 3 \end{bmatrix}$.

- Example: The Hermitian matrix $A = \begin{bmatrix} 6 & 2-i \\ 2+i & 2 \end{bmatrix}$ has eigenvalues $\lambda = 1, 7$ and has $A = UDU^{-1}$ where $D = \begin{bmatrix} 1 & 0 \\ 0 & 7 \end{bmatrix}$ and $U = \frac{1}{\sqrt{30}} \begin{bmatrix} 5 & 2-i \\ 2+i & -5 \end{bmatrix}$.

- Note that U being unitary is equivalent to saying that its columns form an orthonormal basis for V , so, up to possibly including a reflection along one axis, U represents a rotation in space around the origin.
 - Therefore, we may interpret the spectral theorem geometrically as saying that we can decompose any Hermitian transformation of V into a sequence of a rotation of the coordinate axes (applying U^*) followed by a scaling along each coordinate axis (applying D), and then undoing the rotation (applying U).
- As a corollary we obtain the following extremely useful computational fact:
 - Corollary: Every real symmetric matrix has real eigenvalues and is diagonalizable over the real numbers.
 - Proof: This follows immediately from the spectral theorem on $V = \mathbb{R}^n$ since a real symmetric matrix is Hermitian.

- To find the required diagonalization, we need only compute an orthonormal basis for each eigenspace.

- Example: For $A = \begin{bmatrix} 3 & 2 & -2 \\ 2 & 2 & 0 \\ -2 & 0 & 4 \end{bmatrix}$, find a diagonal matrix D and a unitary matrix U such that $A = UDU^{-1}$.

- First, we find the eigenvalues of A . The characteristic polynomial is $p(t) = \det(tI - A) = t^3 - 9t^2 + 18t = t(t-3)(t-9)$ so the eigenvalues are $\lambda = 0, 3, 6$ and the eigenspaces are all 1-dimensional.

- A short calculation then yields the orthonormal bases $\frac{1}{3} \begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix}$, $\frac{1}{3} \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$, and $\frac{1}{3} \begin{bmatrix} 2 \\ 1 \\ -2 \end{bmatrix}$ for the 0-, 3-, and 6-eigenspaces respectively.

- Then the desired matrices are $D = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 6 \end{bmatrix}$ and $U = \frac{1}{3} \begin{bmatrix} 2 & 1 & 2 \\ -2 & 2 & 1 \\ 1 & 2 & -2 \end{bmatrix}$.

- We will remark that although real symmetric matrices are diagonalizable, and complex Hermitian matrices are diagonalizable, it is *not* true that complex symmetric matrices are always diagonalizable.
 - For example, the complex symmetric matrix $\begin{bmatrix} 1 & i \\ i & -1 \end{bmatrix}$ is not diagonalizable. This follows from the observation that its trace and determinant are both zero, but since it is not the zero matrix, the only possibility for its Jordan form is $\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$.
- We also remark that most of these results also extend to the class of skew-Hermitian operators (having the property that $T^* = -T$), with appropriate minor modifications.
 - For example, every eigenvalue of a skew-Hermitian operator is a pure imaginary number (i.e., of the form ai for some real number a), and every skew-Hermitian operator is diagonalizable over \mathbb{C} via an orthonormal basis of eigenvectors.
 - All of these statements follow immediately from the simple observation that T is skew-Hermitian if and only if iT is Hermitian.
- Additionally, the converse of the spectral theorem not quite true: if V has an orthonormal basis of eigenvectors of T , then T is not necessarily Hermitian.
 - The correct general converse theorem is that V has an orthonormal basis of eigenvectors of T if and only if T is a normal operator, meaning that $T^*T = TT^*$.

4.4.3 Stochastic Matrices and Markov Chains

- 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% change 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 stochastic matrix (or a transition matrix).
 - Equivalently, a square matrix M is a stochastic 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 stochastic 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 stochastic matrices:
- Theorem (Markov Chains): If M is a stochastic 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 1-eigenvector of 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.

