Math 121B — Linear Algebra

Neil Donaldson

Fall 2022

Linear Algebra, Stephen Friedberg, Arnold Insel & Lawrence Spence, 4th Ed 2003, Prentice Hall.

Review from 121A

We begin by recalling a few basic notions and notations.

Vector Spaces   Bold-face $\mathbf{v}$ denotes a *vector* in a *vector space V* over a *field* $\mathbb{F}$. A vector space is closed under *vector addition* and *scalar multiplication*

$$\forall \mathbf{v}_1, \mathbf{v}_2 \in V, \ \forall \lambda_1, \lambda_2 \in \mathbb{F}, \quad \lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2 \in V$$

Examples.   Here are four (families of) vector spaces over the field $\mathbb{R}$.

- $\mathbb{R}^2 = \{x\mathbf{i} + y\mathbf{j} : x, y \in \mathbb{R}\} = \{\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right) : x, y \in \mathbb{R}\}$ is a vector space over the field $\mathbb{R}$.

- $P_n(\mathbb{R})$; *polynomials* with degree $\leq n$ and coefficients in $\mathbb{R}$

- $P(\mathbb{R})$; polynomials over $\mathbb{R}$ with any degree.

- $C(\mathbb{R})$; continuous functions from $\mathbb{R}$ to $\mathbb{R}$.

Linear Combinations and Spans   Let $\beta \subseteq V$ be a subset of a vector space V over $\mathbb{F}$. A *linear combination* of vectors in β is any *finite* sum

$$\lambda_1 \mathbf{v}_1 + \cdots + \lambda_n \mathbf{v}_n$$

where $\lambda_j \in \mathbb{F}$ and $\mathbf{v}_j \in \beta$. The *span* of β comprises all linear combinations: this is a *subspace* of V.

Bases and Co-ordinates   A set $\beta \subseteq V$ is a *basis* of V if it has two properties:

Linear Independence Any linear combination yielding the zero vector is trivial; for distinct $\mathbf{v}_j \in \beta$,

$$\lambda_1 \mathbf{v}_1 + \cdots + \lambda_n \mathbf{v}_n = \mathbf{0} \implies \forall j, \lambda_j = 0$$

Spanning Set $V = \operatorname{Span} \beta$; *every* vector in V is a (*finite!*) linear combination of elements of β.

> **Theorem.** β is a basis of V $\iff$ every $\mathbf{v} \in V$ is a unique linear combination *of elements of β.*
> *The cardinality of all basis sets is identical; this is the* dimension $\dim_{\mathbb{F}} V$.

Example. $P_2(\mathbb{R})$ has standard basis $\beta = \{1, x, x^2\}$: every degree ≤ 2 polynomial is a unique as a linear combination $p(x) = a + bx + cx^2$ and so $\dim P_2(\mathbb{R}) = 3$. The real numbers a, b, c are the *co-ordinates* of p with respect to β; the *co-ordinate vector* of p is written

$$[p]_\beta = \begin{pmatrix} a \\ b \\ c \end{pmatrix}$$

Linearity and Linear Maps A function $T : V \to W$ between vector spaces V, W over the *same field* $\mathbb{F}$ is *($\mathbb{F}$-)linear* if it respects the linearity properties of V, W

$$\forall \mathbf{v}_1, \mathbf{v}_2 \in V, \ \forall \lambda_1, \lambda_2 \in \mathbb{F}, \quad T(\lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2) = \lambda_1 T(\mathbf{v}_1) + \lambda_2 T(\mathbf{v}_2)$$

We write $\mathcal{L}(V, W)$ for the set (indeed vector space!) of linear maps from V to W: this is shortened to $\mathcal{L}(V)$ if $V = W$. An *isomorphism* is an invertible/bijective linear map.

Theorem. *If* $\dim_{\mathbb{F}} V = n$ *and* β *is a basis of* V, *then the co-ordinate map* $\mathbf{v} \mapsto [\mathbf{v}]_\beta$ *is an isomorphism of vector spaces* $V \to \mathbb{F}^n$.

Matrices and Linear Maps If V, W are finite-dimensional, then any linear map $T : V \to W$ can be described using matrix multiplication.

Example. If $A = \begin{pmatrix} 2 & -1 \\ 0 & -1 \\ -4 & 3 \end{pmatrix}$, then the linear map $L_A : \mathbb{R}^2 \to \mathbb{R}^3$ (*left-multiplication by A*) is

$$L_A \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2 & -1 \\ 0 & -1 \\ -4 & 3 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2x - y \\ -y \\ 3y - 4x \end{pmatrix}$$

The linear map in fact *defines* the matrix A; we recover the columns of the matrix by feeding the standard basis vectors to the linear map.

$$\begin{pmatrix} 2 \\ 0 \\ -4 \end{pmatrix} = L_A \begin{pmatrix} 1 \\ 0 \end{pmatrix} \qquad \begin{pmatrix} -1 \\ -1 \\ 3 \end{pmatrix} = L_A \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

