

Contents

5 Eigenvalues and Diagonalization	1
5.1 Eigenvalues, Eigenvectors, and The Characteristic Polynomial	1
5.1.1 Eigenvalues and Eigenvectors	2
5.1.2 Eigenvalues and Eigenvectors of Matrices	3
5.1.3 Eigenspaces	6
5.2 Diagonalization	9
5.2.1 Criterion for Diagonalizability	10
5.2.2 Calculating Diagonalizations	12
5.2.3 The Spectral Theorem for Symmetric Matrices	14
5.3 Applications of Diagonalization	17
5.3.1 Transition Matrices and Markov Chains	17
5.3.2 Bilinear Forms and Quadratic Forms	19
5.3.3 Applications of Quadratic Forms, Definiteness	22
5.3.4 Singular Values and Singular Value Decomposition	25

5 Eigenvalues and Diagonalization

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

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

5.1 Eigenvalues, Eigenvectors, and The Characteristic Polynomial

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

5.1.1 Eigenvalues and Eigenvectors

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

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

5.1.2 Eigenvalues and Eigenvectors of Matrices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- By expanding along the top row,

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

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

5.1.3 Eigenspaces

- Using the characteristic polynomial, we can find all the eigenvalues of a matrix A without actually determining the associated eigenvectors. However, we often also want to find the eigenvectors associated to each eigenvalue.
- We might hope that there is a straightforward way to describe all the eigenvectors, and (conveniently) there is: the set of all eigenvectors with a particular eigenvalue λ has a vector space structure.
- **Proposition** (Eigenspaces): If $T : V \rightarrow V$ is linear, then for any fixed value of λ , the set E_λ of vectors in V satisfying $T(\mathbf{v}) = \lambda\mathbf{v}$ is a subspace of V . This space E_λ is called the eigenspace associated to the eigenvalue λ , or more simply the λ -eigenspace.
 - Notice that E_λ is precisely the set of eigenvectors with eigenvalue λ , along with the zero vector.
 - The eigenspaces for a matrix A are defined in the same way: E_λ is the space of vectors \mathbf{v} such that $A\mathbf{v} = \lambda\mathbf{v}$.
 - **Proof:** By definition, E_λ is the kernel of the linear transformation $\lambda I - T$, and is therefore a subspace of V .
- **Example:** Find the 1-eigenspaces, and their dimensions, for $A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$.
 - For the 1-eigenspace of A , we want to find all vectors with $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$.
 - Clearly, all vectors satisfy this equation, so the 1-eigenspace of A is the set of all vectors $\begin{bmatrix} a \\ b \end{bmatrix}$, and has dimension 2.
 - For the 1-eigenspace of B , we want to find all vectors with $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$, or equivalently, $\begin{bmatrix} a+b \\ b \end{bmatrix} = \begin{bmatrix} a \\ b \end{bmatrix}$.
 - The vectors satisfying the equation are those with $b = 0$, so the 1-eigenspace of B is the set of vectors of the form $\begin{bmatrix} a \\ 0 \end{bmatrix}$, and has dimension 1.
 - Notice that the characteristic polynomial of each matrix is $(t - 1)^2$, since both matrices are upper-triangular, and they both have a single eigenvalue $\lambda = 1$ of multiplicity 2. Nonetheless, the matrices do not have the same eigenvectors, and the dimensions of their 1-eigenspaces are different.
- To compute a basis for the λ -eigenspace we must solve 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 with basis $\begin{bmatrix} -2 \\ 3 \end{bmatrix}$.
 - 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 with basis $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.

- 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 with basis $\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$.

◦ 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 with basis $\begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix}$.

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

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

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

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

$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$, so the 0-eigenspace is 1-dimensional with basis $\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$.

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

$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & -i \\ 0 & 0 & 0 \end{bmatrix}$, so the i -eigenspace is 1-dimensional with basis $\begin{bmatrix} 0 \\ i \\ 1 \end{bmatrix}$.

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

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

- 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 $\overline{AB} = \bar{A}\bar{B}$.
- Thus, if $A\mathbf{v} = \lambda\mathbf{v}$, taking complex conjugates gives $\bar{A}\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$, and since $\bar{A} = A$ because A is a real matrix, we see $A\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$: thus, $\bar{\mathbf{v}}$ is an eigenvector with eigenvalue $\bar{\lambda}$.
- The second statement follows from the first, since complex conjugation does not affect linear independence or dimension.

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

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

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

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

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

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

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

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

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

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

- By expanding along the top row, we can compute

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

- To find the eigenvalues, we wish to solve the cubic equation $t^3 + 2t^2 - t - 2 = 0$.
- By the rational root test, if the polynomial has a rational root then it must be an integer dividing -2 : that is, one of $\pm 1, \pm 2$. Testing the possibilities reveals that $t = 1, t = -1$, and $t = -2$ are each roots, from which we obtain the factorization $(t - 1)(t + 1)(t + 2) = 0$.
- Thus, the eigenvalues are $t = -2, -1, 1$.
- We see that $\text{tr}(A) = 2 + (-1) + (-3) = -2$, while the sum of the eigenvalues is $(-2) + (-1) + 1 = -2$.
- Also, $\det(A) = 2$, and the product of the eigenvalues is $(-2)(-1)(1) = 2$.
- In all of the examples above, the dimension of each eigenspace was less than or equal to the multiplicity of the eigenvalue as a root of the characteristic polynomial. This is true in general:
- **Theorem** (Eigenvalue Multiplicity): If λ is an eigenvalue of the matrix A which appears exactly k times as a root of the characteristic polynomial, then the dimension of the eigenspace corresponding to λ is at least 1 and at most k .
 - **Remark**: The number of times that λ appears as a root of the characteristic polynomial is sometimes called the “algebraic multiplicity” of λ , and the dimension of the eigenspace corresponding to λ is sometimes called the “geometric multiplicity” of λ . In this language, the theorem above says that the geometric multiplicity is less than or equal to the algebraic multiplicity.
 - **Example**: If the characteristic polynomial of a matrix is $(t - 1)^3(t - 3)^2$, then the eigenspace for $\lambda = 1$ is at most 3-dimensional, and the eigenspace for $\lambda = 3$ is at most 2-dimensional.
 - **Proof**: The statement that the eigenspace has dimension at least 1 is immediate, because (by assumption) λ is a root of the characteristic polynomial and therefore has at least one nonzero eigenvector associated to it.
 - For the other statement, observe that the dimension of the λ -eigenspace is the dimension of the solution space of the homogeneous system $(\lambda I - A)\mathbf{x} = \mathbf{0}$. (Equivalently, it is the dimension of the nullspace of the matrix $\lambda I - A$.)
 - If λ appears k times as a root of the characteristic polynomial, then when we put the matrix $\lambda I - A$ into its reduced row-echelon form B , we claim that B must have at most k rows of all zeroes.
 - Otherwise, the matrix B (and hence $\lambda I - A$ too, since the nullity and rank of a matrix are not changed by row operations) would have 0 as an eigenvalue more than k times, because B is in echelon form and therefore upper-triangular.
 - But the number of rows of all zeroes in a square matrix in reduced row-echelon form is the same as the number of nonpivotal columns, which is the number of free variables, which is the dimension of the solution space.
 - So, putting all the statements together, we see that the dimension of the eigenspace is at most k .

5.2 Diagonalization

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

5.2.1 Criterion for 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 the linear transformation T given by left-multiplication by A .
 - We then would like to say that A is diagonalizable when T is diagonalizable.
 - By our results on change of basis, this is equivalent to saying that there exists an invertible matrix Q , namely the change-of-basis matrix $Q = [I]_{\gamma}^{\beta}$, for which $Q^{-1}AQ = [I]_{\gamma}^{\beta}[T]_{\gamma}^{\gamma}[I]_{\beta}^{\gamma} = [T]_{\beta}^{\beta}$ is a diagonal matrix.
- **Definition:** An $n \times n$ matrix A is diagonalizable if there exists an invertible $n \times n$ matrix Q for which $Q^{-1}AQ$ is a diagonal matrix.
 - Recall that we say two $n \times n$ matrices A and B are similar if there exists an invertible $n \times n$ matrix Q such that $B = Q^{-1}AQ$.
- Our goal is to study and then characterize diagonalizable linear transformations, which (per the above discussion) is equivalent to characterizing diagonalizable matrices.
- **Proposition** (Characteristic Polynomials and Similarity): If A and B are similar, then they have the same characteristic polynomial, determinant, trace, and eigenvalues (and their eigenvalues have the same multiplicities).
 - **Proof:** Suppose $B = Q^{-1}AQ$. For the characteristic polynomial, we simply compute $\det(tI - B) = \det(Q^{-1}(tI)Q - Q^{-1}AQ) = \det(Q^{-1}(tI - A)Q) = \det(Q^{-1}) \det(tI - A) \det(Q) = \det(tI - A)$.
 - The determinant and trace are both coefficients (up to a factor of ± 1) of the characteristic polynomial, so they are also equal.
 - Finally, the eigenvalues are the roots of the characteristic polynomial, so they are the same and occur with the same multiplicities for A and B .
- The eigenvectors for similar matrices are also closely related:
- **Proposition** (Eigenvectors and Similarity): If $B = Q^{-1}AQ$, then \mathbf{v} is an eigenvector of B with eigenvalue λ if and only if $Q\mathbf{v}$ is an eigenvector of A with eigenvalue λ .
 - **Proof:** Since Q is invertible, $\mathbf{v} = \mathbf{0}$ if and only if $Q\mathbf{v} = \mathbf{0}$. Now assume $\mathbf{v} \neq \mathbf{0}$.
 - First suppose \mathbf{v} is an eigenvector of B with eigenvalue λ . Then $A(Q\mathbf{v}) = Q(Q^{-1}AQ)\mathbf{v} = Q(B\mathbf{v}) = Q(\lambda\mathbf{v}) = \lambda(Q\mathbf{v})$, meaning that $Q\mathbf{v}$ is an eigenvector of A with eigenvalue λ .
 - Conversely, if $Q\mathbf{v}$ is an eigenvector of A with eigenvalue λ . Then $B\mathbf{v} = Q^{-1}A(Q\mathbf{v}) = Q^{-1}\lambda(Q\mathbf{v}) = \lambda(Q^{-1}Q\mathbf{v}) = \lambda\mathbf{v}$, so \mathbf{v} is an eigenvector of B with eigenvalue λ .
- **Corollary:** If $B = Q^{-1}AQ$, then the eigenspaces for B have the same dimensions as the eigenspaces for A .
- As we have essentially worked out already, diagonalizability is equivalent to the existence of a basis of eigenvectors:
- **Theorem** (Diagonalizability): A linear operator $T : V \rightarrow V$ is diagonalizable if and only if there exists a basis β of V consisting of eigenvectors of T .
 - **Proof:** First suppose that V has a basis of eigenvectors $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ with respective eigenvalues $\lambda_1, \dots, \lambda_n$. Then by hypothesis, $T(\mathbf{v}_i) = \lambda_i\mathbf{v}_i$, and so $[T]_{\beta}^{\beta}$ is the diagonal matrix with diagonal entries $\lambda_1, \dots, \lambda_n$.
 - Conversely, suppose T is diagonalizable and let $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be a basis such that $[T]_{\beta}^{\beta}$ is a diagonal matrix whose diagonal entries are $\lambda_1, \dots, \lambda_n$. Then by hypothesis, each \mathbf{v}_i is nonzero and $T(\mathbf{v}_i) = \lambda_i\mathbf{v}_i$, so each \mathbf{v}_i is an eigenvector of T .