4.4.4 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

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

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 $A = \begin{bmatrix} a_{1,1}(x) & \cdots & a_{1,n}(x) \\ \vdots & \ddots & \vdots \\ a_{n,1}(x) & \cdots & a_{n,n}(x) \end{bmatrix}$, $\mathbf{q} =$

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

- We say that the system is homogeneous if $\mathbf{q} = \mathbf{0}$, and it is nonhomogeneous otherwise.
- 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 (Existence-Uniqueness for Homogeneous Systems): If the $n \times n$ coefficient matrix A is constant and I is any interval containing a , then there exists a unique solution to the homogeneous initial value problem $\mathbf{y}' = A\mathbf{y}$ with $\mathbf{y}(a) = \mathbf{y}_0$ on I . As a consequence, the vector space of solutions to $\mathbf{y}' = A\mathbf{y}$ on I is an n -dimensional vector space.
 - This existence and uniqueness parts of the theorem are analytic in nature. The fact that the vector space of solutions is n -dimensional follows by noting that the existence-uniqueness statement implies that the vector space of solutions to $\mathbf{y}' = A\mathbf{y}$ is isomorphic to the vector space of possible initial condition vectors \mathbf{y}_0 , which is clearly n -dimensional.
 - We, of course, would actually like to write down the solutions explicitly. The key observation is that if \mathbf{v} is an eigenvector of A with eigenvalue λ , then $\mathbf{y} = e^{\lambda x}\mathbf{v}$ is a solution to $\mathbf{y}' = A\mathbf{y}$.
 - This follows simply by differentiating $\mathbf{y} = e^{\lambda x}\mathbf{v}$ with respect to x : we see $\mathbf{y}' = \lambda e^{\lambda x}\mathbf{v} = \lambda\mathbf{y} = A\mathbf{y}$.
 - 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. In fact, they will always give us a basis for the solution space:
 - 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 $\mathbf{y}' = A\mathbf{y}$ 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 $\mathbf{y}' = A\mathbf{y}$. We claim that they are a basis for the solution space. Since the solution space is n -dimensional, it suffices to show that these solutions are linearly independent.

- For this, we simply compute the determinant of the matrix W whose columns are these n vectors: after factoring out the exponentials from each column, we obtain $\det(W) = e^{(\lambda_1 + \dots + \lambda_n)x} \det(M)$, 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, $\det(W)$ 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 $\mathbf{y}' = A\mathbf{y}$ 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 is real but has nonreal eigenvalues, by taking an appropriate linear combination we can produce real-valued solution vectors.

- Explicitly, suppose A has a complex eigenvalue $\lambda = a + bi$ with associated eigenvector $\mathbf{v} = \mathbf{w}_1 + i\mathbf{w}_2$. Then $\bar{\lambda} = a - bi$ has an eigenvector $\bar{\mathbf{v}} = \mathbf{w}_1 - i\mathbf{w}_2$ (the conjugate of \mathbf{v}), so we obtain the two solutions $e^{\lambda x} \mathbf{v}$ and $e^{\bar{\lambda} x} \bar{\mathbf{v}}$ to the system $\mathbf{y}' = A\mathbf{y}$.

- Then the linear combinations $\frac{1}{2}(e^{\lambda x} \mathbf{v} + e^{\bar{\lambda} x} \bar{\mathbf{v}}) = e^{ax}(\mathbf{w}_1 \cos(bx) - \mathbf{w}_2 \sin(bx))$ and $\frac{1}{2i}(e^{\lambda x} \mathbf{v} - e^{\bar{\lambda} x} \bar{\mathbf{v}}) = e^{ax}(\mathbf{w}_1 \sin(bx) + \mathbf{w}_2 \cos(bx))$ are both real-valued.

- Thus, to obtain real-valued solutions, we can replace the two complex-valued solutions $e^{\lambda x} \mathbf{v}$ and $e^{\bar{\lambda} x} \bar{\mathbf{v}}$ with the two real-valued solutions $e^{ax}(\mathbf{w}_1 \cos(bx) - \mathbf{w}_2 \sin(bx))$ and $e^{ax}(\mathbf{w}_1 \sin(bx) + \mathbf{w}_2 \cos(bx))$, which are simply the real and imaginary parts of $e^{\lambda x} \mathbf{v}$ respectively.

- Example: Find all real-valued functions y_1 and y_2 such that
$$\begin{aligned} y_1' &= 3y_1 - 2y_2 \\ y_2' &= y_1 + y_2 \end{aligned}.$$

- The coefficient matrix is $A = \begin{bmatrix} 3 & -2 \\ 1 & 1 \end{bmatrix}$, whose characteristic polynomial is $\det(tI - A) = \begin{vmatrix} t-3 & 2 \\ -1 & t-1 \end{vmatrix} = t^2 - 4t + 5$ so the eigenvalues are $\lambda = 2 \pm i$.