More generally, if $T \in \mathcal{L}(V, W)$ and $\beta = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ and $\gamma = \{\mathbf{w}_1, \ldots, \mathbf{w}_m\}$ are bases of V, W respectively, then the *matrix of T with respect to β and γ* is

$$[T]_\beta^\gamma = \left([T(\mathbf{v}_1)]_\gamma \ \cdots \ [T(\mathbf{v}_n)]_\gamma \right) \in M_{m \times n}(\mathbb{F})$$

whose j^{th} column is obtained by feeding the j^{th} basis vector of β to T and taking its co-ordinate vector with respect to γ. This fits naturally with the co-ordinate isomorphisms

$$T(\mathbf{v}) = \mathbf{w} \iff [T]_\beta^\gamma [\mathbf{v}]_\beta = [\mathbf{w}]_\gamma$$

There are two special cases when $V = W$:

- If $\beta = \gamma$, then we simply write $[T]_\beta$ instead of $[T]_\beta^\beta$.

- If $T = I$ is the identity map, then $Q_\beta^\gamma := [I]_\beta^\gamma$ is the *change of co-ordinate matrix* from β to γ.

Being able to convert linear maps into matrix multiplication is a central skill in linear algebra. Test your comfort by working through the following; if everything feels familiar, you should consider yourself in a good place as far as pre-requisites are concerned!

Example. Let $T : P_2(\mathbb{R}) \to P_1(\mathbb{R})$ be the linear map defined by *differentiation*

$$T(a + bx + cx^2) = b + 2cx \tag{$*$}$$

The standard bases of $P_2(\mathbb{R})$ and $P_1(\mathbb{R})$ are, respectively, $\beta = \{1, x, x^2\}$ and $\gamma = \{1, x\}$. Observe that

$$[T(1)]_\gamma = [0]_\gamma = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \qquad [T(x)]_\gamma = [1]_\gamma = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad [T(x^2)]_\gamma = [2x]_\gamma = \begin{pmatrix} 0 \\ 2 \end{pmatrix}$$

$$\implies [T]_\beta^\gamma = \left([T(1)]_\gamma \; [T(x)]_\gamma \; [T(x^2)]_\gamma \right) = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

Written in co-ordinates, we see the original linear map $(*)$

$$\left[T(a + bx + cx^2) \right]_\gamma = [T]_\beta^\gamma [a + bx + cx^2]_\beta = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} b \\ 2c \end{pmatrix} = [b + 2cx]_\gamma \tag{$\dagger$}$$

1. $\eta = \{1 + x, x + x^2, x^2 + 1\}$ is also a basis of $P_2(\mathbb{R})$. Show that

$$[T]_\eta^\gamma = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 2 & 2 \end{pmatrix}$$

2. As in $(\dagger)$ above, the matrix multiplication

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 2 & 2 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} a + b \\ 2b + 2c \end{pmatrix}$$

corresponds to an equation $[T(p)]_\gamma = [T]_\eta^\gamma [p]_\eta$ for some polynomial $p(x)$; what is $p(x)$ in terms of a, b, c?

3. Find the change of co-ordinate matrix Q_η^β and check that the matrices of T are related by

$$[T]_\eta^\gamma = [T]_\beta^\gamma Q_\eta^\beta$$

3

1 Diagonalizability & the Cayley–Hamilton Theorem

1.1 Eigenvalues, Eigenvectors & Diagonalization (Review)

> **Definition 1.1.** Suppose V is a vector space over $\mathbb{F}$ and $T \in \mathcal{L}(V)$. A *non-zero* $\mathbf{v} \in V$ is an *eigenvector* of T with *eigenvalue* $\lambda \in \mathbb{F}$ (together an *eigenpair*) if
>
> $$T(\mathbf{v}) = \lambda \mathbf{v}$$
>
> For matrices, the eigenvalues/vectors of $A \in M_n(\mathbb{F})$ are precisely those of $L_A \in \mathcal{L}(\mathbb{F}^n)$.
>
> Suppose λ is an eigenvalue of T;
>
> 1. The *eigenspace* of λ is the nullspace $E_\lambda := \mathcal{N}(T - \lambda I)$.
> 2. The *geometric multiplicity* of λ of the dimension $\dim E_\lambda$.
>
> We say that T is *diagonalizable* if there exists a basis of eigenvectors; an *eigenbasis*.

We start by recalling a couple of basic facts, the first of which is easily proved by induction.

> **Lemma 1.2.** If $\mathbf{v}_1, \ldots, \mathbf{v}_k$ are eigenvectors corresponding to distinct eigenvalues, then $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is linearly independent.
>
> Moreover, if $\dim_\mathbb{F} V = n$ and $T \in \mathcal{L}(V)$ has n distinct eigenvalues, then T is diagonalizable.