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

5.2.2 Calculating 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 $p(x)$ 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 $p(A) = Q\mathbf{0}Q^{-1} = \mathbf{0}$, as claimed.
 - **Remark:** In the case where A is not diagonalizable, the proof of the Cayley-Hamilton theorem is more difficult².

5.2.3 The Spectral Theorem for Symmetric Matrices

- An important computational result is a result known as the real spectral theorem, which says that every real symmetric matrix is diagonalizable. We will prove this result and then give some of its applications.
- **Definition:** If A is an $n \times n$ matrix, we say A is symmetric if $A^T = A$.
 - We will also make use of the fact that if \mathbf{v} and \mathbf{w} are column vectors in \mathbb{R}^n , then we can express the dot product $\mathbf{v} \cdot \mathbf{w}$ as the matrix product $\mathbf{v}^T \mathbf{w}$.
 - In particular, note that $\mathbf{v}^T \mathbf{v} = \|\mathbf{v}\|^2$, so $\mathbf{v}^T \mathbf{v} = 0$ if and only if $\mathbf{v} = \mathbf{0}$.
- We will also use a few observations about orthogonal matrices:
- **Definition:** An $n \times n$ matrix U such that $U^{-1} = U^T$ (equivalently, $U^T U = I_n$) is called an orthogonal matrix.
 - Observe that an $n \times n$ matrix U is orthogonal if and only if its columns are an orthonormal basis for \mathbb{R}^n , since the dot product of its i th column with its j th column is the (i, j) -entry in $U^T U$.
- **Proposition** (Properties of Symmetric Matrices): Suppose A is an $n \times n$ symmetric real matrix. Then the following properties hold:
 1. Eigenvectors of A with different eigenvalues are orthogonal.
 - **Proof:** Suppose that $A\mathbf{v}_1 = \lambda_1\mathbf{v}_1$ and $A\mathbf{v}_2 = \lambda_2\mathbf{v}_2$ where $\lambda_1 \neq \lambda_2$.
 - Then, since $A = A^T$, we have $\lambda_2\mathbf{v}_1^T \mathbf{v}_2 = \mathbf{v}_1^T (\lambda_2\mathbf{v}_2) = \mathbf{v}_1^T (A\mathbf{v}_2) = \mathbf{v}_1^T A^T \mathbf{v}_2 = (A\mathbf{v}_1)^T \mathbf{v}_2 = (\lambda_1\mathbf{v}_1)^T \mathbf{v}_2 = \lambda_1\mathbf{v}_1^T \mathbf{v}_2$ since λ_1 is real.
 - But since $\lambda_1 \neq \lambda_2$, this means $\mathbf{v}_1^T \mathbf{v}_2 = 0$, which is to say, $\mathbf{v}_1 \cdot \mathbf{v}_2 = 0$, so \mathbf{v}_1 and \mathbf{v}_2 are orthogonal.
 2. All eigenvalues of A are real numbers.

²One way to proceed is to note that even if A is a non-diagonalizable matrix with real entries, it is the limit of a sequence of diagonalizable matrices (this is true because we may always perturb A slightly to make its eigenvalues unequal, which will yield a diagonalizable matrix). Then since the characteristic polynomial is a continuous function of the entries of the matrix, passing it through the limit shows that $p(A)$ is still zero even for non-diagonalizable A .

- Proof: Suppose $\lambda = a + bi$ is an eigenvalue of A with eigenvector $\mathbf{x} + \mathbf{y}i$ where \mathbf{x}, \mathbf{y} are real vectors.
 - Then since A is a real matrix, $\bar{\lambda} = a - bi$ is also an eigenvalue of A with eigenvector $\mathbf{x} - \mathbf{y}i$.
 - If λ is not real, then by (1), since $\lambda \neq \bar{\lambda}$ we see that $(\mathbf{x} + \mathbf{y}i) \cdot (\mathbf{x} - \mathbf{y}i) = 0$. Expanding out yields $\mathbf{x} \cdot \mathbf{x} - i(\mathbf{x} \cdot \mathbf{y}) + i(\mathbf{y} \cdot \mathbf{x}) - i^2(\mathbf{y} \cdot \mathbf{y}) = 0$, which yields $\mathbf{x} \cdot \mathbf{x} + \mathbf{y} \cdot \mathbf{y} = 0$. But since \mathbf{x} and \mathbf{y} are real vectors, this forces $\mathbf{x} = \mathbf{y} = \mathbf{0}$, which is impossible since $\mathbf{0}$ is not an eigenvector.
 - Therefore, λ must be real, as claimed.
3. The matrix A is diagonalizable, and so \mathbb{R}^n has a basis consisting of eigenvectors of A .
- Proof: Suppose that λ is an eigenvalue of A , which must be real by (2). Choose a unit eigenvector \mathbf{e}_1 and then extend it to an orthonormal basis $\beta = \{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n\}$ for \mathbb{R}^n .
 - Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the linear transformation with $T(\mathbf{v}) = A\mathbf{v}$ and let γ be the standard basis: then $A = [T]_\gamma^\gamma$.
 - If $S = [T]_\beta^\beta$, then since the columns $\mathbf{e}_1, \dots, \mathbf{e}_n$ of S are orthonormal, we see that S is an orthogonal matrix, so $S^{-1} = S^T$. By our results on change of basis, we have $[T]_\beta^\beta = SAS^{-1} = SAS^T$. Therefore, the transpose of $[T]_\beta^\beta$ is $(SAS^T)^T = (S^T)^T A^T S^T = SA^T S^T = SAS^T$, so it is also a symmetric matrix.
 - Now, since $T(\mathbf{e}_1) = \lambda \mathbf{e}_1$, the first column of $[T]_\beta^\beta$ is simply $(\lambda, 0, \dots, 0)^T$. But since $[T]_\beta^\beta$ is symmetric, it actually has the form $\begin{bmatrix} \lambda & 0 \\ 0 & B \end{bmatrix}$ for some $1 \times (n-1)$ matrix A and some symmetric $(n-1) \times (n-1)$ matrix B .
 - By iterating this argument repeatedly on the smaller matrix B (which is still symmetric), we eventually obtain a diagonalization of A .
 - Since the eigenvalues and eigenvectors of A are real, this also means that \mathbb{R}^n has a basis consisting of eigenvectors of A , as claimed.
- In fact, we can strengthen this argument to show that symmetric matrices are diagonalizable in a particularly nice way:
 - Theorem (Real Spectral Theorem): If A is a real symmetric matrix, then \mathbb{R}^n has an orthonormal basis β of eigenvectors of A . Therefore, A may be written in the form $A = UDU^{-1}$ where D is a diagonal matrix with real entries and U is an orthogonal matrix with real entries (i.e., satisfying $U^{-1} = U^T$).
 - More succinctly, this result says that real symmetric matrices are orthogonally diagonalizable.
 - Proof: As noted above, A has a basis of eigenvectors and is therefore diagonalizable.
 - To get the more specific statement here, start with a basis for each eigenspace, and then apply Gram-Schmidt, yielding an orthonormal basis for each eigenspace.
 - Since A 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 property (1) above, so we obtain an orthonormal basis β of eigenvectors.
 - If U is the matrix whose columns are the vectors in the orthonormal basis β , then as we have repeatedly noted, we have $A = UDU^{-1}$ where D is diagonal. Furthermore, because the columns of U are orthonormal, as remarked above this means U is orthogonal, so $U^{-1} = U^T$.
 - Remark: The set of eigenvalues of A is called the spectrum of A . The spectral theorem shows that the behavior of A on \mathbb{R}^n can be decomposed into its eigenspaces where A acts very simply (as scalar multiplication), with one piece coming from each piece of the spectrum. (This is the reason for the name of the theorem.)
 - Per the theorem above, if A is a real symmetric matrix, to find the decomposition $A = UDU^{-1}$ where D is diagonal and U is orthogonal, we simply take U to be the matrix whose columns are an orthonormal basis of eigenvectors of A and D to be the diagonal matrix with the corresponding eigenvalues.
 - Example: For $A = \begin{bmatrix} 3 & 6 \\ 6 & 8 \end{bmatrix}$, find a diagonal matrix D and an orthogonal matrix U such that $A = UDU^{-1}$.
 - First, we find the eigenvalues of A . The characteristic polynomial is $p(t) = \det(tI - A) = t^2 - 11t + 12 = (t + 1)(t - 12)$ so the eigenvalues are $\lambda = -1, 12$.

