Apprentice Linear Algebra, 4th lecture, 07/08/05

REU 2005

Instructor: László Babai
Scribe: Wei Ren

4 Eigenvalues and Eigenvectors

4.1 System of Linear Equations

Recall that a system of k linear equations in ℓ unknowns can be represented concisely as

$$A\mathbf{x} = \mathbf{b}, \tag{1}$$

where A is a $k \times \ell$ matrix. We say the system is *homogeneous* if $\mathbf{b} = \mathbf{0}$, and we have proven that for a homogeneous system of linear equations, the set of solutions U is a subspace of $\mathbb{R}^\ell$, and in fact the basic fact on systems of linear equations tells us $\dim(U) = \ell - \operatorname{rk}(A)$. Consequently, we have the following theorem:

Theorem 4.1. *A system of homogeneous linear equations has a non-trivial solution iff* $\operatorname{rk}(A) < \ell$*; i.e., the matrix A does not have full column-rank (its rank is less than the number fo columns).*

Proof: By the basic fact, if $\operatorname{rk}(A) < \ell$, then $\dim(U) > 0$. This means U has a non-zero vector, i.e., a non-trivial solution. □

Since the row rank is equal to the column rank for any matrix, it immediately follows that

Corollary 4.2. *If $k < \ell$ (there are fewer equations than unknowns), then the system has a nontrivial solution.*

Now we consider non-homogeneous systems, where $\mathbf{b} \neq \mathbf{0}$. We want to know when such a system has a solution. Recall that if $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_\ell \end{bmatrix}$, and the matrix A is written as $A = \begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_\ell \end{bmatrix}$, where $\mathbf{a}_i$ is the i^{th} column of the matrix, then $A\mathbf{x} = \sum_{i=1}^{i=\ell} x_i \mathbf{a}_i$ is simply a linear combination of the columns. Therefore if $A\mathbf{x} = \mathbf{b}$ holds, then $\mathbf{b} \in \operatorname{Span}\{\mathbf{a}_1, \dots, \mathbf{a}_\ell\}$. Recall that the span of the a_i is called the *column space* of the matrix A. It follows that $\operatorname{rk}(A) = \operatorname{rk}([A|\mathbf{b}])$, where the *augmented* matrix $[A|\mathbf{b}]$ is defined as $\begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_\ell & \mathbf{b} \end{bmatrix}$. Conversely, if we know $\operatorname{rk}(A) = \operatorname{rk}([A|\mathbf{b}])$, then $\mathbf{b}$ is in the column space of A. We have thus shown the following:

Theorem 4.3. *The system of linear equations $A\mathbf{x} = \mathbf{b}$ has a solution iff* $\mathrm{rk}(A) = \mathrm{rk}([A|\mathbf{b}])$.

Using the above theorem, we can determine whether or not a given system has a solution. To completely describe the set of solutions for $A\mathbf{x} = \mathbf{b}$, we consider the related homogeneous system, $A\mathbf{x} = \mathbf{0}$, and let U be the its set of the solutions. We have the following description:

Theorem 4.4. *If $A\mathbf{x} = \mathbf{b}$ has a solution $\mathbf{w}$, then the solution set for the system is*

$$U + \mathbf{w} := \{\mathbf{x} + \mathbf{w} \,|\, \mathbf{x} \in U\}.$$

Proof: We have $A\mathbf{x} = \mathbf{b} \iff A\mathbf{x} = A\mathbf{w} \iff A(\mathbf{x} - \mathbf{w}) = \mathbf{0} \iff \mathbf{x} - \mathbf{w} \in U$. $\square$

Note that the theorem states the solution set for any system of linear equations is either empty or it is a translate of a subspace (the solution space for the corresponding homogeneous system).