Eigenvalues and Eigenvectors in finite dimensions

If $\dim_\mathbb{F} V = n$ and ϵ is a basis, then the eigenvector definition is equivalent to a matrix equation

$$[T]_\epsilon [\mathbf{v}]_\epsilon = \lambda [\mathbf{v}]_\epsilon$$

In such a situation, T being diagonalizable means $\exists \beta$ such that $[T]_\beta$ is a *diagonal matrix*

$$[T]_\beta = \begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{pmatrix}$$

Thankfully there is systematic way to find eigenvalues and eigenvectors in finite-dimensions:

1. Choose any basis ϵ of V and compute the matrix $A = [T]_\epsilon \in M_n(\mathbb{F})$.

2. Observe that

$$\begin{aligned} \lambda \in \mathbb{F} \text{ is an eigenvalue} &\iff \exists [\mathbf{v}]_\epsilon \in \mathbb{F}^n \setminus \{\mathbf{0}\} \text{ such that } A[\mathbf{v}]_\epsilon = \lambda [\mathbf{v}]_\epsilon \\ &\iff \exists [\mathbf{v}]_\epsilon \in \mathbb{F}^n \setminus \{\mathbf{0}\} \text{ such that } (A - \lambda I)[\mathbf{v}]_\epsilon = \mathbf{0} \\ &\iff \det(A - \lambda I) = 0 \end{aligned}$$

This last is a degree n polynomial equation whose roots are the eigenvalues.

3. For each eigenvalue λ_j, compute the eigenspace $E_{\lambda_j} = \mathcal{N}(T - \lambda_j I)$ to find the eigenvectors. Remember that E_{λ_j} is a subspace of the original vector space V, so translate back if necessary!

Definition 1.3. The *characteristic polynomial of* $T \in \mathcal{L}(V)$ is the degree-n polynomial

$$p(t) := \det(T - tI)$$

The eigenvalues of T are precisely the solutions to the *characteristic equation* $p(t) = 0$.

Examples 1.4. 1. $A = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ has characteristic polynomial $p(t) = t^2 + 1 = (t + i)(t - i)$. As a linear map $L_A \in \mathcal{L}(\mathbb{R}^2)$, A has *no eigenvalues* and *no eigenvectors*!

As a linear map $L_A \in \mathcal{L}(\mathbb{C}^2)$, we have two eigenvalues $\pm i$. Indeed

$$(A - iI)\mathbf{v} = \begin{pmatrix} -i & -1 \\ 1 & -i \end{pmatrix} \mathbf{v} \implies E_i = \text{Span} \begin{pmatrix} i \\ 1 \end{pmatrix}$$

and similarly $E_{-i} = \text{Span} \left(\begin{smallmatrix} -i \\ 1 \end{smallmatrix}\right)$. We therefore have an eigenbasis $\beta = \{\left(\begin{smallmatrix} i \\ 1 \end{smallmatrix}\right), \left(\begin{smallmatrix} -i \\ 1 \end{smallmatrix}\right)\}$ (of $\mathbb{C}^2$), with respect to which

$$[L_A]_\beta = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$$

2. Let $T \in \mathcal{L}(P_2(\mathbb{R}))$ be defined by

$$T(f)(x) = f(x) + (x - 1)f'(x)$$

With respect to the standard basis $\epsilon = \{1, x, x^2\}$, we have the non-diagonal matrix

$$A = [T]_\epsilon = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 2 & -2 \\ 0 & 0 & 3 \end{pmatrix} \implies p(t) = \det(A - tI) = (1 - t)(2 - t)(3 - t)$$

With three distinct eigenvalues, T is diagonalizable. To find the eigenvectors, compute the nullspaces:

$$\lambda_1 = 1: \quad \mathbf{0} = (A - \lambda_1 I)[\mathbf{v}_1]_\epsilon = \left(\begin{smallmatrix} 0 & -1 & 0 \\ 0 & 1 & -2 \\ 0 & 0 & 2 \end{smallmatrix}\right) [\mathbf{v}_1]_\epsilon \implies [\mathbf{v}_1]_\epsilon \in \text{Span} \left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix}\right) \implies E_1 = \text{Span}\{1\}$$

$$\lambda_2 = 2: \quad A - \lambda_2 I = \left(\begin{smallmatrix} -1 & -1 & 0 \\ 0 & 0 & -2 \\ 0 & 0 & 1 \end{smallmatrix}\right) \implies [\mathbf{v}_2]_\epsilon \in \text{Span} \left(\begin{smallmatrix} 1 \\ -1 \\ 0 \end{smallmatrix}\right) \implies E_2 = \text{Span}\{1 - x\}$$

$$\lambda_3 = 3: \quad A - \lambda_3 I = \left(\begin{smallmatrix} -2 & -1 & 0 \\ 0 & -1 & -2 \\ 0 & 0 & 0 \end{smallmatrix}\right) \implies [\mathbf{v}_3]_\epsilon \in \text{Span} \left(\begin{smallmatrix} 1 \\ -2 \\ 1 \end{smallmatrix}\right) \implies E_3 = \text{Span}\{1 - 2x + x^2\}$$

Making a sensible choice of non-zero eigenvectors, we obtain an eigenbasis, with respect to which the linear map is necessarily diagonal

$$\beta = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\} = \{1, 1 - x, 1 - 2x + x^2\} = \{1, 1 - x, (1 - x)^2\}$$

$$T\big(a + b(1 - x) + c(1 - x)^2\big) = a + 2b(1 - x) + 3c(1 - x)^2$$

$$[T]_\beta = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix}$$

Conditions for diagonalizability of finite-dimensional operators

We now borrow a little terminology from the theory of polynomials.