- Next, we find an orthonormal basis for each eigenspace. A short calculation shows that $\begin{bmatrix} -3 \\ 2 \end{bmatrix}$ is a basis for the (-1) -eigenspace and $\begin{bmatrix} 2 \\ 3 \end{bmatrix}$ is a basis for the 12-eigenspace, so normalizing these vectors yields an orthonormal basis of eigenvectors $\frac{1}{\sqrt{13}} \begin{bmatrix} -3 \\ 2 \end{bmatrix}, \frac{1}{\sqrt{13}} \begin{bmatrix} 2 \\ 3 \end{bmatrix}$.
 - Then the desired matrices are $D = \begin{bmatrix} -1 & 0 \\ 0 & 12 \end{bmatrix}$ and $U = \frac{1}{\sqrt{13}} \begin{bmatrix} -3 & 2 \\ 2 & 3 \end{bmatrix}$.
 - Indeed, since the columns of U are orthonormal, we see $U^{-1} = U^T = \frac{1}{\sqrt{13}} \begin{bmatrix} -3 & 2 \\ 2 & 3 \end{bmatrix}$, and then we can compute $UDU^{-1} = \frac{1}{\sqrt{13}} \begin{bmatrix} -3 & 2 \\ 2 & 3 \end{bmatrix} \begin{bmatrix} -1 & 0 \\ 0 & 12 \end{bmatrix} \frac{1}{\sqrt{13}} \begin{bmatrix} -3 & 2 \\ 2 & 3 \end{bmatrix} = \begin{bmatrix} 3 & 6 \\ 6 & 8 \end{bmatrix} = A$, as required.
- Example: For $A = \begin{bmatrix} 3 & 2 & -2 \\ 2 & 2 & 0 \\ -2 & 0 & 4 \end{bmatrix}$, find a diagonal matrix D and an orthogonal 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 $\lambda = 0, 3, 6$.
 - Next, we find an orthonormal basis for each eigenspace. As
 - * 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 .
 - A short calculation shows that $\begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix}$ is a basis for the 0-eigenspace, $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ is a basis for the 3-eigenspace, and $\begin{bmatrix} 2 \\ 1 \\ -2 \end{bmatrix}$ is a basis for the 6-eigenspace.
 - Normalizing yields an orthonormal basis of eigenvectors $\frac{1}{3} \begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix}, \frac{1}{3} \begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}, \frac{1}{3} \begin{bmatrix} 2 \\ 1 \\ -2 \end{bmatrix}$.
 - 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, 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 eigenvalues are 0 and 0, but the 0-eigenspace only has dimension 1.
 - The correct statement for complex matrices is that complex matrices with $\overline{A^T} = A$ are diagonalizable (i.e., we must use the conjugate-transpose rather than the conjugate).
 - Matrices with $\overline{A^T} = A$ are called Hermitian matrices, and satisfy analogues of the properties we listed above for symmetric matrices: their eigenvalues are all real, there is an orthonormal basis of \mathbb{C}^n consisting of eigenvectors of A , and A is diagonalizable via a unitary matrix U with $\overline{U^T} = U$.

5.3 Applications of Diagonalization

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

5.3.1 Transition Matrices and 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 transition matrix (or a stochastic matrix).
 - Equivalently, a square matrix M is a transition matrix precisely when $M^T \mathbf{v} = \mathbf{v}$, where \mathbf{v} is the column vector of all 1s.
 - From this description, we can see that \mathbf{v} is an eigenvector of M^T of eigenvalue 1, and since M^T and M have the same characteristic polynomial, we conclude that M has 1 as an eigenvalue.
 - If it were true that M were diagonalizable and every eigenvalue of M had absolute value less than 1 (except for the eigenvalue 1), then we could apply the same argument as we did in the example to conclude that the powers of M approached a limit.
 - Unfortunately, this is not true in general: the transition matrix $M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ has M^2 equal to the identity matrix, so odd powers of M are equal to M while even powers are equal to the identity. (In this case, the eigenvalues of M are 1 and -1 .)
 - Fortunately, the argument does apply to a large class of transition matrices:
- Theorem (Markov Chains): If M is a transition matrix, then every eigenvalue λ of M has $|\lambda| \leq 1$. Furthermore, if some power of M has all entries positive, then the only eigenvalue of M of absolute value 1 is $\lambda = 1$, and the 1-eigenspace has dimension 1. In such a case, the “matrix limit” $\lim_{k \rightarrow \infty} M^k$ exists and has all columns equal to a “steady-state” solution of the Markov chain whose transition matrix is M .
 - We will not prove this theorem, although most of the arguments (when M is diagonalizable) are similar to the computations we did in the example above.
- Example: In a certain country, there are two cities, City A and City B. Each year, $1/5$ of the residents of City A move to City B, and $1/4$ of the residents of City B move to City A; the remaining residents stay in their current city. If in year 0 the populations of Cities A and B are 9000 and 18000 residents respectively, find the cities’ populations in year n , and identify the behavior as $n \rightarrow \infty$.
 - If the populations of the cities are currently $\begin{bmatrix} a \\ b \end{bmatrix}$, then the populations one year later are $\begin{bmatrix} 4/5 & 1/4 \\ 1/5 & 3/4 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix}$.
 - The given information implies that the populations are given by the Markov process whose transition matrix is $M = \begin{bmatrix} 4/5 & 1/4 \\ 1/5 & 3/4 \end{bmatrix}$, so we can analyze the system by diagonalizing M .
 - The characteristic polynomial of M is $p(t) = t^2 - \frac{31}{20}t + \frac{11}{20} = (t - 1)(t - \frac{11}{20})$, so the eigenvalues are 1 and $\frac{11}{20}$, with respective eigenspaces spanned by $\begin{bmatrix} 5/4 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$.
 - So with $Q = \begin{bmatrix} 5/4 & -1 \\ 1 & 1 \end{bmatrix}$ we get $M^n = Q \begin{bmatrix} 1 & 0 \\ 0 & (11/20)^n \end{bmatrix} Q^{-1} = \frac{1}{9} \begin{bmatrix} 5 + 4(11/20)^n & 5 - 5(11/20)^n \\ 4 - 4(11/20)^n & 4 + 5(11/20)^n \end{bmatrix}$.
 - Then the populations in year n are $M^n \begin{bmatrix} 3000 \\ 6000 \end{bmatrix} = \begin{bmatrix} 15000 - 6000(11/20)^n \\ 12000 + 6000(11/20)^n \end{bmatrix}$. So as $n \rightarrow \infty$, the populations approach 15000 in city A and 12000 in city B.
 - Note that the limiting behavior of the populations approaches the steady-state solution $\begin{bmatrix} 15000 \\ 12000 \end{bmatrix}$, as predicted by the theorem.

5.3.2 Bilinear Forms and Quadratic Forms

- Up until now, we have exclusively focused our attention on linear phenomena (i.e., linear transformations and matrices). But we can also use many of the ideas we have developed to study quadratic phenomena as well.
 - In fact, we have already encountered a general example of a quadratic function: if V is an inner product space, then $Q(\mathbf{v}) = \langle \mathbf{v}, \mathbf{v} \rangle = \|\mathbf{v}\|^2$ behaves like a quadratic function in many instances.
 - For example, for $V = \mathbb{R}^n$ with the dot product, we can see $Q(x_1, \dots, x_n) = (x_1, x_2, \dots, x_n) \cdot (x_1, x_2, \dots, x_n) = x_1^2 + x_2^2 + \dots + x_n^2$ is clearly quadratic.
 - Some other examples of quadratic functions of this general type are $Q(x, y) = x^2 + 3xy + 2y^2$ and $Q(x, y, z) = xz + yz$.
- The precise definition of a quadratic form is a bit more complicated than our definition of a linear transformation, but it relies on a type of pairing that is quite similar to an inner product:
- **Definition:** If V is a (real) vector space, a function $\Phi : V^2 \rightarrow \mathbb{R}$ is a symmetric bilinear form on V if it satisfies the following two properties:
 - [B1]** Linearity in both arguments: $\Phi(\mathbf{v}_1 + c\mathbf{v}_2, \mathbf{w}) = \Phi(\mathbf{v}_1, \mathbf{w}) + c\Phi(\mathbf{v}_2, \mathbf{w})$ and $\Phi(\mathbf{v}, \mathbf{w}_1 + c\mathbf{w}_2) = \Phi(\mathbf{v}, \mathbf{w}_1) + c\Phi(\mathbf{v}, \mathbf{w}_2)$ for any scalar c .
 - [B2]** Symmetry: $\Phi(\mathbf{v}, \mathbf{w}) = \Phi(\mathbf{w}, \mathbf{v})$.
 - Notice that a symmetric bilinear form is simply an inner product with the positive-definiteness axiom [I3] removed. Therefore, any inner product $\langle \cdot, \cdot \rangle$ on V is automatically a symmetric bilinear form on V .
 - **Example:** The dot product $\Phi(\mathbf{v}, \mathbf{w}) = \mathbf{v} \cdot \mathbf{w}$ on \mathbb{R}^n is a symmetric bilinear form, as is the inner product $\Phi(f, g) = \int_a^b f(x)g(x) dx$ on the space of continuous functions on the interval $[a, b]$.
- A large class of examples of symmetric bilinear forms arise as follows: if $V = \mathbb{R}^n$, then for any symmetric matrix $A \in M_{n \times n}(\mathbb{R})$, the map $\Phi_A(\mathbf{v}, \mathbf{w}) = \mathbf{v}^T A \mathbf{w}$ is a bilinear form on V .
 - **Example:** The matrix $A = \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix}$ yields the bilinear form $\Phi_A \left(\begin{bmatrix} x_1 \\ y_1 \end{bmatrix}, \begin{bmatrix} x_2 \\ y_2 \end{bmatrix} \right) = [x_1 \ y_1] \begin{bmatrix} 1 & 2 \\ 2 & 4 \end{bmatrix} \begin{bmatrix} x_2 \\ y_2 \end{bmatrix} = x_1x_2 + 2x_1y_2 + 2x_2y_1 + 4y_1y_2$.
 - It is not hard to check that this function $\Phi_A(\mathbf{v}, \mathbf{w})$ is linear in both \mathbf{v} and \mathbf{w} : we have $\Phi_A(\mathbf{v}_1 + c\mathbf{v}_2, \mathbf{w}) = (\mathbf{v}_1 + c\mathbf{v}_2)^T A \mathbf{w} = \mathbf{v}_1^T A \mathbf{w} + c\mathbf{v}_2^T A \mathbf{w} = \Phi_A(\mathbf{v}_1, \mathbf{w}) + c\Phi_A(\mathbf{v}_2, \mathbf{w})$, and similarly Φ_A is linear in the second argument.
 - For symmetry, we have $\Phi_A(\mathbf{w}, \mathbf{v}) = \mathbf{w}^T A \mathbf{v} = (\mathbf{v}^T A^T \mathbf{w})^T = (\mathbf{v}^T A \mathbf{w})^T = \mathbf{v}^T A \mathbf{w}$ where we used the facts that $A^T = A$ since A is symmetric and that the matrix $(\mathbf{v}^T A^T \mathbf{w})^T$ equals its own transpose since it is a 1×1 matrix.
 - In fact, one may show that every symmetric bilinear form on \mathbb{R}^n is of the form $\Phi_A(\mathbf{v}, \mathbf{w}) = \mathbf{v}^T A \mathbf{w}$, in the same way that every linear transformation on \mathbb{R}^n is of the form $T(\mathbf{v}) = A\mathbf{v}$ for an appropriate $n \times n$ matrix A .
- Like with linear transformations, we may also associate matrices to bilinear forms:
- **Definition:** If V is a finite-dimensional vector space, $\beta = \{\beta_1, \dots, \beta_n\}$ is a basis of V , and Φ is a symmetric bilinear form on V , the associated matrix of Φ with respect to β is the matrix $[\Phi]_\beta \in M_{n \times n}(F)$ whose (i, j) -entry is the value $\Phi(\beta_i, \beta_j)$.
 - This is the natural analogue of the matrix associated to a linear transformation, for bilinear forms.
- **Example:** For the bilinear form $\Phi((a, b), (c, d)) = 2ac + 4ad + 4bc - bd$ on \mathbb{R}^2 , find $[\Phi]_\beta$ for the standard basis $\beta = \{(1, 0), (0, 1)\}$.
 - We simply calculate the four values $\Phi(\beta_i, \beta_j)$ for $i, j \in \{1, 2\}$, where $\beta_1 = (1, 0)$ and $\beta_2 = (0, 1)$.
 - This yields $\Phi(\beta_1, \beta_1) = 2$, $\Phi(\beta_1, \beta_2) = 4$, $\Phi(\beta_2, \beta_1) = 4$, and $\Phi(\beta_2, \beta_2) = -1$.