- By row-reducing we see that the $(2+i)$ -eigenspace is spanned by $\begin{bmatrix} 1+i \\ 1 \end{bmatrix}$, while the $(2-i)$ -eigenspace is spanned by $\begin{bmatrix} 1-i \\ 1 \end{bmatrix}$.

- This yields the complex-valued basis $\begin{bmatrix} 1+i \\ 1 \end{bmatrix} e^{(2+i)x}$, $\begin{bmatrix} 1-i \\ 1 \end{bmatrix} e^{(2-i)x}$ for the solution space.

- We want real-valued solutions, so extracting the real and imaginary parts as above yields the equivalent basis $e^{2x} \left(\begin{bmatrix} 1 \\ 1 \end{bmatrix} \cos(x) - \begin{bmatrix} 1 \\ 0 \end{bmatrix} \sin(x) \right)$ and $e^{2x} \left(\begin{bmatrix} 1 \\ 1 \end{bmatrix} \sin(x) + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \cos(x) \right)$.

- The general real-valued solution is then
$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \boxed{C_1 e^{2x} \begin{bmatrix} \cos(x) - \sin(x) \\ \cos(x) \end{bmatrix} + C_2 e^{2x} \begin{bmatrix} \sin(x) + \cos(x) \\ \sin(x) \end{bmatrix}}.$$

- When the coefficient matrix is not diagonalizable, we can adapt the eigenvalue method to generate a basis for the solution space using chains of generalized eigenvectors.

- The main observation is the following: if $\{\mathbf{v}_{k-1}, \mathbf{v}_{k-2}, \dots, \mathbf{v}_1, \mathbf{v}_0\}$ is a chain of k generalized λ -eigenvectors above the λ -eigenvector \mathbf{v} , where $\mathbf{v}_i = (A - \lambda I)^i \mathbf{v}$ for each i , then $e^{\lambda x} \mathbf{v}_0, e^{\lambda x}(\mathbf{v}_1 + x\mathbf{v}_0), e^{\lambda x}(\mathbf{v}_2 + x\mathbf{v}_1 + \frac{x^2}{2}\mathbf{v}_0), \dots, e^{\lambda x}(\mathbf{v}_{k-1} + x\mathbf{v}_{k-2} + \dots + \frac{x^{k-2}}{(k-2)!}\mathbf{v}_1 + \frac{x^{k-1}}{(k-1)!}\mathbf{v}_0)$ yield k linearly independent solutions to the system $\mathbf{y}' = A\mathbf{y}$.

- To see this, observe that $\frac{d}{dx} \left[\frac{x^d}{d!} e^{\lambda x} \mathbf{v}_i \right] = \frac{x^{d-1}}{(d-1)!} e^{\lambda x} \mathbf{v}_i + \lambda \frac{x^d}{d!} e^{\lambda x} \mathbf{v}_i$, while $(A - \lambda I) \left[\frac{x^d}{d!} e^{\lambda x} \mathbf{v}_i \right] = \frac{x^d}{d!} e^{\lambda x} \mathbf{v}_{i-1}$ hence $A \left[\frac{x^d}{d!} e^{\lambda x} \mathbf{v}_i \right] = \frac{x^d}{d!} e^{\lambda x} \mathbf{v}_{i-1} + \lambda \frac{x^d}{d!} e^{\lambda x} \mathbf{v}_i$; then summing from $d = 0$ to $d = k - 1$ and reindexing the sum on the first term shows that $\frac{d}{dx} \left[\mathbf{v}_{k-1} + \dots + \frac{x^{k-1}}{(k-1)!} \mathbf{v}_0 \right]$ is equal to $A[\mathbf{v}_{k-1} + \dots + \frac{x^{k-1}}{(k-1)!} \mathbf{v}_0]$, as desired.

- These solutions are (trivially) linearly independent since $\{\mathbf{v}_{k-1}, \mathbf{v}_{k-2}, \dots, \mathbf{v}_1, \mathbf{v}_0\}$ is linearly independent, and each solution contains one more of the vectors \mathbf{v}_i than the previous solution.
- Since we can always find a basis of generalized eigenvectors, we can always construct a solution basis in this manner.

- Example: Find all functions y_1 and y_2 such that
$$\begin{aligned} y_1' &= 5y_1 - 9y_2 \\ y_2' &= 4y_1 - 7y_2 \end{aligned}$$

- The coefficient matrix is $A = \begin{bmatrix} 5 & -9 \\ 4 & -7 \end{bmatrix}$, whose characteristic polynomial is $\det(tI - A) = (t + 1)^2$, so the eigenvalues of A are $\lambda = 1, 1$.