4.2 Eigenvectors

Let $M_n(\mathbb{F})$ denote the set of all $n \times n$ matrices. First note that it is a ring, because we have the 3 usual arithmetic operations (addition, subtraction and multiplication). In addition, $M_n(\mathbb{F})$ is a vector space over $\mathbb{F}$, of dimension n^2. To see this, define E_{ij} to be the matrix whose only nonzero entry is the ij^{th} entry, and its value is 1. For any matrix $A = (a_{ij})$, we have $A = \sum a_{ij}E_{ij}$. The representation is unique and this shows the set $\{E_{ij}\}_{\substack{i=1\ldots n\\ j=1\ldots n}}$ form a basis for $M_n(\mathbb{F})$.

Definition 4.5. A ring $\mathcal{A}$ is an *algebra* over $\mathbb{F}$ if $\mathcal{A}$ is also a vector space over $\mathbb{F}$ and it satisfies the identity

$$(\forall \lambda \in \mathbb{F})(\forall a, b \in \mathcal{A})((\lambda a)b = \lambda(ab)).$$

Examples 4.6. The field of complex numbers $\mathbb{C}$ is a field extension of $\mathbb{Q}$, so it is a vector space over $\mathbb{Q}$. It also has a ring structure. Therefore $\mathbb{C}$ is an (infinitely dimensional) algebra over $\mathbb{Q}$. It is also a 2−dimensional algebra over $\mathbb{R}$.

The next definition is one of the most important in all of linear algebra.

Definition 4.7. Let $A \in M_n(\mathbb{F})$. A vector $\mathbf{x} \in \mathbb{F}^n$ is an *eigenvector* of A if

(i) $\mathbf{x} \neq \mathbf{0}$,

(ii) $A\mathbf{x} = \lambda\mathbf{x}$ for some $\lambda \in \mathbb{F}$.

Eigenvectors can also be defined for linear transformations.

Definition 4.8. Let $f : V \to V$ be a linear transformation. A vector $\mathbf{x} \in V$ is an *eigenvector* of f if

(i) $\mathbf{x} \neq \mathbf{0}$,

(ii) $f(\mathbf{x}) = \lambda \mathbf{x}$ for some $\lambda \in \mathbb{F}$.

Naturally we also have the following definition:

Definition 4.9. We say $\lambda \in \mathbb{F}$ is an *eigenvalue* of the matrix A (resp. of the linear map f) if there is an eigenvector $\mathbf{x} \in \mathbb{F}^n$ (resp. $\mathbf{x} \in V$) s.t. $A\mathbf{x} = \lambda \mathbf{x}$ (resp. $f(\mathbf{x}) = \lambda \mathbf{x}$).

Examples 4.10. Consider the geometric transformation that reflects vectors in 3−space in a given plane. Any nonzero vector that is orthogonal to the plane is an eigenvector, with eigenvalue -1. Any nonzero vector in the plane is also an eigenvector, with eigenvalue 1.

Examples 4.11. We now define a geometric transformation in the plane that has eigenvalues 2 and 3. Take any two nonzero vectors that are not parallel to each other, say $\mathbf{v}_1$ and $\mathbf{v}_2$. We know the set $\{\mathbf{v}_1, \mathbf{v}_2\}$ is a basis for the plane. We can define a linear map by specifying its actions on a basis. Define L by $L(\mathbf{v}_1) = 2\mathbf{v}_1, L(\mathbf{v}_2) = 3\mathbf{v}_2$, i.e., the map L stretches vectors by a factor of 2 in the direction of $\mathbf{v}_1$, and it stretches in the direction of $\mathbf{v}_2$ by a factor of 3. Then the vector $\mathbf{v}_1$ is an eigenvector with eigenvalue 2, and $\mathbf{v}_2$ is an eigenvector with eigenvalue 3.

The following exercise puts a bound on the maximal number of eigenvalues we can have.

Exercise 4.12. Prove that if $\mathbf{e}_1, \dots, \mathbf{e}_k$ are eigenvectors to distinct eigenvalues, then they are linearly independent.

An immediate corollary to this exercise is:

Corollary 4.13. *The number of eigenvalues is $\leq n$, or* $\dim(V)$.

Definition 4.14. An *eigenbasis* is a basis consisting of eigenvectors.