Definition 1.5. Let $\mathbb{F}$ be a field and $p(t)$ a polynomial with coefficients in $\mathbb{F}$.

1. Let $\lambda \in \mathbb{F}$ be a root; $p(\lambda) = 0$. The *algebraic multiplicity* $\mathrm{mult}(\lambda)$ is the largest power of $\lambda - t$ to divide $p(t)$. Otherwise said, there exists[1] some polynomial $q(t)$ such that

$$p(t) = (\lambda - t)^{\mathrm{mult}(\lambda)} q(t) \quad \text{and} \quad q(\lambda) \neq 0$$

2. We say that $p(t)$ *splits* over $\mathbb{F}$ if it factorizes completely into linear factors; equivalently $\exists a, \lambda_1, \ldots, \lambda_k \in \mathbb{F}$ such that

$$p(t) = a(\lambda_1 - t)^{m_1} \cdots (\lambda_k - t)^{m_k}$$

When $p(t)$ splits, the algebraic multiplicities sum to the degree n of the polynomial

$$n = m_1 + \cdots + m_k$$

Of course, we are most interested when $p(t)$ is the *characteristic polynomial* of a linear map $\mathrm{T} \in \mathcal{L}(V)$. If such a polynomial splits, then $a = 1$ and $\lambda_1, \ldots, \lambda_k$ are necessarily the (distinct) eigenvalues of T.

Example 1.6. The field matters! For instance $p(t) = t^2 + 1 = (t - i)(t + i) = -(i - t)(-i - t)$ splits over $\mathbb{C}$ but not over $\mathbb{R}$. Its roots are plainly $\pm i$.

For the purposes of review, we state the main result; this will be proved in the next section.

Theorem 1.7. *Let V be finite-dimensional. A linear map $\mathrm{T} \in \mathcal{L}(V)$ is diagonalizable if and only if,*

1. *Its characteristic polynomial splits over $\mathbb{F}$, and,*

2. *The geometric and algebraic multiplicities of each eigenvalue are equal; $\dim E_{\lambda_j} = \mathrm{mult}(\lambda_j)$.*

Example 1.8. The matrix $A = \left(\begin{smallmatrix} 3 & 1 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 5 \end{smallmatrix}\right)$ is easily seen to have eigenvalues $\lambda_1 = 3$ and $\lambda_2 = 5$. Indeed

$$p(t) = (3 - t)^2(5 - t), \qquad \mathrm{mult}(3) = 2, \qquad \mathrm{mult}(5) = 1$$

$$E_3 = \mathrm{Span}\begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \qquad E_5 = \mathrm{Span}\begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \qquad \dim E_3 = \dim E_5 = 1$$

This matrix is *non-diagonalizable* since $\dim E_3 = 1 \neq 2 = \mathrm{mult}(3)$.

Everything prior to this *should* be review. If it feels very unfamiliar, revisit your notes from 121A, particularly sections 5.1 and 5.2 of the textbook.

---

[1] The existence follows from Descartes *factor theorem* and the *division algorithm* for polynomials.

Exercises 1.1 1. For each matrix over $\mathbb{R}$; find its characteristic polynomial, its eigenvalues/spaces, and its algebraic and geometric multiplicities; decide if it is diagonalizable.

(a) $A = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 1 \\ 0 & 0 & 0 & 3 \end{pmatrix}$ (b) $B = \begin{pmatrix} -1 & 6 & 0 & 0 \\ -2 & 6 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix}$

2. Suppose A is a *real* matrix with eigenpair $(\lambda, \mathbf{v})$. If $\lambda \notin \mathbb{R}$ show that $(\overline{\lambda}, \overline{\mathbf{v}})$ is also an eigenpair.

3. Show that the characteristic polynomial of $A = \begin{pmatrix} 3 & -4 \\ 4 & 3 \end{pmatrix}$ does not split over $\mathbb{R}$. Diagonalize A over $\mathbb{C}$.

4. Give an example of a 2×2 matrix whose entries are *rational numbers* and whose characteristic polynomial splits over $\mathbb{R}$, but not over $\mathbb{Q}$.

5. Diagonalize $L_C \in \mathcal{L}(\mathbb{C}^2)$ where $C = \begin{pmatrix} 2i & 1 \\ 2 & 0 \end{pmatrix}$.

6. If $p(t)$ splits, explain why

$$\det T = \lambda_1^{\text{mult}(\lambda_1)} \cdots \lambda_k^{\text{mult}(\lambda_k)}$$

where $\lambda_1, \ldots, \lambda_k$ are the distinct eigenvalues of T.