- The 1-eigenspace is 1-dimensional and spanned by the eigenvector $\mathbf{v}_0 = \begin{bmatrix} 3 \\ 2 \end{bmatrix}$. To construct a chain above this vector, we solve $(A - \lambda I)\mathbf{v}_1 = \mathbf{v}_0$ to obtain a solution $\mathbf{v}_1 = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$.

- The procedure above then yields the two linearly independent solutions $\begin{bmatrix} 3 \\ 2 \end{bmatrix} e^{-t}$ and $\begin{bmatrix} 3 \\ 2 \end{bmatrix} t e^{-t} +$

$\begin{bmatrix} 2 \\ 1 \end{bmatrix} e^{-t}$, so the general solution to the system is
$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = C_1 \begin{bmatrix} 3 \\ 2 \end{bmatrix} e^{-x} + C_2 \left(\begin{bmatrix} 3 \\ 2 \end{bmatrix} x e^{-x} + \begin{bmatrix} 2 \\ 1 \end{bmatrix} e^{-x} \right)$$

- As a final remark, we will note that there exists a method known as variation of parameters for solving a non-homogeneous system of linear differential equations if the homogeneous system can be solved.

- Explicitly, suppose $\mathbf{y}_1, \dots, \mathbf{y}_n$ are the n linearly independent solutions to the homogeneous equation $\mathbf{y}' = A\mathbf{y}$ and we want to solve the nonhomogeneous equation $\mathbf{y}' = A\mathbf{y} + \mathbf{q}$ where $\mathbf{q} = (q_1, \dots, q_n)$.

- We look for functions $c_1(x), \dots, c_n(x)$ making $\tilde{\mathbf{y}} = c_1(x)\mathbf{y}_1 + \dots + c_n(x)\mathbf{y}_n$ a solution to the nonhomogeneous equation $\mathbf{y}' = A\mathbf{y} + \mathbf{q}$.

- Differentiating $\tilde{\mathbf{y}}$ via the product rule and using the fact that $\mathbf{y}'_i = A\mathbf{y}_i$ for each i yields $\tilde{\mathbf{y}}' = (c_1\mathbf{y}'_1 + \dots + c_n\mathbf{y}'_n) + (c'_1\mathbf{y}_1 + \dots + c'_n\mathbf{y}_n) = A(c_1\mathbf{y}_1 + \dots + c_n\mathbf{y}_n) + (c'_1\mathbf{y}_1 + \dots + c'_n\mathbf{y}_n) = A\tilde{\mathbf{y}} + (c'_1\mathbf{y}_1 + \dots + c'_n\mathbf{y}_n)$.

- Therefore, we simply need to take c'_1, \dots, c'_n to satisfy the equation $c'_1\mathbf{y}_1 + \dots + c'_n\mathbf{y}_n = \mathbf{q}$, which is merely a matrix equation $Y\mathbf{c}' = \mathbf{q}$ where Y is the matrix whose columns are $\mathbf{y}_1, \dots, \mathbf{y}_n$ and \mathbf{c}' is the column vector (c'_1, \dots, c'_n) .

- In fact, since Y is invertible because its columns $\mathbf{y}_1, \dots, \mathbf{y}_n$ are linearly independent, the system has a unique solution, and we may then integrate the resulting solution vector to obtain the functions c_1, \dots, c_n . Including the arbitrary constants of integration there, in fact, will give the general solution to the nonhomogeneous system $\mathbf{y}' = A\mathbf{y} + \mathbf{q}$.

- Example: Find all functions y_1 and y_2 such that
$$\begin{aligned} y_1' &= y_2 \\ y_2' &= -y_1 + \sec(x) \end{aligned}$$