◦ Thus, the associated matrix is $[\Phi]_\beta = \begin{bmatrix} 2 & 4 \\ 4 & -1 \end{bmatrix}$.

• Example: For the bilinear form $\Phi(p, q) = \int_0^1 p(x)q(x) dx$ on $P_2(\mathbb{R})$, find $[\Phi]_\beta$ for the basis $\beta = \{1, x, x^2\}$.

◦ We simply calculate the nine values $\Phi(\beta_i, \beta_j)$ for $i, j \in \{1, 2, 3\}$, where $\beta_1 = 1, \beta_2 = x, \beta_3 = x^2$.

◦ For example, $\Phi(\beta_1, \beta_3) = \int_0^1 1 \cdot x^2 dx = \frac{1}{3}$ and $\Phi(\beta_3, \beta_2) = \int_0^1 x^2 \cdot x dx = \frac{1}{4}$.

◦ The resulting associated matrix is $[\Phi]_\beta = \begin{bmatrix} 1 & 1/2 & 1/3 \\ 1/2 & 1/3 & 1/4 \\ 1/3 & 1/4 & 1/5 \end{bmatrix}$.

• Like with matrices associated with linear transformations, we can describe how the matrices associated to bilinear forms relate to coordinate vectors and how they change under a change of basis:

• Proposition (Associated Matrices): Suppose that V is a finite-dimensional vector space, $\beta = \{\beta_1, \dots, \beta_n\}$ is an ordered basis of V , and Φ is a bilinear form on V . Then the following hold:

1. If \mathbf{v} and \mathbf{w} are any vectors in V , then $\Phi(\mathbf{v}, \mathbf{w}) = [\mathbf{v}]_\beta^T [\Phi]_\beta [\mathbf{w}]_\beta$.

◦ Proof: If $\mathbf{v} = \beta_i$ and $\mathbf{w} = \beta_j$ then the result follows immediately from the definition of matrix multiplication and the matrix $[\Phi]_\beta$. The result for arbitrary \mathbf{v} and \mathbf{w} then follows by linearity.

2. The association $\Phi \mapsto [\Phi]_\beta$ of a symmetric bilinear form with its matrix representation yields an isomorphism of the space $\mathcal{B}(V)$ of bilinear forms on V with the symmetric $n \times n$ matrices.

◦ Proof: The inverse map is defined as follows: given a symmetric matrix A , define a bilinear form Φ_A via $\Phi_A(\mathbf{v}, \mathbf{w}) = [\mathbf{v}]_\beta^T A [\mathbf{w}]_\beta$.

◦ It is easy to verify that this map is a well-defined linear transformation and that it is inverse to the map given above.

3. If γ is another basis of V and $P = [I]_\beta^\gamma$ is the change-of-basis matrix from β to γ , then $[\Phi]_\gamma = P^T [\Phi]_\beta P$.

◦ Proof: By definition, $[\mathbf{v}]_\gamma = P [\mathbf{v}]_\beta$. Hence $[\mathbf{v}]_\gamma^T P^T [\Phi]_\beta P [\mathbf{w}]_\beta = [\mathbf{v}]_\beta^T [\Phi]_\beta [\mathbf{w}]_\beta$.

◦ This means that $P^T [\Phi]_\beta P$ and $[\Phi]_\gamma$ agree, as bilinear forms, on all pairs of vectors $[\mathbf{v}]_\beta$ and $[\mathbf{w}]_\beta$ in \mathbb{R}^n , so they are equal.

• Now that we have discussed bilinear forms, we can define their associated quadratic forms:

• Definition: If Φ is a symmetric bilinear form on V , the function $Q : V \rightarrow \mathbb{R}$ given by $Q(\mathbf{v}) = \Phi(\mathbf{v}, \mathbf{v})$ is called the quadratic form associated to Φ . If β is a basis of V , then the associated matrix of Q with respect to β is the matrix $[\Phi]_\beta$ associated to the bilinear form Φ .

◦ From the definition, we can see that $Q(r\mathbf{v}) = \Phi(r\mathbf{v}, r\mathbf{v}) = r^2 \Phi(\mathbf{v}, \mathbf{v}) = r^2 Q(\mathbf{v})$: thus, Q scales quadratically with its input vector (thereby justifying the name “quadratic form”).

◦ Also, by setting $\alpha = 0$ we see $Q(\mathbf{0}) = 0$, and by setting $\alpha = -1$ we see $Q(-\mathbf{v}) = Q(\mathbf{v})$.

◦ Example: If Φ is an inner product, then the associated quadratic form $Q(\mathbf{v}) = \langle \mathbf{v}, \mathbf{v} \rangle = \|\mathbf{v}\|^2$ is the square of the norm of \mathbf{v} .

• We are mainly interested in the situation where $V = \mathbb{R}^n$, in which case we have the quadratic form $Q_A(\mathbf{v}) = \mathbf{v}^T A \mathbf{v}$ associated to a symmetric $n \times n$ matrix A .

◦ If β is the standard basis of \mathbb{R}^n , then the matrix associated to Q_A with respect to β is simply A .

◦ When we multiply out the product $Q_A(\mathbf{v}) = \mathbf{v}^T A \mathbf{v}$ where $\mathbf{v} = [x_1 \ x_2 \ \dots \ x_n]^T$, we obtain a sum of the form $Q(x_1, x_2, \dots, x_n) = \sum_{1 \leq i \leq j \leq n} c_{i,j} x_i x_j$.

◦ The associated matrix A then has entries $a_{i,j} = a_{j,i} = \begin{cases} c_{i,i} & \text{for } i = j \\ c_{i,j}/2 & \text{for } i \neq j \end{cases}$, as can be seen by comparing coefficients for each monomial $x_i x_j$.

- Example: For $A = \begin{bmatrix} 7 & -2 \\ -2 & 3 \end{bmatrix}$, the associated quadratic form on \mathbb{R}^2 is $Q_A(x, y) = [x \ y] \begin{bmatrix} 7 & -2 \\ -2 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = 7x^2 - 4xy + 3y^2$.
- Example: For $B = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & \frac{3}{2} \\ 1 & \frac{3}{2} & 4 \end{bmatrix}$, the associated quadratic form is $Q_B(x, y, z) = [x \ y \ z] \begin{bmatrix} 1 & 0 & 1 \\ 0 & 0 & \frac{3}{2} \\ 1 & \frac{3}{2} & 4 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = x^2 + 2xz + 3yz + 4z^2$.

- Example: Find the associated matrix M for the quadratic form $Q(x, y, z) = x^2 - 4xy + 6xz + 9y^2 - yz + 8z^2$ on \mathbb{R}^3 .

- We simply read off the entries of the matrix from the coefficients: the diagonal entries are the coefficients of x^2, y^2, z^2 while the off-diagonal entries are half the corresponding coefficients of xy, xz, yz .

- We obtain the matrix $M = \begin{bmatrix} 1 & -2 & 3 \\ -2 & 9 & -1/2 \\ 3 & -1/2 & 8 \end{bmatrix}$.

- Indeed, we can check that $[x \ y \ z] \begin{bmatrix} 1 & -2 & 3 \\ -2 & 9 & -1/2 \\ 3 & -1/2 & 8 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = [x \ y \ z] \begin{bmatrix} x - 2y + 3z \\ -2x + 9y - z/2 \\ 3x - y/2 + 8z \end{bmatrix} = x^2 - 4xy + 6xz + 9y^2 - yz + 8z^2$ as required.

- Example: Find the quadratic form $Q(a, b, c, d)$ on \mathbb{R}^4 associated to the matrix $C = \begin{bmatrix} 3 & 0 & 1 & -2 \\ 0 & 4 & 1 & 8 \\ 1 & 1 & 0 & 0 \\ -2 & 8 & 0 & 2 \end{bmatrix}$.