7. Suppose $T \in \mathcal{L}(V)$ is invertible with eigenvalue λ. Prove that λ^{-1} is an eigenvalue of T^{-1} with the same eigenspace E_λ. If T is diagonalizable, prove that T^{-1} is also diagonalizable.

8. If V is finite-dimensional and $T \in \mathcal{L}(V)$, we may define $\det T$ to equal $\det[T]_\beta$, where β is any basis of V. Explain why the choice of basis does not matter; that is, if γ is any other basis of V, we have $\det[T]_\gamma = \det[T]_\beta$.

1.2 Invariant Subspaces and the Cayley–Hamilton Theorem

The proof of Theorem 1.7 is facilitated by a new concept, of which eigenspaces are a special case.

Definition 1.9. Suppose $T \in \mathcal{L}(V)$. A subspace W of V is T-*invariant* if $T(W) \subseteq W$. In such a case, the *restriction* of T to W is the linear map

$$T_W : W \to W : \mathbf{w} \mapsto T(\mathbf{w})$$

Examples 1.10. 1. The trivial subspace $\{\mathbf{0}\}$ and the entire vector space V are invariant for any linear map $T \in \mathcal{L}(V)$.

2. Every eigenspace is invariant; if $\mathbf{v} \in E_\lambda$, then $T(\mathbf{v}) = \lambda\mathbf{v} \in E_\lambda$.

3. Continuing Example 1.8, if $A = \left(\begin{smallmatrix} 3 & 1 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 5 \end{smallmatrix} \right)$ then $W = \operatorname{Span}\{\mathbf{i}, \mathbf{j}\}$ is an invariant subspace for the linear map L_A. Indeed

$$A(x\mathbf{i} + y\mathbf{j}) = (3x + y)\mathbf{i} + 3y\mathbf{j} \in W$$

W is an example of a *generalized eigenspace*; we'll study these properly at the end of term.

To prove our diagonalization criterion, we need to see how to factorize the characteristic polynomial. It turns out that factors of $p(t)$ correspond to T-invariant subspaces!

Example 1.11. $W = \operatorname{Span}\{\mathbf{i}, \mathbf{j}\}$ is an invariant subspace of $A = \left(\begin{smallmatrix} 1 & 2 & 4 \\ 0 & 3 & 1 \\ 0 & 0 & 2 \end{smallmatrix} \right) \in M_3(\mathbb{R})$. With respect to the standard basis, the restriction $[L_A]_W$ has matrix $\left(\begin{smallmatrix} 1 & 2 \\ 0 & 3 \end{smallmatrix} \right)$. The characteristic polynomial $p_W(t)$ of the restriction is plainly a factor of the whole,

$$p(t) = (1 - t)(2 - t)(3 - t) = (2 - t)p_W(t)$$

Theorem 1.12. *Suppose* $T \in \mathcal{L}(V)$*, that* $\dim V$ *is finite and that* W *is a T-invariant subspace of* V. *Then the characteristic polynomial of the restriction* T_W *divides that of* T.

The proof simply abstracts the approach of the example.

Proof. Extend a basis β_W of W to a basis β of V. Since $T(\mathbf{w}) \in \operatorname{Span}\beta_W$ for each $\mathbf{w} \in W$, we see that the matrix of $[T]$ has block form

$$[T]_\beta = \left(\begin{array}{c|c} A & B \\ \hline O & C \end{array} \right) \implies p(t) = \det(A - tI)\det(C - tI) = p_W(t)\det(C - tI)$$

where $p_W(t)$ is the characteristic polynomial, and $A = [T_W]_{\beta_W}$ the matrix of the restriction T_W. ∎

Corollary 1.13. *If* λ *is an eigenvalue of* T*, then* $T_{E_\lambda} = \lambda I_{E_\lambda}$ *is a multiple of the identity, whence,*

1. *The characteristic polynomial of the restriction* T_{E_λ} *is* $p_\lambda(t) = (\lambda - t)^{\dim E_\lambda}$.

2. $p_\lambda(t)$ *divides the characteristic polynomial of* T*. In particular* $\dim E_\lambda \le \operatorname{mult}(\lambda)$.

We are now in a position to state and prove an extended version of Theorem 1.7.