- The coefficient matrix for the homogeneous system is $A = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$ with eigenvalues $\lambda = \pm i$. Row-reducing to find eigenvectors yields the complex-valued solution basis $\begin{bmatrix} -i \\ 1 \end{bmatrix} e^{ix}$, $\begin{bmatrix} i \\ 1 \end{bmatrix} e^{-ix}$, and then extracting real and imaginary parts yields the equivalent real-valued solution basis $\begin{bmatrix} \sin(x) \\ \cos(x) \end{bmatrix}$, $\begin{bmatrix} -\cos(x) \\ \sin(x) \end{bmatrix}$.
- We then want $\tilde{\mathbf{y}} = c_1(x) \begin{bmatrix} \sin(x) \\ \cos(x) \end{bmatrix} + c_2(x) \begin{bmatrix} -\cos(x) \\ \sin(x) \end{bmatrix}$ where $\begin{bmatrix} \sin(x) & -\cos(x) \\ \cos(x) & \sin(x) \end{bmatrix} \begin{bmatrix} c_1'(x) \\ c_2'(x) \end{bmatrix} = \begin{bmatrix} 0 \\ \sec(x) \end{bmatrix}$.
- Left-multiplying by $\begin{bmatrix} \sin(x) & \cos(x) \\ -\cos(x) & \sin(x) \end{bmatrix}$ yields $\begin{bmatrix} c_1'(x) \\ c_2'(x) \end{bmatrix} = \begin{bmatrix} \sin(x) & \cos(x) \\ -\cos(x) & \sin(x) \end{bmatrix} \begin{bmatrix} 0 \\ \sec(x) \end{bmatrix} = \begin{bmatrix} 1 \\ \tan(x) \end{bmatrix}$ and now taking antiderivatives yields $c_1(x) = C_1 + x$ and $c_2(x) = C_2 + \ln(\sec(x))$.
- The general solution is therefore $\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \boxed{(C_1 + x) \begin{bmatrix} \sin(x) \\ \cos(x) \end{bmatrix} + (C_2 + \ln(\sec(x))) \begin{bmatrix} -\cos(x) \\ \sin(x) \end{bmatrix}}$.

4.4.5 Matrix Exponentials and the Jordan Form

- There is also another, quite different, method for using diagonalization and the Jordan canonical form 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 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 $A \in M_{n \times n}(\mathbb{C})$, we define the exponential of A as the infinite sum $e^A = \sum_{n=0}^{\infty} A^n/n!$.
 - The definition is motivated by the Taylor series for the real exponential $e^z = \sum_{n=0}^{\infty} z^n/n!$, but in order for this definition to make sense, we need to know that the infinite sum actually converges.
- Theorem (Exponential Solutions): For any matrix A , the infinite series $\sum_{n=0}^{\infty} \sum_{n=0}^{\infty} A^n/n!$ converges absolutely, in the sense that the series in each of the entries of the matrix converges absolutely. Furthermore, the unique solution to the initial value problem $\mathbf{y}' = \mathbf{A}\mathbf{y}$ with $\mathbf{y}(a) = \mathbf{y}_0$ is given by $\mathbf{y}(t) = e^{A(t-a)}\mathbf{y}_0$.
 - Proof: Define the “matrix norm” $\|M\|$ to be the sum of the absolute values of the entries of M .
 - Observe that $\|A + B\| \leq \|A\| + \|B\|$ for any matrices A and B : this simply follows by applying the triangle inequality in each entry of $A + B$.
 - Likewise, we also have $\|AB\| \leq \|A\| \cdot \|B\|$ for any matrices A and B : this follows by observing that the entries of the product matrix are a sum of products of entries from A and entries from B and applying the triangle inequality.
 - Then $\left\| \sum_{n=0}^k \frac{A^n}{n!} \right\| \leq \sum_{n=0}^k \frac{\|A^n\|}{n!} \leq \sum_{n=0}^k \frac{\|A\|^n}{n!} \leq \sum_{n=0}^{\infty} \frac{\|A\|^n}{n!} = e^{\|A\|}$, so each entry in any partial sum of the infinite series $\sum_{n=0}^{\infty} \frac{A^n}{n!}$ has absolute value at most $e^{\|A\|}$.
 - Thus, the infinite series converges absolutely, so we can differentiate term-by-term to see that $\frac{d}{dx}[e^{Ax}] = \frac{d}{dx} \left[\sum_{n=0}^{\infty} \frac{A^n}{n!} x^n \right] = \sum_{n=0}^{\infty} \frac{A^n}{(n-1)!} x^{n-1} = A \left[\sum_{n=0}^{\infty} \frac{A^n}{n!} x^n \right] = Ae^{Ax}$.
 - Therefore, we see that $\mathbf{y}(t) = e^{A(t-a)}\mathbf{y}_0$ is a solution to the initial value problem (since it satisfies the differential equation and the initial condition). The uniqueness part of the existence-uniqueness theorem guarantees it is the only solution.
- The theorem above tells us that we can use matrix exponentials to write down the solutions of initial value problems. 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 $A = Q^{-1}DQ$, then $e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!} = \sum_{n=0}^{\infty} \frac{(Q^{-1}DQ)^n}{n!} = \sum_{n=0}^{\infty} \frac{Q^{-1}D^nQ}{n!} = Q^{-1} \left[\sum_{n=0}^{\infty} \frac{D^n}{n!} \right] Q = Q^{-1}e^DQ$.
- Furthermore, again from the power series definition, if D is diagonal with diagonal entries $\lambda_1, \dots, \lambda_n$, then e^D is diagonal with diagonal entries $e^{\lambda_1}, \dots, e^{\lambda_n}$.
- 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 $Q^{-1}AQ = D = \begin{bmatrix} 1 & 0 \\ 0 & 4 \end{bmatrix}$.
 - Thus, $e^{Ax} = Qe^{Dx}Q^{-1} = Q \begin{bmatrix} e^x & 0 \\ 0 & e^{4x} \end{bmatrix} Q^{-1} = \frac{1}{3} \begin{bmatrix} 2e^x + e^{4x} & e^x - e^{4x} \\ 2e^x - 2e^{4x} & e^x + 2e^{4x} \end{bmatrix}$.
 - Then $\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \frac{1}{3} \begin{bmatrix} 2e^x + e^{4x} & e^x - e^{4x} \\ 2e^x - 2e^{4x} & e^x + 2e^{4x} \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \end{bmatrix}$ for arbitrary constants C_1 and C_2 .
- If the matrix is not diagonalizable, we must use the Jordan canonical form. By the same calculation as given above for the diagonalization, it suffices to compute the exponential of each Jordan block separately.