A linear transformation may or may not have an eigenbasis, as the examples below illustrate.

Examples 4.15. Consider rotation of a plane by an angle of α, where α is not a multiple of π. Then this transformation has no eigenvectors at all.

Examples 4.16. Consider the shearing transformation of the plane, which in coordinates can be defined by $(x, y) \mapsto (x + y, y)$. We see geometrically that all eigenvectors lie on the line defined by $(1, 0)$. Therefore we do not have an eigenbasis.

Examples 4.17. The transformation that reflects a vector in $\mathbb{R}^3$ across a given plane does have an eigenbasis. As we have observed before, all vectors in the plane, as well as vectors orthogonal to the plane, are eigenvectors. Therefore to form a basis, we simply pick one vector that is orthogonal, and then two non-parallel vectors from the plane.

The next theorem is one of the most widely applied mathematical results not only within mathematics but in science and engineering. We will prove it later in the course.

Theorem 4.18 (The Spectral Theorem). *If A is a symmetric real matrix (i.e., $A^T = A$), then A has an orthogonal eigenbasis.*

An orthogonal eigenbasis means an eigenbasis where the basis vectors are pairwise orthogonal (perpendicular).

Examples 4.19. One application of the Spectral Theorem spawned a branch of Statistics called "Factor Analysis." One can prove, using the Spectral Theorem, that for any n given random variables, there exists a set of pairwise uncorrelated random variables such that the original random variables and the new (uncorrelated) random variables have the same span, i. e., each of the original random variables can be expressed as a linear combination of the new (uncorrelated) variables (the "factors") and each new variable can be expressed as a linear combination of the original ones.

Examples 4.20. Another example, in mechanics, is the theorem that every solid body has 3 perpendicular axes of inertia (if rotated about one of these axes and then left alone, the body will continue to rotate about the axis). In other words, from the point of view of rigid motion, each solid body, no matter how irregularly shaped, can be replaced by an ellipsoid.

Next we study the equation

$$A\mathbf{x} = \lambda\mathbf{x}.$$

We want to know what values of λ will get us nontrivial solutions for $\mathbf{x}$, which will then be eigenvectors. By rearranging the terms, we get the following system:

$$(\lambda I - A)\mathbf{x} = \mathbf{0}.$$

Here I denotes the identity matrix. By theorem (4.1), we see that this system has a nontrivial solution iff the matrix $(\lambda I - A)$ does not have full rank.

We have the following definition:

Definition 4.21. A matrix $A \in M_n(\mathbb{F})$ is *singular* if it does not have full rank. Equivalently, it is singular if $\det(A) = 0$.

The discussion before the definition now shows

Theorem 4.22. *The value λ is an eigenvalue of A iff* $\det(\lambda I - A) = 0$.

Examples 4.23. We consider the 2×2 matrices. Let $A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$. Then we get

$$\begin{aligned} \det(\lambda I - A) &= \begin{vmatrix} \lambda - a & b \\ c & \lambda - d \end{vmatrix} \\ &= (\lambda - a)(\lambda - d) - bc \\ &= \lambda^2 - (a+d)\lambda + (ad - bc). \end{aligned}$$

Note that in the last equation, we have $(a+d) = \mathrm{tr}(A), (ad-bc) = \det(A)$. Solutions to this quadratic equation (in λ) are the eigenvalues of A.

Inspired by the preceding example, we have the following definition:

Definition 4.24. The *characteristic polynomial* of A is defined to be

$$f_A(x) = \det(xI - A).$$

Observe that for an $n \times n$ matrix A, its characteristic polynomial is of degree n (in x). Moreover, we can express the coefficients of this polynomial as follows:

$$\begin{aligned}
&a_n = 1;\\
&a_{n-1} = -Tr(A);\\
&a_{n-2} = \sum \det(2 \times 2 \text{ symmetric minors});\\
&a_{n-3} = -\sum \det(3 \times 3 \text{ symmetric minors});\\
&\vdots\\
&a_0 = f_A(0) = \det(-A) = (-1)^n \det(A).
\end{aligned}$$

In summary, we have

$$a_{n-i} = (-1)^i \sum \det(i \times i \text{ symmetric minors}), \tag{2}$$

where the sum is taken over all such symmetric minors. (An $i \times i$ symmetric minor is an $i \times i$ submatrix of A that is symmetric along the diagonal.) Note that there are $\binom{n}{i}$ terms in this sum.