Theorem 1.14. *Suppose* $\dim_{\mathbb{F}} V = n$ *and that* $T \in \mathcal{L}(V)$ *has distinct eigenvalues* $\lambda_1, \ldots, \lambda_k$. *The following are equivalent:*

1. T *is diagonalizable.*

2. *The characteristic polynomial splits over* $\mathbb{F}$ *and* $\dim E_{\lambda_j} = \text{mult}(\lambda_j)$ *for each* j*; indeed*

$$p(t) = p_{\lambda_1}(t) \cdots p_{\lambda_k}(t) = (\lambda_1 - t)^{\dim E_{\lambda_1}} \cdots (\lambda_k - t)^{\dim E_{\lambda_k}}$$

3. $\sum_{j=1}^{k} \dim E_{\lambda_j} = n$

4. $V = E_{\lambda_1} \oplus \cdots \oplus E_{\lambda_k}$

Example 1.15. $A = \begin{pmatrix} 7 & 0 & -12 \\ 0 & 1 & 0 \\ 2 & 0 & -3 \end{pmatrix}$ is diagonalizable. Indeed $p(t) = (1-t)^2(3-t)$ splits, and we have

λ	1	3
$\text{mult}(\lambda)$	2	1
E_λ	$\text{Span}\left\{ \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \right\}$	$\text{Span} \begin{pmatrix} 3 \\ 0 \\ 1 \end{pmatrix}$
$\dim E_\lambda$	2	1

and $\mathbb{R}^4 = E_1 \oplus E_3$

With respect to the eigenbasis $\beta = \left\{ \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ 0 \\ 1 \end{pmatrix} \right\}$, the map is diagonal $[L_A]_\beta = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 3 \end{pmatrix}$

Proof. $(1 \Rightarrow 2)$ If T is diagonalizable with eigenbasis β, then $[T]_\beta$ is diagonal. But then

$$p(t) = (\lambda_1 - t)^{m_1} \cdots (\lambda_k - t)^{m_k}$$

splits and $\sum \text{mult}(\lambda_i) = n$. The cardinality n of an eigenbasis is *at most* $\sum \dim E_{\lambda_i}$ since every element is an (independent) eigenvector. By Corollary 1.13 $(\dim E_{\lambda_j} \leq \text{mult}(\lambda_j))$ we see that

$$n \leq \sum \dim E_{\lambda_j} \leq \sum \text{mult}(\lambda_j) = n \implies \forall j, \dim E_{\lambda_j} = \text{mult}(\lambda_j)$$

whence the inequalities are *equalities* with each pair equal $\dim E_{\lambda_j} = \text{mult}(\lambda_j)$

$(2 \Rightarrow 3)$ $p(t)$ splits $\implies n = \sum \text{mult}(\lambda_j) = \sum \dim E_{\lambda_j}$

$(3 \Rightarrow 4)$ Assume $E_{\lambda_1} \oplus \cdots \oplus E_{\lambda_j}$ exists.[2] If $(\lambda_{j+1}, \mathbf{v}_{j+1})$ is an eigenpair, then $\mathbf{v}_{j+1} \notin E_{\lambda_1} \oplus \cdots \oplus E_{\lambda_j}$ for otherwise this would contradict Lemma 1.2.

By induction, $E_{\lambda_1} \oplus \cdots \oplus E_{\lambda_k}$ exists; by assumption it has dimension $n = \dim V$ and therefore *equals* V.

$(4 \Rightarrow 1)$ For each j, choose a basis β_j of E_{λ_j}. Then $\beta := \beta_1 \cup \cdots \cup \beta_k$ is a basis of V consisting of eigenvectors of T; an eigenbasis. $\blacksquare$

---

[2] Distinct eigenspaces have trivial intersection: $i_1 \neq i_2 \leq j \implies E_{i_1} \cap E_{i_2} = \{\mathbf{0}\}$.

T-cyclic Subspaces and the Cayley–Hamilton Theorem

We finish this chapter by introducing a general family of invariant subspaces and using them to prove a startling result.

> **Definition 1.16.** Let $T \in \mathcal{L}(V)$ and let $\mathbf{v} \in V$. The T-*cyclic subspace* generated by $\mathbf{v}$ is the span
>
> $$\langle \mathbf{v} \rangle = \text{Span}\{\mathbf{v}, T(\mathbf{v}), T^2(\mathbf{v}), \ldots\}$$

Example 1.17. Recalling Example 1.10.3 Let $A = \left(\begin{smallmatrix} 3 & 1 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 5 \end{smallmatrix} \right)$, and $\mathbf{v} = \mathbf{i} + \mathbf{k}$. It is easy to see that

$$A\mathbf{v} = 3\mathbf{i} + 5\mathbf{k}, \quad A^2\mathbf{v} = 9\mathbf{i} + 25\mathbf{k}, \quad \ldots, \quad A^m\mathbf{v} = 3^m\mathbf{i} + 5^m\mathbf{k}$$

all of which lie in $\text{Span}\{\mathbf{i}, \mathbf{k}\}$. Plainly this is the L_A-cyclic subspace $\langle \mathbf{i} + \mathbf{k} \rangle$.

The proof of the following basic result is left as an exercise.