- We can read off the coefficients of each of the monomials from the entries of the matrix: the diagonal entries are the coefficients of a^2, b^2, c^2, d^2 , while the off-diagonal entries are half the corresponding coefficients of ab, ac, ad, \dots, cd .

- We obtain the quadratic form $Q(a, b, c, d) = 3a^2 + 2ac - 4ad + 4b^2 + 2bc + 16bd + 2d^2$.

- Alternatively, we could just multiply out the product $[a \ b \ c \ d] \begin{bmatrix} 3 & 0 & 1 & -2 \\ 0 & 4 & 1 & 8 \\ 1 & 1 & 0 & 0 \\ -2 & 8 & 0 & 2 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix}$, which yields $3a^2 + 2ac - 4ad + 4b^2 + 2bc + 16bd + 2d^2$ after simplification.

- When working with quadratic forms, much as when we are working with linear transformations, we often want to express them in as simple a form as possible. Pleasantly, quadratic forms behave quite nicely under change of basis.

- Explicitly, suppose we are considering the quadratic form $Q_A(\mathbf{v}) = \mathbf{v}^T A \mathbf{v}$, where we think of A as being described in terms of the standard basis β of \mathbb{R}^n .

- If γ is another basis of V and $P = [I]_\beta^\gamma$ is the change-of-basis matrix from β to γ , then as shown above, we have $[\Phi]_\gamma = P^T [\Phi]_\beta P = P^T A P$.

- Therefore, after we change basis from β to γ , we obtain the quadratic form $Q_B(\mathbf{v}) = \mathbf{v}^T B \mathbf{v}$ where $B = P^T A P$.

- The most convenient possible outcome would be if we can change basis to make the resulting matrix $B = P^T A P$ diagonal. Quite pleasantly, the spectral theorem guarantees that we can always do this!

- Explicitly, since A is symmetric, the spectral theorem says that it is orthogonally diagonalizable; in other words, there exists an orthogonal matrix Q (with $Q^T = Q^{-1}$) such that $QSQ^{-1} = D$ is diagonal.

- But since $Q^T = Q^{-1}$, if we take $P = Q^T$ then this condition is the same as saying $P^T A P = B$ is diagonal.

- Therefore, we may always change basis to diagonalize a symmetric bilinear form over \mathbb{R} by computing the (regular) diagonalization of its associated matrix A , which is quite efficient as it only requires finding the eigenvalues and eigenvectors.
- The corresponding diagonalization represents “completing the square” in the quadratic form via a change of variables that is orthogonal (i.e., arises from an orthonormal basis), which corresponds geometrically to a rotation of the standard coordinate axes, possibly also with a reflection.
- **Example:** Find the matrix associated to the quadratic form $Q(x, y) = 4x^2 - 4xy + 7y^2$, and also find an orthonormal basis of \mathbb{R}^2 that diagonalizes Q .

- We can read off the associated matrix from the coefficients as $A = \begin{bmatrix} 4 & -2 \\ -2 & 7 \end{bmatrix}$.
- To diagonalize Q , we diagonalize A by finding the eigenvalues and eigenvectors of A .
- The characteristic polynomial is $p(t) = (t-4)(t-7) - 4 = t^2 - 11t + 24 = (t-3)(t-8)$, so the eigenvalues are $\lambda = 3, 8$.
- We can compute eigenvectors $(2, 1)$ and $(-1, 2)$ for $\lambda = 3, 8$ respectively, so upon normalizing these eigenvectors, we see that we can take $\gamma = \left\{ \frac{1}{\sqrt{5}}(2, 1), \frac{1}{\sqrt{5}}(-1, 2) \right\}$
- Explicitly, with $Q = \frac{1}{\sqrt{5}} \begin{bmatrix} 2 & -1 \\ 1 & 2 \end{bmatrix}$, we have $[\Phi]_\gamma = Q^T [\Phi]_\beta Q = \begin{bmatrix} 3 & 0 \\ 0 & 8 \end{bmatrix}$.
- In terms of the quadratic form, this says for $x' = \frac{1}{\sqrt{5}}(2x + y)$ and $y' = \frac{1}{\sqrt{5}}(-x + 2y)$, we have $Q(x, y) = 4x^2 - 4xy + 7y^2 = 3(x')^2 + 8(y')^2$. Note that by changing basis in this manner, we have eliminated the cross-term $-4xy$ in the original quadratic form Q .

- **Example:** Find an orthogonal change of basis that diagonalizes the quadratic form $Q(x, y, z) = 5x^2 + 4xy + 6y^2 + 4yz + 7z^2$ over \mathbb{R}^3 .

- We simply diagonalize the matrix for the corresponding bilinear form, which is $A = \begin{bmatrix} 5 & 2 & 0 \\ 2 & 6 & 2 \\ 0 & 2 & 7 \end{bmatrix}$. The characteristic polynomial is $p(t) = \det(tI_3 - A) = t^3 - 18t^2 + 99t - 162 = (t-3)(t-6)(t-9)$, so the eigenvalues are $\lambda = 3, 6, 9$.
- Computing a basis for each eigenspace yields eigenvectors $\begin{bmatrix} 2 \\ -2 \\ 1 \end{bmatrix}$, $\begin{bmatrix} -2 \\ -1 \\ 2 \end{bmatrix}$, $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ for $\lambda = 3, 6, 9$.
- Hence we may take $Q = \frac{1}{3} \begin{bmatrix} 2 & -2 & 1 \\ -2 & -1 & 2 \\ 1 & 2 & 2 \end{bmatrix}$, so that $Q^T = Q^{-1}$ and $QAQ^{-1} = \begin{bmatrix} 3 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & 9 \end{bmatrix} = D$.
- Therefore the desired change of basis is $x' = \frac{1}{3}(2x - 2y + z)$, $y' = \frac{1}{3}(-2x - y + 2z)$, $z' = \frac{1}{3}(x + 2y + 2z)$, and with this change of basis it is not hard to verify that, indeed, $Q(x, y, z) = 3(x')^2 + 6(y')^2 + 9(z')^2$.

5.3.3 Applications of Quadratic Forms, Definiteness

- One application of diagonalizing quadratic forms is that it allows us to describe the shape of the graph of the quadratic form.
 - For conics in \mathbb{R}^2 , the general equation is $Ax^2 + Bxy + Cy^2 + Dx + Ey + F = 0$. By diagonalizing, we may eliminate the xy term, and so the quadratic term can be put into the form $Ax^2 + Cy^2$. We then have various cases depending on the signs of A and C .
 - If A and C are both zero then the conic degenerates to a line. If one is zero and the other is not, then by rescaling and swapping variables we may assume $A = 1$ and $C = 0$, in which case the equation $x^2 + Dx + Ey + F = 0$ yields a parabola upon solving for y .

- If both A, C are nonzero, then we may complete the square to eliminate the linear terms, and then rescale so that $F = -1$. The resulting equation then has the form $A'x^2 + C'y^2 = 1$. If A', C' have the same sign, then the curve is an ellipse, while if A', C' have the opposite sign, the curve is a hyperbola.
- **Example:** Diagonalize the quadratic form $Q(x, y) = 2x^2 - 4xy - y^2$. Use the result to describe the shape of the conic section $Q(x, y) = 1$ in \mathbb{R}^2 .
 - The matrix associated to the corresponding bilinear form is $A = \begin{bmatrix} 2 & -2 \\ -2 & -1 \end{bmatrix}$.
 - The characteristic polynomial is $p(t) = \det(tI_2 - A) = t^3 - t + 6$ with eigenvalues $\lambda = 3, -2$.
 - We need to diagonalize A using an orthonormal basis of eigenvectors. Since the eigenvalues are distinct, we simply compute a basis for each eigenspace: doing so yields eigenvectors $(-2, 1)$ and $(1, 2)$ for $\lambda = 3, -2$ respectively.
 - Thus, we may diagonalize A via the orthogonal matrix $Q = \frac{1}{\sqrt{5}} \begin{bmatrix} -2 & 1 \\ 1 & 2 \end{bmatrix}$, and the resulting diagonalization is $Q(x, y, z) = 3(x')^2 - 2(y')^2$.
 - In particular, since the change of basis is orthonormal, in the new coordinate system the equation $Q(x, y, z) = 1$ reads simply as $3(x')^2 - 2(y')^2 = 1$, which is the graph of a hyperbola.
- One of the main properties of a quadratic form that determines its behavior is whether it takes positive values, negative values, or both:
- **Definition:** A quadratic form Q is positive definite if $Q(\mathbf{v}) > 0$ for every nonzero vector $\mathbf{v} \in V$, it is negative definite if $Q(\mathbf{v}) < 0$ for every nonzero vector $\mathbf{v} \in V$, and it is indefinite if Q takes both positive and negative values.
 - **Example:** If V is a real inner product space, then the square of the norm function $\|\mathbf{v}\|^2 = \langle \mathbf{v}, \mathbf{v} \rangle$ is a positive-definite quadratic form on V . Indeed, it is not hard to see that the underlying bilinear pairing Φ associated with Q is an inner product precisely when Q is a positive-definite quadratic form.
 - **Example:** The quadratic form $Q(x, y) = x^2 + 2y^2$ is positive definite, since $Q(x, y) > 0$ for all $(x, y) \neq (0, 0)$.
 - **Example:** The quadratic form $Q(x, y, z) = -2x^2 - 2xy - 5y^2 = -(x - y)^2 - (x + 2y)^2$ is negative definite, since the second expression shows that $Q(x, y) < 0$ for all $(x, y) \neq (0, 0)$.
 - **Example:** The quadratic form $Q(x, y) = xy$ is indefinite, since $Q(1, 1) = 1$ and $Q(-1, 1) = -1$, so Q takes both positive and negative values.
 - There are also a moderately useful weaker versions of these conditions: we say Q is positive semidefinite if $Q(\mathbf{v}) \geq 0$ for all $\mathbf{v} \in V$ and negative semidefinite if $Q(\mathbf{v}) \leq 0$ for all $\mathbf{v} \in V$.
 - **Example:** The quadratic form $Q(x, y) = x^2$ is positive semidefinite, since $Q(x, y) \geq 0$ for all (x, y) , but it is not positive definite because $Q(0, 1) = 0$.
 - It is easy to see that Q is positive (semi)definite if and only if $-Q$ is negative (semi)definite, so for example by the above we see that $Q(x, y) = -x^2$ is negative semidefinite.
- By diagonalizing a quadratic form, we can easily determine its definiteness:
- **Proposition (Definiteness and Eigenvalues):** If Q is a quadratic form on a finite-dimensional vector space V with associated matrix A , then Q is positive definite if and only if all eigenvalues of A are positive, Q is positive semidefinite if and only if all eigenvalues of A are nonnegative, Q is negative definite if and only if all eigenvalues of A are negative, Q is negative semidefinite if and only if all eigenvalues of A are nonpositive, and Q is indefinite if and only if it has both a positive and a negative eigenvalue.
 - **Proof:** Observe that definiteness is unaffected by changing basis, because each of the definiteness conditions $Q(\mathbf{v}) > 0$, $Q(\mathbf{v}) \geq 0$, $Q(\mathbf{v}) < 0$, $Q(\mathbf{v}) \leq 0$ are statements about all vectors \mathbf{v} in the vector space V .
 - Therefore, we may diagonalize Q without affecting its definiteness. After diagonalizing, we have an expression of the form $Q(x_1, x_2, \dots, x_n) = a_1x_1^2 + a_2x_2^2 + \dots + a_nx_n^2$.