Exercise 4.25. Verify equation (2).

4.3 Matrices, Bases and Linear Transformations

In the previous sectin, we defined eigenvectors and eigenvalues separately for matrices and linear maps. In this section, we will show that every linear transformation can be represented as a matrix, by choosing a basis.

Suppose $\varphi : V \to W$ is a linear map, where V has dimension ℓ and W has dimension k. Recall that any vector $\mathbf{x} \in V$ can be uniquely represented by its coordinates in a given basis. In other words, if $\underline{\mathbf{e}} := \{\mathbf{e}_1, \ldots, \mathbf{e}_\ell\}$ is a basis for V, then we have $\mathbf{x} = \sum_{i=1}^{i=\ell} \alpha_i \mathbf{e}_i$, for $\alpha_i \in \mathbb{F}$. We write $[\mathbf{x}]_{\underline{\mathbf{e}}} = \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_\ell \end{pmatrix}$.

Observe that map $\mathbf{x} \mapsto [\mathbf{x}]_{\underline{\mathbf{e}}}$ is a one-to-one and onto map between V and $\mathbb{F}^n$. We thus have an isomorphism $V \approx \mathbb{F}^n$.

Now fix this basis $\underline{\mathbf{e}}$ for V, and similarly fix a basis $\underline{\mathbf{f}}$ for W: $\underline{\mathbf{f}} = \{\mathbf{f}_1, \dots, \mathbf{f}_k\}$. Define

$$[\varphi]_{\underline{\mathbf{e}},\underline{\mathbf{f}}} := \begin{bmatrix} [\varphi\mathbf{e}_1]_{\underline{\mathbf{f}}}, & \dots & [\varphi\mathbf{e}_\ell]_{\underline{\mathbf{f}}} \end{bmatrix}.$$

Here each $[\varphi\mathbf{e}_i]_{\underline{\mathbf{f}}}$ is a column vector of height k, and therefore $[\varphi]_{\underline{\mathbf{e}},\underline{\mathbf{f}}}$ is a $k \times \ell$ matrix.

The following example illustrates how to find this matrix we just defined. It is also one of the most delightful matrices.

Examples 4.26. We come back to our favorite example of linear transformation, rotation of a plane by an angle of α. Denote this transformation by φ_α. We think of the plane as $\mathbb{R}^2$. Let $\underline{\mathbf{e}}$ be the standard basis of $\mathbb{R}^2$. For simplicity, we will write $[\varphi_\alpha]_{\underline{\mathbf{e}}}$ instead of $[\varphi_\alpha]_{\underline{\mathbf{e}},\underline{\mathbf{e}}}$. We will use this convention for all the later examples, and we always use the same basis whenever we have a map from a vector into itself.

To find the matrix $[\varphi_\alpha]_{\underline{\mathbf{e}}}$, we must express $\varphi_\alpha(\mathbf{e}_1)$ and $\varphi_\alpha(\mathbf{e}_2)$ as linear combinations of $\mathbf{e}_1$ and $\mathbf{e}_2$. By a simple geometric observation, we get $\varphi_\alpha(\mathbf{e}_1) = (\cos\alpha)\mathbf{e}_1 + (\sin\alpha)\mathbf{e}_2$. For the second basis vector, note that we have the relation $\varphi_{\frac{\pi}{2}}(x, y) = (-y, x)$. Since the composition of rotations is clearly commutative, we have

$$\begin{aligned} \varphi_\alpha(\mathbf{e}_2) &= \varphi_\alpha \circ \varphi_{\frac{\pi}{2}}(\mathbf{e}_1) \\ &= \varphi_{\frac{\pi}{2}} \circ \varphi_\alpha(\mathbf{e}_1) \\ &= \varphi_{\frac{\pi}{2}}((\cos\alpha)\mathbf{e}_1 + (\sin\alpha)\mathbf{e}_2) \\ &= (-\sin\alpha)\mathbf{e}_1 + (\cos\alpha)\mathbf{e}_2. \end{aligned}$$

Therefore now by definition, we get

$$[\varphi_\alpha]_{\underline{\mathbf{e}}} = \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix}.$$

The determinant of this matrix is $(\cos\alpha)^2 + (\sin\alpha)^2 = 1$.