> **Lemma 1.18.** $\langle \mathbf{v} \rangle$ *is the smallest* T-*invariant subspace of* V *containing* $\mathbf{v}$, *specifically:*
>
> 1. $\langle \mathbf{v} \rangle$ *is* T-*invariant.*
> 2. *If* $W \leq V$ *is* T-*invariant and* $\mathbf{v} \in W$, *then* $\langle \mathbf{v} \rangle \leq W$.
> 3. $\dim \langle \mathbf{v} \rangle = 1 \iff \mathbf{v}$ *is an eigenvector of* T.

We were lucky in the example that the general form $A^m\mathbf{v}$ was so clear. It is helpful to develop a more precise test for identifying the dimension and a basis of a T-cyclic subspace.

Suppose a T-cyclic subspace $\langle \mathbf{v} \rangle = \text{Span}\{\mathbf{v}, T(\mathbf{v}), T^2(\mathbf{v}), \ldots\}$ has finite dimension.[3] Let $k \geq 1$ be maximal such that the set

$$\{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\}$$

is linearly independent.

- If k doesn't exist, the *infinite* linearly independent set $\{\mathbf{v}, T(\mathbf{v}), \ldots\}$ contradicts $\dim \langle \mathbf{v} \rangle < \infty$.

- By the maximality of k, $T^k(\mathbf{v}) \in \text{Span}\{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\}$; by induction this extends to

$$j \geq k \implies T^j(\mathbf{v}) \in \text{Span}\{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\}$$

It follows that $\langle \mathbf{v} \rangle = \text{Span}\{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\}$, and we've proved a useful criterion.

> **Theorem 1.19.** *Suppose* $\mathbf{v} \neq \mathbf{0}$, *then*
>
> $$\dim \langle \mathbf{v} \rangle = k \iff \{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\} \text{ is a basis of } \langle \mathbf{v} \rangle$$
> $$\iff k \text{ is maximal such that } \{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\} \text{ is linearly independent}$$
> $$\iff k \text{ is minimal such that } T^k(\mathbf{v}) \in \text{Span}\{\mathbf{v}, T(\mathbf{v}), \ldots, T^{k-1}(\mathbf{v})\}$$

---

[3]Necessarily the situation if $\dim V < \infty$, when we are thinking about characteristic polynomials.

Examples 1.20. 1. According to the Theorem, in Example 1.17 we need only have noticed

- $\mathbf{v} = \mathbf{i} + \mathbf{k}$ and $A\mathbf{v} = 3\mathbf{i} + 5\mathbf{k}$ are linearly independent.
- That $A^2(\mathbf{i} + \mathbf{k}) = 9\mathbf{i} + 25\mathbf{k} \in \text{Span}\{\mathbf{v}, A\mathbf{v}\}$.

We could then conclude that $\langle \mathbf{v} \rangle = \text{Span}\{\mathbf{v}, A\mathbf{v}\}$ has dimension 2.

2. Let $T(p(x)) = 3p(x) - p''(x)$ viewed as a linear map $T \in \mathcal{L}(P_2(\mathbb{R}))$ and consider the T-cyclic subspace generated by the polynomial $p(x) = x^2$

$$T(x^2) = 3x^2 - 2, \quad T^2(x^2) = T(3x^2 - 2) = 3(3x^2 - 2) - 6 = 9x^2 - 12, \quad \ldots$$

Observe that $\{x^2, T(x^2)\}$ is linearly independent, but that

$$T^2(x^2) = 9x^2 - 12 = -9x^2 + 6(3x^2 - 2) \in \text{Span}\{x^2, T(x^2)\}$$

We conclude that $\dim \langle x^2 \rangle = 2$. An alternative basis for $\langle x^2 \rangle$ is plainly $\{1, x^2\}$.

We finish by considering the interaction of a T-cyclic subspace with the characteristic polynomial. Surprisingly, the coefficients of the characteristic polynomial and the linear combination coincide.

Continuing the Example, if $W = \langle x^2 \rangle$ and $\beta_W = \{x^2, T(x^2)\} = \{x^2, 3x^2 - 2\}$, then

$$[T_W]_{\beta_W} = \begin{pmatrix} 0 & -9 \\ 1 & 6 \end{pmatrix} \implies p_W(t) = t^2 - 6t + 9$$

Theorem 1.21. *Let* $T \in \mathcal{L}(V)$ *and suppose* $W = \langle \mathbf{w} \rangle$ *has* $\dim W = k$ *with basis*

$$\beta_W = \{\mathbf{w}, T(\mathbf{w}), \ldots, T^{k-1}(\mathbf{w})\}$$

in accordance with Theorem 1.19, then

1. *If* $T^k(\mathbf{w}) + a_{k-1}T^{k-1}(\mathbf{w}) + \cdots + a_0\mathbf{w} = \mathbf{0}$*, then the characteristic polynomial of* T_W *is*

$$p_W(t) = (-1)^k \left(t^k + a_{k-1}t^{k-1} + \cdots + a_1 t + a_0 \right)$$

2. $p_W(T_W) = 0$ *is the zero map on* W*.*

Proof. 1. This is an exercise.