- If any of the coefficients are negative, then Q necessarily takes negative values (specifically, if $a_i < 0$, then Q evaluated at the i th standard unit vector will be a_i).
- Likewise, if any of the coefficients are positive then Q necessarily takes positive values, and if any coefficients are zero or have opposite signs then Q takes the value 0 at some nonzero vector.
- Assuming we use an orthogonal diagonalization, then since the coefficients a_i are simply the eigenvalues of A , all of the claimed statements then follow immediately.
- Explicitly, if Q takes only positive values on nonzero vectors then no coefficients a_i can be zero or negative (so they are all positive), if Q takes only nonnegative values then no coefficients a_i can be negative, and likewise in the other two cases.

• **Example:** Determine the definiteness of the quadratic form $Q(x, y) = 2x^2 + 4xy + 5y^2$ on \mathbb{R}^2 .

- The associated matrix is $\begin{bmatrix} 2 & 2 \\ 2 & 5 \end{bmatrix}$, whose characteristic polynomial is $p(t) = (t - 2)(t - 5) - (2)(2) = t^2 - 7t + 6 = (t - 6)(t - 1)$, so its eigenvalues are $\lambda = 1, 6$.
- Since both eigenvalues are positive, Q is positive definite.

• **Example:** Determine the definiteness of the quadratic form $Q(x, y) = 3x^2 - 6xy - 5y^2$ on \mathbb{R}^2 .

- The associated matrix is $\begin{bmatrix} 3 & -3 \\ -3 & -5 \end{bmatrix}$, whose characteristic polynomial is $p(t) = (t - 3)(t + 5) - (-3)(-3) = t^2 - 2t - 24 = (t + 6)(t - 4)$, so its eigenvalues are $\lambda = -6, 4$.
- Since one eigenvalue is positive and the other is negative, Q is indefinite.

• **Example:** Determine the definiteness of the quadratic form $Q(x, y, z) = 3x^2 - 2xy + 4xz + 3y^2 - 4yz + 2z^2$ on \mathbb{R}^3 .

- The associated matrix is $\begin{bmatrix} 3 & -1 & 2 \\ -1 & 3 & -2 \\ 2 & -2 & 2 \end{bmatrix}$, whose characteristic polynomial is $p(t) = t^3 - 8t^2 + 12t = t(t - 2)(t - 6)$, so its eigenvalues are $\lambda = 0, 2, 6$.
- Since one eigenvalue is zero and the others are positive, Q is positive semidefinite.

• We can also view definiteness as a property of symmetric matrices themselves by considering the definiteness of the associated quadratic form. In this lens, we can give another way to identify definiteness using determinants:

• **Theorem (Sylvester's Criterion):** Suppose A is an $n \times n$ real matrix. For each $1 \leq k \leq n$, define the k th principal minor A_k to be the upper-left $k \times k$ corner of A . Then A is positive definite if and only if $\det(A_k) > 0$ for all k , and A is positive semidefinite if and only if $\det(A_k) \geq 0$ for all k .

- Since A is positive (semi)definite if and only if $-A$ is negative (semi)definite, one can also use Sylvester's criterion to identify negative definite and negative semidefinite matrices.
- We will not prove Sylvester's criterion, although it is not hard to see that the given condition is necessary, since if A is positive definite we must have $\mathbf{x}^T A \mathbf{x} > 0$ for all vectors $\mathbf{x} = [x_1 \ x_2 \ \cdots \ x_k \ 0 \ \cdots \ 0]$: this means the matrix A_k must also be positive definite and therefore must have positive determinant. (A similar observation holds when A is positive semidefinite.)

• **Example:** Determine the definiteness of the matrix $A = \begin{bmatrix} 2 & -1 & 2 \\ -1 & 4 & 2 \\ 2 & 2 & 5 \end{bmatrix}$.

- The principal minors are $[2]$, $\begin{bmatrix} 2 & -1 \\ -1 & 4 \end{bmatrix}$, and $\begin{bmatrix} 2 & -1 & 2 \\ -1 & 4 & 2 \\ 2 & 2 & 5 \end{bmatrix}$, whose determinants respectively are 2, 7, and 3.
- Since all of the principal minors have positive determinants, the given matrix is positive definite.

- Remark: To four decimal places, the eigenvalues are 6.8004, 4.0917, and 0.1078, so we see A is indeed positive definite.
- One extremely useful practical application of classifying the definiteness of a quadratic form is that it allows us to classify critical points of functions of several variables in many cases.
 - This type of calculation arises quite often in numerical optimization problems that employ step methods to search for maximum or minimum values of functions.
 - Such methods will usually terminate by finding a critical point of the underlying function, at which point it becomes necessary to classify the critical point's type in order to determine whether an actual minimum or maximum has been found.
- This particular result from multivariable calculus is often called the “second derivatives test” :
- Theorem (Second Derivatives Test in \mathbb{R}^n): Suppose f is a function of n variables x_1, \dots, x_n that is twice-differentiable and P is a critical point of f , so that $f_{x_i}(P) = 0$ for each i . Let H be the Hessian matrix, whose (i, j) -entry is the second-order partial derivative $f_{x_i x_j}(P)$. If H is positive definite then f has a local minimum at P , if H is negative definite then f has a local maximum at P , and if H is indefinite then f has a saddle point at P . If H is positive or negative semidefinite, then the test is inconclusive.
 - The idea of the proof (which we only outline) is to observe that because all of the partial derivatives of f vanish at P , the value $f(x_1, \dots, x_n) - f(P)$ is closely approximated by the quadratic Taylor expansion of f at P .
 - This quadratic Taylor expansion is precisely the quadratic form whose associated matrix is H , and so the behavior of f near P will be determined by the definiteness of H .
 - When H is positive-definite, this says $f(x_1, \dots, x_n) - f(P)$ is positive near P , meaning that P is a local minimum. Likewise, when H is negative-definite, then $f(x_1, \dots, x_n) - f(P)$ is negative near P , meaning that P is a local maximum, and when H is indefinite, P is a saddle point since f takes values above and below $f(P)$.
 - The only minor issue occurs when one of the eigenvalues is zero: in that case, the quadratic Taylor approximation does not determine the behavior of f as one approaches P along the direction of the corresponding eigenvector (it would be necessary to look at higher derivatives of f).
- Example: Classify the critical point at $(0, 0)$ for the function $f(x, y) = x^3 + 2x^2 + xy + 4y^2$.
 - We compute the Hessian matrix: we have $f_{xx} = 6x + 4$, $f_{xy} = f_{yx} = 1$, and $f_{yy} = 8$, so evaluating these at $(0, 0)$ yields $H = \begin{bmatrix} 4 & 1 \\ 1 & 8 \end{bmatrix}$.
 - The characteristic polynomial of H is $p(t) = \det(tI_2 - H) = t^2 - 12t + 31$, whose roots are $\lambda = 6 \pm \sqrt{5}$.
 - Since the eigenvalues are both positive, the critical point is a local minimum.
 - Alternatively, we could have used Sylvester's criterion, since the principal minors have determinants 4 and 31. Since these are both positive, we see that the Hessian matrix is positive-definite, so the critical point is a local minimum.

5.3.4 Singular Values and Singular Value Decomposition

- Diagonalization is a very useful tool, but it suffers from two main drawbacks: first, not all linear transformations $T : V \rightarrow V$ are diagonalizable, and second, we cannot diagonalize general linear transformations $T : V \rightarrow W$ when V and W are different.
 - There are various tools, such as the Jordan canonical form, that can give a “near diagonalization” for non-diagonalizable linear transformations $T : V \rightarrow V$.
- We will instead discuss another decomposition that is in some sense a hybrid between QR factorization and diagonalization, called singular value decomposition.