We can try to find the eigenvalues of this matrix. Geometrically, it is clear that unless α is a multiple of π, we do not have eigenvectors. We compute the characteristic polynomial of this matrix:

$$\begin{aligned} f_{\varphi_\alpha}(x) &= \det(xI - [\varphi_\alpha]_{\underline{\mathbf{e}}}) \\ &= \begin{vmatrix} x - \cos\alpha & \sin\alpha \\ -\sin\alpha & x - \cos\alpha \end{vmatrix} \\ &= x^2 - (2\cos\alpha)x + 1. \end{aligned}$$

Over the complex numbers, we have two solutions to this equation: $\lambda = \cos\alpha \pm i\sin\alpha$. They are the eigenvalues of the matrix $[\varphi_\alpha]_{\underline{\mathbf{e}}}$ over $\mathbb{C}$, and we see that we indeed have no real eigenvalues unless $\sin\alpha = 0$, i.e. α is a multiple of π.

Now we carry out this computation all over again. This time we choose a different basis for $\mathbb{R}^2$. Let $\underline{\mathbf{e}}' = \{\mathbf{e}_1, \mathbf{e}_3\}$, where $\mathbf{e}_1$ is the same as before, and $\mathbf{e}_3$ is $\varphi_\alpha \mathbf{e}_1$. Assuming α is not a multiple of π, we have formed a basis. We must compute $\mathbf{w} = \varphi_\alpha \mathbf{e}_3$. Using geometry, we see that $\mathbf{e}_1 + \mathbf{w} = \gamma \mathbf{e}_3$, where $\gamma = 2\cos\alpha$. This shows $\mathbf{w} = (2\cos\alpha)\mathbf{e}_3 - \mathbf{e}_1$, and we get

$$[\varphi_\alpha]_{\underline{\mathbf{e}}'} = \begin{bmatrix} 0 & -1 \\ 1 & 2\cos\alpha \end{bmatrix}.$$

The characteristic polynomial of this matrix is

$$f_{[\varphi_\alpha]_{\underline{\mathbf{e}}'}} = x^2 - (2\cos\alpha)x + 1.$$

This is the exact same polynomial we got previously, using the standard basis.

This example suggests that the characteristic polynomial is invariant under basis change. This is indeed the case, and we will prove the fact later. We formulate this fact in the following theorem.

Theorem 4.27. *If $\varphi : V \to V$ is a linear transformation, then $f_{[\phi]_{\mathbf{B}}}(x)$ does not depend on the choice of the basis $\mathbf{B}$ for V.*

This theorem leads to the following definition:

Definition 4.28. The characteristic polynomial of the linear transformation $\varphi : V \to V$ is

$$f_\varphi(x) = f_{[\phi]_{\mathbf{B}}}(x),$$

where $\mathbf{B}$ is any basis of V.

The next exercise ensures us the representation of linear maps using matrices is consistent with the definition of the matrix multiplications.

Exercise 4.29. Suppose $\varphi : V \to W$ is a linear map, and we fix bases $\underline{\mathbf{e}}$ and $\underline{\mathbf{f}}$ for V and W, respectively. Prove that

$$[\varphi \mathbf{x}]_{\underline{\mathbf{f}}} = [\varphi]_{\underline{\mathbf{e}},\underline{\mathbf{f}}}[x]_{\underline{\mathbf{e}}}.$$

The multiplication on the right hand side is the usual matrix multiplication.