- **Proposition (Exponential of Jordan Block):** We have $e^{Jx} = \begin{bmatrix} e^{\lambda x} & xe^{\lambda x} & \frac{x^2}{2}e^{\lambda x} & \dots & \frac{x^{d-1}}{(d-1)!}e^{\lambda x} \\ & e^{\lambda x} & xe^{\lambda x} & \ddots & \vdots \\ & & \ddots & \ddots & \frac{x^2}{2}e^{\lambda x} \\ & & & e^{\lambda x} & xe^{\lambda x} \\ & & & & e^{\lambda x} \end{bmatrix}$, where

J is the $d \times d$ Jordan block matrix with eigenvalue λ .

- **Proof:** Write $J = \lambda I + N$. As we showed earlier, N^d is the zero matrix, and $NI = IN$ since I is the identity matrix.
- Applying the binomial expansion yields $(Jx)^k = x^k(\lambda I + N)^k = x^k \left[\lambda^k I + \binom{k}{1} \lambda^{k-1} N^1 + \dots + \binom{k}{k-d} \lambda^{k-d} N^d + \dots \right]$, but since N^d is the zero matrix, only the terms up through N^{d-1} are nonzero. (Note that we are using the fact that $IN = NI$, since the binomial theorem does not hold for general matrices.)
- It is then a straightforward (if somewhat lengthy) computation to plug these expressions into the infinite sum defining e^{Jx} and evaluate the infinite sum to obtain the stated result.

- **Example:** Solve the system of linear differential equations $\mathbf{y}'(x) = \begin{bmatrix} 2 & 1 & 0 & 0 \\ 0 & 2 & 1 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{y}$, where $\mathbf{y}(0) = \begin{bmatrix} 1 \\ 2 \\ -4 \\ 3 \end{bmatrix}$.

- Observe that the coefficient matrix A is already in Jordan canonical form.

- Hence $e^{Ax} = \begin{bmatrix} e^{2x} & xe^{2x} & x^2e^{2x}/2 & 0 \\ 0 & e^{2x} & xe^{2x} & 0 \\ 0 & 0 & e^{2x} & 0 \\ 0 & 0 & 0 & e^x \end{bmatrix}$, so the solution is $\mathbf{y}(t) = e^{Ax} \begin{bmatrix} 1 \\ 2 \\ -4 \\ 3 \end{bmatrix} = \begin{bmatrix} e^{2x} + 2xe^{2x} + 2x^2e^{2x} \\ 2e^{2x} - 4xe^{2x} \\ -4e^{2x} \\ 3e^x \end{bmatrix}$.

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

Copyright notice: This material is copyright Evan Dummit, 2012-2023. You may not reproduce or distribute this material without my express permission.