- The main idea is as follows: if we have a linear transformation $T : V \rightarrow W$ where V and W are finite-dimensional inner product spaces, then we may construct orthonormal bases β of V and γ of W such that the associated matrix $A = [T]_{\beta}^{\gamma}$ has its only nonzero entries in the “diagonal” positions $a_{i,i}$.
 - This procedure combines the ideas of diagonalization, in that we obtain a representation of T by an essentially diagonal matrix (up to not being square), and the QR factorization, in that we perform orthonormal changes of basis to simplify the form of the transformation.
 - The main idea is to note that if A is any $m \times n$ matrix, then $A^T A$ is a symmetric $m \times m$ matrix (since the transpose $(A^T A)^T = A^T (A^T)^T = A^T A$ again).
 - Moreover, the quadratic form $Q_{A^T A}(\mathbf{v}) = \mathbf{v}^T (A^T A) \mathbf{v}$ is positive semidefinite, since the quantity $\mathbf{v}^T (A^T A) \mathbf{v} = (A\mathbf{v})^T (A\mathbf{v}) = \|A\mathbf{v}\|^2$ is necessarily nonnegative, so from our characterization of the definiteness of quadratic forms, we see that all of the eigenvalues of $A^T A$ are nonnegative.
 - Therefore, since $A^T A$ is a symmetric $n \times n$ matrix with nonnegative eigenvalues, it can be orthogonally diagonalized, and the diagonal entries of its diagonalization D are $\sigma_1^2, \sigma_2^2, \dots, \sigma_n^2$ for some nonnegative real numbers $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$.
 - We can then use the orthonormal basis of eigenvectors of $A^T A$ to write down the desired decomposition of A .
- To make this precise, we introduce some terminology:
 - **Definition:** If A is any $m \times n$ matrix, the singular values of A are the nonnegative real numbers $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$ such that $\sigma_1^2, \sigma_2^2, \dots, \sigma_n^2$ are the eigenvalues of $A^T A$.
 - By the above discussion, since $A^T A$ is a real symmetric $n \times n$ matrix that is positive semidefinite, its eigenvalues are nonnegative real numbers. The (nonnegative) square roots of these numbers are the singular values of A .
 - We have previously shown that the rank of $A^T A$ is the same as the rank of A . Therefore, if A has rank r , the singular values $\sigma_1, \dots, \sigma_r$ will be positive and the remaining ones $\sigma_{r+1}, \dots, \sigma_n$ will be zero.
 - **Example:** Find the singular values of the matrix $A = \begin{bmatrix} 2 & 2 \\ 2 & 2 \\ -1 & 1 \\ 1 & -1 \end{bmatrix}$.
 - We have $A^T A = \begin{bmatrix} 2 & 2 & -1 & 1 \\ 2 & 2 & 1 & -1 \end{bmatrix} \begin{bmatrix} 2 & 2 \\ 2 & 2 \\ -1 & 1 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} 10 & 6 \\ 6 & 10 \end{bmatrix}$ with characteristic polynomial $p(t) = \det(tI_2 - A) = (t - 10)(t - 10) - (6)(6) = t^2 - 20t + 64 = (t - 4)(t - 16)$.
 - Since the eigenvalues of $A^T A$ are $\lambda = 4, 16$, we see that the singular values of A are $\boxed{4, 2}$.
 - **Example:** Find the singular values of the matrix $A = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 0 & 1 \end{bmatrix}$.
 - We have $A^T A = \begin{bmatrix} 1 & 1 & 0 \\ 1 & -1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix}$ with characteristic polynomial $p(t) = \det(tI_2 - A) = (t - 2)(t - 3)$.
 - Since the eigenvalues of $A^T A$ are $\lambda = 2, 3$, we see that the singular values of A are $\boxed{\sqrt{3}, \sqrt{2}}$.
 - **Example:** Find the singular values of the matrix $A = \begin{bmatrix} 1 & 2 & 1 \\ 2 & 1 & 1 \end{bmatrix}$.

- We have $A^T A = \begin{bmatrix} 1 & 2 \\ 2 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 5 & 4 & 3 \\ 4 & 5 & 3 \\ 3 & 3 & 2 \end{bmatrix}$ with characteristic polynomial $p(t) = \det(tI_3 - A) = (t-5)[t^2 - 7t + 1] - (-4)[-4t - 1] + (-3)[3t - 3] = t^3 - 12t^2 + 11t = t(t-1)(t-11)$.

- Since the eigenvalues of $A^T A$ are $\lambda = 0, 1, 11$, we see that the singular values of A are $\boxed{\sqrt{11}, 1, 0}$.

- Our general result is that we can use the singular values of a matrix to write down a matrix associated to T with respect to orthonormal bases that is particularly nice:

- **Theorem** (Singular Value Basis): Suppose $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear transformation of rank r with singular values $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$. Then there exist orthonormal bases $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of \mathbb{R}^n and $\{\mathbf{w}_1, \dots, \mathbf{w}_m\}$ of \mathbb{R}^m such that $T(\mathbf{v}_1) = \sigma_1 \mathbf{w}_1$, $T(\mathbf{v}_2) = \sigma_2 \mathbf{w}_2$, ..., $T(\mathbf{v}_r) = \sigma_r \mathbf{w}_r$, and $T(\mathbf{v}_{r+1}) = T(\mathbf{v}_{r+2}) = \dots = T(\mathbf{v}_n) = \mathbf{0}$.

- The main idea is to use the orthonormal basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of eigenvectors of $A^T A$, where A is the matrix associated to T .

- **Proof:** Suppose the matrix associated to T with respect to the standard bases of \mathbb{R}^n and \mathbb{R}^m is A .

- As noted in the discussion earlier, $A^T A$ is a positive-semidefinite symmetric $n \times n$ matrix of rank r . By the real spectral theorem, \mathbb{R}^n has an orthonormal basis of eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_r, \mathbf{v}_{r+1}, \dots, \mathbf{v}_n$ with associated eigenvalues $\sigma_1^2, \sigma_2^2, \dots, \sigma_r^2, 0, \dots, 0$.

- Now note that $\langle T(\mathbf{v}_i), T(\mathbf{v}_j) \rangle = \langle A\mathbf{v}_i, A\mathbf{v}_j \rangle = (A\mathbf{v}_i)^T (A\mathbf{v}_j) = \mathbf{v}_i^T (A^T A) \mathbf{v}_j = \sigma_i^2 \langle \mathbf{v}_i, \mathbf{v}_j \rangle$.

- Therefore, for $i \neq j$, since $\mathbf{v}_i, \mathbf{v}_j$ are orthonormal, the last term is zero, so $T(\mathbf{v}_i)$ and $T(\mathbf{v}_j)$ are orthonormal. Furthermore, for $i = j$, we see that $\|T(\mathbf{v}_i)\| = \sigma_i$.

- This means the vectors $\mathbf{w}_1 = T(\mathbf{v}_1)/\sigma_1$, $\mathbf{w}_2 = T(\mathbf{v}_2)/\sigma_2$, ..., and $\mathbf{w}_r = T(\mathbf{v}_r)/\sigma_r$ are orthonormal. By extending this set to an orthonormal basis $\{\mathbf{w}_1, \dots, \mathbf{w}_m\}$ of \mathbb{R}^m , we obtain the claimed result.

- We can recast the theorem above in terms of matrices, as follows:

- **Theorem** (Singular Value Decomposition): Suppose A is an $m \times n$ real matrix of rank r . If $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r > 0$ are the nonzero singular values of A , then A can be written as a matrix product $A = U\Sigma V^T$ where U is an orthogonal $n \times n$ matrix, V is an orthogonal $m \times m$ matrix, and Σ is the $n \times m$ matrix whose first r diagonal entries are $\sigma_1, \dots, \sigma_r$ and whose remaining entries are 0.

- **Proof:** Let $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be the linear transformation with $T(\mathbf{v}) = A\mathbf{v}$.

- By the theorem above, we have orthonormal bases $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of \mathbb{R}^n and $\gamma = \{\mathbf{w}_1, \dots, \mathbf{w}_m\}$ of \mathbb{R}^m such that $T(\mathbf{v}_1) = \sigma_1 \mathbf{w}_1$, $T(\mathbf{v}_2) = \sigma_2 \mathbf{w}_2$, ..., $T(\mathbf{v}_r) = \sigma_r \mathbf{w}_r$, and $T(\mathbf{v}_{r+1}) = T(\mathbf{v}_{r+2}) = \dots = T(\mathbf{v}_n) = \mathbf{0}$.

- This means the associated matrix $[T]_\beta^\gamma$ is the $n \times m$ matrix Σ whose first r diagonal entries are $\sigma_1, \dots, \sigma_r$ and whose remaining entries are 0.

- Now let α be the standard basis of \mathbb{R}^n and δ be the standard basis of \mathbb{R}^m and note that $[T]_\alpha^\delta = A$. Furthermore, since β is orthonormal the change-of-basis matrix $[I]_\beta^\alpha$ is orthogonal hence so is its inverse (also equal to its transpose) $V = [I]_\alpha^\beta$, and since γ is orthonormal the change-of-basis matrix $U = [I]_\gamma^\delta$ is also orthogonal.

- Then $A = [T]_\alpha^\delta = [I]_\gamma^\delta [T]_\beta^\gamma [I]_\alpha^\beta = U\Sigma V^T$, as claimed.

- To calculate the singular value decomposition $A = U\Sigma V^T$ of an $m \times n$ matrix A , follow these steps:

- **Step 1:** Find an orthonormal basis $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of \mathbb{R}^n consisting of unit eigenvectors for the symmetric matrix $A^T A$ with corresponding eigenvalues $\sigma_1^2, \dots, \sigma_n^2$.

- **Step 2:** Calculate the unit vectors $\mathbf{w}_i = A\mathbf{v}_i/\sigma_i$, and then use Gram-Schmidt (if necessary) to extend this set to an orthonormal basis of \mathbb{R}^m .

- **Step 3:** Write down the matrices U whose columns are the vectors \mathbf{w}_i , Σ whose diagonal elements are the singular values σ_i , and V whose columns are the vectors \mathbf{v}_i . Then $A = U\Sigma V^T$.

- **Example:** Find the singular values, and a singular value decomposition, of $A = \begin{bmatrix} 0 & 6 \\ 6 & 5 \end{bmatrix}$.