We next two exercises will show the definition of matrix multiplication is entirely natural.

Exercise 4.30. If A is a $k \times \ell$ matrix, and $A\mathbf{x} = \mathbf{0}$ for all $\mathbf{x} \in \mathbb{F}^\ell$, then A is the zero matrix.

An immediate corollary is the following:

Corollary 4.31. *If $A\mathbf{x} = B\mathbf{x}$ for all vectors $\mathbf{x} \in \mathbb{F}^\ell$, then $A = B$.*

Exercise 4.32. Suppose we have two maps $\varphi : V \to W$, $\psi : W \to T$. We choose bases $\underline{\mathbf{e}}, \underline{\mathbf{f}}, \underline{\mathbf{g}}$ for V, W, and T respectively. Prove that

$$[\varphi\psi]_{\underline{\mathbf{e}},\underline{\mathbf{g}}} = [\psi]_{\underline{\mathbf{f}},\underline{\mathbf{g}}}[\varphi]_{\underline{\mathbf{e}},\underline{\mathbf{f}}}.$$

(Hint: Use Corollary 4.31 and Exercise 4.29.)

We now see that the seemingly involved definition of matrix multiplication was constructed to satisfy the composition rules of linear maps. Since the composition of any maps is associative, we at once get the following result:

Corollary 4.33. *Matrix multiplication is associative.*

4.4 Change of Basis

Now that we can represent linear maps with matrices by choosing bases, we want to know the relations between matrix representations of the same linear map under different bases.

First we relate the coordinate vectors $[\mathbf{x}]_{\underline{\mathbf{e}}}$ and $[\mathbf{x}]_{\underline{\mathbf{e}}'}$, where $\underline{\mathbf{e}} = \{\mathbf{e}_1, \dots, \mathbf{e}_\ell\}$ and $\underline{\mathbf{e}}' = \{\mathbf{e}'_1, \dots, \mathbf{e}'_\ell\}$ are two bases for V. We define a "basis change transformation" $\sigma : V \to V$, given by

$$\sigma(\mathbf{e}_i) = \mathbf{e}'_i.$$

We know this map is unique and invertible ($\sigma^{-1} : \mathbf{e}' \mapsto \mathbf{e}$).

If $\mathbf{x} = \sum_{i=1}^{\ell} \alpha_i \mathbf{e}_i$ then we have

$$\sigma\mathbf{x} = \sigma(\sum_{i=1}^{\ell} \alpha_i \mathbf{e}_i) = \sum_{i=1}^{\ell} \alpha_i \sigma(\mathbf{e}_i) = \sum_{i=1}^{\ell} \alpha_i \mathbf{e}'_i.$$

This shows $[\sigma(\mathbf{x})]_{\underline{\mathbf{e}}'} = [\mathbf{x}]_{\underline{\mathbf{e}}}$, (as vectors of $\mathbb{F}^\ell$). By the relation $[\sigma(\mathbf{x})]_{\underline{\mathbf{e}}'} = [\sigma]_{\underline{\mathbf{e}}'}[\mathbf{x}]_{\underline{\mathbf{e}}'}$, where $[\sigma]_{\underline{\mathbf{e}}'}$ is the representation of σ as a matrix under the basis $\underline{\mathbf{e}}'$, we get the following formula for basis change:

$$[\mathbf{x}]_{\underline{\mathbf{e}}'} = [\sigma]_{\underline{\mathbf{e}}'}^{-1}[\mathbf{x}]_{\underline{\mathbf{e}}}. \tag{3}$$

Now suppose we have $\varphi : V \to W$, a linear map. Suppose $\underline{\mathbf{e}}$ and $\underline{\mathbf{e}}'$ are two bases for V, while $\underline{\mathbf{f}}$ and $\underline{\mathbf{f}}'$ are two bases for W. Further suppose we have the basis change transformations, $\sigma : V \to V$ and $\tau : W \to W$, where $\sigma(\mathbf{e}_i) = \mathbf{e}'_i$, $\tau(\mathbf{f}_j) = \mathbf{f}'_j$. Using the pairs of bases, $\{\underline{\mathbf{e}}, \underline{\mathbf{f}}\}$ and $\{\underline{\mathbf{e}}', \underline{\mathbf{f}}'\}$, we get two matrix representations for the map: $[\varphi]_{\underline{\mathbf{e}},\underline{\mathbf{f}}}$ and $[\varphi]_{\underline{\mathbf{e}}',\underline{\mathbf{f}}'}$. We will use the more suggestive notations $[\varphi]_{\text{old}}$ for the former matrix, and $[\varphi]_{\text{new}}$ for the latter. We will use the subscript "new" whenever we use the bases $\underline{\mathbf{e}}'$ or $\underline{\mathbf{f}}'$, and we will use "old" whenever we refer to the basis $\underline{\mathbf{e}}, \underline{\mathbf{f}}$. Using the formula we just derived, we get

$$[\varphi\mathbf{x}]_{\text{new}} = [\varphi]_{\text{new}}[\mathbf{x}]_{\text{new}} = [\varphi]_{\text{new}}[\sigma]_{\text{new}}^{-1}[\mathbf{x}]_{\text{old}};$$
$$[\varphi\mathbf{x}]_{\text{new}} = [\tau]_{\text{new}}^{-1}[\varphi\mathbf{x}]_{\text{old}} = [\tau]_{\text{new}}^{-1}[\varphi]_{\text{old}}[\mathbf{x}]_{\text{old}}.$$

This implies the following equality:

$$[\tau]_{\text{new}}^{-1}[\varphi]_{\text{old}}[\mathbf{x}]_{\text{old}} = [\varphi]_{\text{new}}[\sigma]_{\text{new}}^{-1}[\mathbf{x}]_{\text{old}}.$$

Since this holds for all vectors $\mathbf{x}$, by Corollary 4.31, we get

$$[\tau]_{\text{new}}^{-1}[\varphi]_{\text{old}} = [\varphi]_{\text{new}}[\sigma]_{\text{new}}^{-1}.$$

By multiplying both sides with $[\sigma]_{\text{new}}$, we get

$$[\varphi]_{\text{new}} = [\tau]_{\text{new}}^{-1}[\varphi]_{\text{old}}[\sigma]_{\text{new}}. \tag{4}$$

Since σ itself is a linear map (from V to itself), we can apply this formula and conclude

$$[\sigma]_{\text{new}} = [\sigma]_{\text{new}}^{-1}[\sigma]_{\text{old}}[\sigma]_{\text{new}}.$$

The basis change matrix is invertible as the corresponding linear map has an inverse, so we can cancel the $[\sigma]_{\text{new}}$ factor and derive $[\sigma]_{\text{new}}^{-1}[\sigma]_{\text{old}} = I$. qWe have shown

Corollary 4.34. *The basis change matrices are the same in either basis; i.e. we have* $[\sigma]_{new} = [\sigma]_{old}$.

We denote the matrix in the corollary by S. Similarly, we define $T := [\tau]_{\text{new}} = [\tau]_{\text{old}}$.

Now equation (4) has the form

$$[\varphi]_{\text{new}} = T^{-1}[\varphi]_{\text{old}}S.$$

This is the sought-after basis change formula for linear maps.

As a special case, we have:

Corollary 4.35. *If we have* $\varphi : V \to V$ *a linear transformation, then* $[\varphi]_{new} = S^{-1}[\varphi]_{old}S$.

This motivates the following definition:

Definition 4.36. Suppose $A, B \in M_n(\mathbb{F})$. We say A and B are *similar* (written $A \sim B$) if there is an invertible matrix S s.t. $B = S^{-1}AS$.

The preceding discussions show

Theorem 4.37. *The matrices associated with a linear transformation in different bases are similar.*

Examples 4.38. As in example (4.26), the following two matrices are similar, because they represent the same linear transformation:

$$\begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix} \sim \begin{bmatrix} 0 & -1 \\ 1 & 2\cos\alpha \end{bmatrix}.$$

Now we see that Theorem (4.27) is equivalent to the following statement:

Theorem 4.39. *Similar matrices have the same characteristic polynomial.*

Exercise 4.40. Prove this theorem.