- We have $A^T A = \begin{bmatrix} 36 & 30 \\ 30 & 61 \end{bmatrix}$ with characteristic polynomial $p(t) = (t - 36)(t - 61) - (30)(30) = t^2 - 97t + 1296 = (t - 16)(t - 81)$, so the singular values of A are $\sigma_1 = \sqrt{81} = 9$ and $\sigma_2 = \sqrt{16} = 4$.
- We can then find a basis for the 81-eigenspace of $A^T A$ as $\left\{ \begin{bmatrix} 2 \\ 3 \end{bmatrix} \right\}$ and a basis for the 16-eigenspace of $A^T A$ as $\left\{ \begin{bmatrix} -3 \\ 2 \end{bmatrix} \right\}$, so after normalizing we can take $\mathbf{v}_1 = \frac{1}{\sqrt{13}} \begin{bmatrix} 2 \\ 3 \end{bmatrix}$ and $\mathbf{v}_2 = \frac{1}{\sqrt{13}} \begin{bmatrix} -3 \\ 2 \end{bmatrix}$.
- We also have $\mathbf{w}_1 = \frac{1}{9} A \mathbf{v}_1 = \frac{1}{\sqrt{13}} \begin{bmatrix} 2 \\ 3 \end{bmatrix}$ and $\mathbf{w}_2 = \frac{1}{4} A \mathbf{v}_2 = \frac{1}{\sqrt{13}} \begin{bmatrix} 3 \\ -2 \end{bmatrix}$; as expected we see that $\{\mathbf{w}_1, \mathbf{w}_2\}$ is an orthonormal set (and in fact an orthonormal basis) of \mathbb{R}^2 .
- Putting all of this together, we get the matrices $U = \frac{1}{\sqrt{13}} \begin{bmatrix} 2 & 3 \\ 3 & -2 \end{bmatrix}$, $\Sigma = \begin{bmatrix} 9 & 0 \\ 0 & 4 \end{bmatrix}$, and $V^T = \frac{1}{\sqrt{13}} \begin{bmatrix} 2 & -3 \\ 3 & 2 \end{bmatrix}$.

- **Example:** Find a singular value decomposition of $A = \begin{bmatrix} 2 & 2 \\ 2 & 2 \\ -1 & 1 \\ 1 & -1 \end{bmatrix}$.


- We previously found that the eigenvalues of $A^T A = \begin{bmatrix} 10 & 6 \\ 6 & 10 \end{bmatrix}$ are $\lambda = 16, 4$ and so the singular values of A are $\sigma_1 = 4$ and $\sigma_2 = 2$.
- We can then calculate a basis for the 16-eigenspace of $A^T A$ as $\left\{ \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}$ and a basis for the 4-eigenspace of $A^T A$ as $\left\{ \begin{bmatrix} -1 \\ 1 \end{bmatrix} \right\}$, so after normalizing we can take $\mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}$ and $\mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 1 \end{bmatrix}$.
- We also have $\mathbf{w}_1 = \frac{1}{4} A \mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}$ and $\mathbf{w}_2 = \frac{1}{2} A \mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 0 \\ 1 \\ -1 \end{bmatrix}$; as expected we see that $\{\mathbf{w}_1, \mathbf{w}_2\}$ is an orthonormal set.
- By using Gram-Schmidt, we can extend $\{\mathbf{w}_1, \mathbf{w}_2\}$ to an orthonormal basis $\{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \mathbf{w}_4\}$ of \mathbb{R}^4 with $\mathbf{w}_3 = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \\ 0 \\ 0 \end{bmatrix}$ and $\mathbf{w}_4 = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}$.

- Putting all of this together, we get the matrices $U = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & -1 & 0 & 1 \end{bmatrix}$, $\Sigma = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$, and

$$V^T = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}.$$

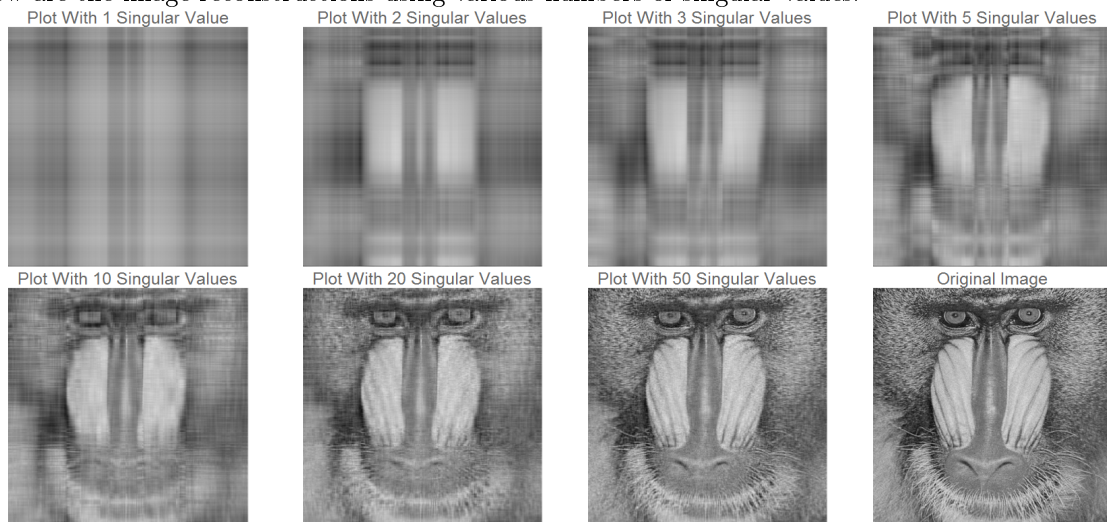
- **Remark:** Note that the singular value decomposition is not unique here, since for example we could choose other vectors $\mathbf{w}_3, \mathbf{w}_4$ to complete the orthonormal basis of \mathbb{R}^4 .
- The singular value basis and associated decomposition have a convenient geometric interpretation in terms of the action of the transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ on the “unit sphere” $\|\mathbf{v}\| = 1$ in \mathbb{R}^n .
 - To illustrate, consider the linear transformation $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ with associated standard matrix $A = \begin{bmatrix} 0 & 6 \\ 6 & 5 \end{bmatrix}$ from the first example above.

- The image of the unit circle $\|\mathbf{v}\| = 1$ (i.e., $x^2 + y^2 = 1$) under T is an ellipse, shown below:
Image of Unit Circle Under T




- We can see quite clearly from the second picture that the vectors $A\mathbf{v}_1$ and $A\mathbf{v}_2$ give the principal axes of the ellipse.
- This observation can be verified algebraically from the facts that $\beta = \{\mathbf{v}_1, \mathbf{v}_2\}$ and $\gamma = \{\mathbf{w}_1, \mathbf{w}_2\}$ are orthonormal bases of \mathbb{R}^2 and the fact that $[T]_\beta^\gamma$ is the matrix $\begin{bmatrix} 9 & 0 \\ 0 & 4 \end{bmatrix}$: then the image $9a\mathbf{w}_1 + 4b\mathbf{w}_2$ of any linear combination $a\mathbf{v}_1 + b\mathbf{v}_2$ on the unit circle (i.e., with $a^2 + b^2 = 1$) has norm $81a^2 + 16b^2$, and the norm is clearly maximized when $b = 0$ and minimized when $a = 0$.
- This means that the major axis of the ellipse is parallel to \mathbf{w}_1 and has length σ_1 , while the minor axis of the ellipse is parallel to \mathbf{w}_2 and has length σ_2 .
- It is not hard to see that analogous results hold in higher dimensions, for the same reasons: in general, the image of the unit sphere $\|\mathbf{v}\| = 1$ under a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^m$ of rank r will be an r -dimensional ellipsoid whose principal axes are the vectors $\sigma_1\mathbf{w}_1, \sigma_2\mathbf{w}_2, \dots, \sigma_r\mathbf{w}_r$.
- This geometric interpretation of singular value decomposition has many practical applications, such as performing principal component analysis and doing data compression.
 - The main idea is that for an $m \times n$ matrix A with singular values $\sigma_1, \dots, \sigma_r$ and corresponding orthonormal bases $\{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\{\mathbf{w}_1, \dots, \mathbf{w}_m\}$, if we multiply out the matrix product $A = U\Sigma V^T$, we can rephrase the singular value decomposition as giving a sum $A = \sigma_1\mathbf{v}_1\mathbf{w}_1^T + \sigma_2\mathbf{v}_2\mathbf{w}_2^T + \dots + \sigma_r\mathbf{v}_r\mathbf{w}_r^T$ of a total of r $m \times n$ matrices $\mathbf{v}_i\mathbf{w}_i^T$ each of which has rank 1.
 - Therefore, if we want to approximate A by a matrix of rank less than r , the best approximation will be given by deleting the terms of the sum above that have the smallest norm, which are the terms with smallest σ_i .
 - In other words, the best approximation to A by a matrix of rank d is obtained by taking the initial terms of the singular value sum above: $\sigma_1\mathbf{v}_1\mathbf{w}_1^T + \sigma_2\mathbf{v}_2\mathbf{w}_2^T + \dots + \sigma_d\mathbf{v}_d\mathbf{w}_d^T$.
 - In the situation where we have a set of data that is high-dimensional (i.e., lies inside \mathbb{R}^n where n is large), this gives an explicit procedure for projecting onto a smaller-dimensional subspace that loses as little information as possible.
- We can illustrate these ideas by calculating the singular value decomposition of a matrix representing the 2-dimensional grid of color intensity from a black-and-white photograph (taken from the standard set of sample data included with Mathematica).
 - The photograph used here is 512 pixels by 512 pixels, corresponding to a 512×512 matrix A .
 - We can then give compressed versions of the image file by taking the initial terms of the singular value sum.

- Below are the image reconstructions using various numbers of singular values:



- The total amount of data required to store the full image is the equivalent to 512^2 data points (one per pixel). To store the decomposition with k singular values, on the other hand, requires only storing about $2k \cdot 512$ data points (each singular value matrix $\mathbf{v}_i \mathbf{w}_i^T$ requires just the values of the vectors \mathbf{v}_i and \mathbf{w}_i).
- So, for example, to store and reconstruct the compressed image using 20 singular values only requires about $40/512 \approx 8\%$ of the total amount of uncompressed data in the original image.
- The reason this sort of procedure works is because most of the information in the image is carried by the first few singular values, which are much larger than the later ones. For this image, the first ten singular values are 66679, 10490, 5904, 4144, 3501, 2853, 2664, 2420, 2384, and 2188, with most of the remaining values being smaller:



- Therefore, taking just the first few singular values will capture the vast majority of information contained in the data set.

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

Copyright notice: This material is copyright Evan Dummit, 2012-2021. You may not reproduce or distribute this material without my express permission.