We are often interested in deciding whether or not a given matrix is similar to a diagonal matrix.

Exercise 4.41. Prove that $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ is not similar to any diagonal matrix. (What type of linear transformation does this matrix describe? We have given it a name.)

Exercise 4.42. (a) If $A \in M_n(\mathbb{F})$, and $f_A(x) = (x - \lambda_1) \cdots (x - \lambda_n)$ and the λ_i are all distinct then $A \sim \operatorname{diag}(\lambda_1, \ldots, \lambda_n)$. (b) Prove that this conclusion will fail if we omit that assumption of distinctness of the λ_i.

This leads to the following important definition:

Definition 4.43. A matrix *diagonalizable* if it is similar to a diagonal matrix.

Exercise 4.44. Prove: the matrix of a linear transformation $\varphi : V \to V$ (with respect to any basis) is diagonalizable if and only if φ has an eigenbasis.

Not every matrix is diagonalizable, as exercise (4.41) shows. We have the following useful sufficient condition, which follows from exercise (4.42):

Exercise 4.45. If $A \in M_n(\mathbb{F})$, and A has n distinct eigenvalues, then A is diagonalizable.

Consider the triangular matrix $A := \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ & a_{22} & \cdots & a_{2n} \\ & & \ddots & \vdots \\ & & & a_{nn} \end{bmatrix}$. Its characteristic polynomial is $f_A(x) = (x - a_{11}) \cdots (x - a_{nn})$. Thus we now know

Theorem 4.46. *The eigenvalues of a triangular matrix are the diagonal entries.*

We have two ways of accounting for multiple eigenvalues.

Definition 4.47. The *algebraic multiplicity* of an eigenvalue λ is k if $(x - \lambda)^k \mid f_A(x)$, and $(x - \lambda)^{k+1} \nmid f_A(x)$.

Definition 4.48. The *geometric multiplicity* of an eigenvalue λ is the number of linearly independent eigenvectors for eigenvalue λ.

Examples 4.49. The central reflection in $\mathbb{R}^3$ has one eigenvalue -1, with geometric multiplicity 3.

Exercise 4.50. Prove that the geometric multiplicity of λ is equal to $\dim(U_\lambda)$ where $U_\lambda := \ker(\lambda I - A)$. We call U_λ the *eigensubspace* corresponding to λ; it consist of $\mathbf{0}$ and the eigenvectors for λ.

Exercise 4.51. Prove the the geometric multiplicity of λ is $\leq$ its algebraic multiplicity.

The next exercise characterizes diagonalizable matrices. Note that the splitting condition is always satisfied if $\mathbb{F} = \mathbb{C}$.

Exercise 4.52. If $f_A(x)$ splits into linear factor over $\mathbb{F}$, then the following are equivalent:

(i) A is diagonalizable;

(ii) for all λ, the algebraic multiplicity of λ is equal to its geometric multiplicity.

This last series of exercises will let us prove the Cayley-Hamilton Theorem.

Theorem 4.53 (Cayley-Hamilton). *If $A \in M_n(\mathbb{F})$, then $f_A(A) = 0$.*

We will prove the statement for matrices over $\mathbb{C}$. In particular, it holds for matrices with integer coefficients. It will then follow that the statements holds for matrices over any field.

Exercise 4.54. (i) Prove: if matrices A and B are similar and the Theorem holds for A then it holds for B.

(ii) Prove the Theorem for diagonal matrices.

(iii) Prove it for diagonalizable matrices.

(iv) Prove that if $\lim_{k\to\infty} A_k = A$, and the Theorem holds for all A_k then it also holds for A.

(v) Prove that every triangular matrix is the limit of a sequence of diagonalizable matrices. (Hint: distinct eigenvalues.)

(vi) Prove the Theorem for triangular matrices.

(vii) Finally, prove that every matrix is similar to a triangular matrix.