2. Write $S \in \mathcal{L}(V)$ for the linear map

$$S := p_W(T) = (-1)^k (T^k + a_{k-1}T^{k-1} + \cdots + a_0 I)$$

Part 1 says $S(\mathbf{w}) = \mathbf{0}$. Since S is a polynomial in T, it commutes with all powers of T:

$$\forall j, \ S(T^j(\mathbf{w})) = T^j(S(\mathbf{w})) = \mathbf{0}$$

Since S is zero on the basis β_W of W, we see that S_W is the zero function. ∎

With a little sneakiness, we can drop the W's in the second part of the Theorem and observe an intimate relation between a linear map and its characteristic polynomial.

> **Corollary 1.22 (Cayley–Hamilton).** *If V is finite-dimensional, then* $\mathrm{T} \in \mathcal{L}(V)$ *satisfies its characteristic polynomial;* $p(\mathrm{T}) = 0$.

Proof. Let $\mathbf{w} \in V$ and consider the cyclic subspace $W = \langle \mathbf{w} \rangle$ generated by $\mathbf{w}$. By Theorem 1.12,

$$p(t) = q_W(t)p_W(t)$$

for some polynomial q_W. But the previous result says that $p_W(\mathrm{T})(\mathbf{w}) = \mathbf{0}$, whence

$$p(\mathrm{T})(\mathbf{w}) = \mathbf{0}$$

Since we may apply this reasoning to *any* $\mathbf{w} \in V$, we conclude that $p(\mathrm{T})$ is the zero function. ∎

Examples 1.23. 1. $A = \begin{pmatrix} 2 & 1 \\ 3 & 4 \end{pmatrix}$ has $p(t) = t^2 - 6t + 5$ and we confirm:

$$A^2 - 6A = \begin{pmatrix} 7 & 6 \\ 18 & 19 \end{pmatrix} - 6 \begin{pmatrix} 2 & 1 \\ 3 & 4 \end{pmatrix} = -5I$$

It may seem like a strange thing to do for this matrix, but the characteristic equation can be used to calculate the inverse of A:

$$A^2 - 6A + 5I = 0 \implies A(A - 6I) = -5I \implies A^{-1} = \frac{1}{5}(6I - A) = \frac{1}{5}\begin{pmatrix} 4 & -1 \\ -3 & 2 \end{pmatrix}$$

2. We use the Cayley–Hamilton Theorem to compute A^4 when

$$A = \begin{pmatrix} 2 & -1 & \frac{8}{3} \\ 0 & 1 & -6 \\ 0 & 0 & 2 \end{pmatrix}$$

The characteristic polynomial is

$$p(t) = (2 - t)^2(1 - t) = 4 - 8t + 5t^2 - t^3$$

By Cayley–Hamilton,

$$
\begin{aligned}
A^4 = AA^3 &= A(5A^2 - 8A + 4I) \\
&= 5A^3 - 8A^2 + 4A = 5(5A^2 - 8A + 4I) - 8A^2 + 4A \\
&= 17A^2 - 36A + 20I = 17\begin{pmatrix} 4 & -3 & \frac{50}{3} \\ 0 & 1 & -18 \\ 0 & 0 & 4 \end{pmatrix} - 36\begin{pmatrix} 2 & -1 & \frac{8}{3} \\ 0 & 1 & -6 \\ 0 & 0 & 2 \end{pmatrix} + 20\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \\
&= \begin{pmatrix} 16 & -15 & \frac{562}{3} \\ 0 & 1 & -90 \\ 0 & 0 & 16 \end{pmatrix}
\end{aligned}
$$

3. Recall Example 1.4.2, where the linear map $T(f(x)) = f(x) + (x-1)f'(x)$ had

$$p(t) = (1-t)(2-t)(3-t) = -t^3 + 6t^2 - 11t + 6$$

By Cayley–Hamilton, $T^3 = 6T^2 - 11T + 6I$. You can check this explicitly, after first computing

$$T^2(f(x)) = f(x) + 3(x-1)f'(x) + (x-1)^2 f''(x), \text{ etc.}$$

Cayley–Hamilton can also be used to simplify higher powers of T and even to compute the inverse!

$$I = \frac{1}{6}(T^3 - 6T^2 + 11T) \implies T^{-1} = \frac{1}{6}(T^2 - 6T + 11I)$$

$$\implies T^{-1}(f(x)) = f(x) - \frac{1}{2}(x-1)f'(x) + \frac{1}{6}(x-1)^2 f''(x)$$

Exercises 1.2   1. For the linear map $T = L_A : \mathbb{R}^3 \to \mathbb{R}^3$ where $A = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 2 & 4 \\ 0 & 0 & 2 \end{pmatrix}$ find the T-cyclic subspace generated by the standard basis vector $\mathbf{e}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$.

2. Let $T = L_A$, where $A = \begin{pmatrix} 1 & 2 & 4 \\ 0 & 3 & 1 \\ 0 & 0 & 2 \end{pmatrix}$ and let $\mathbf{v} = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}$. Compute $T(\mathbf{v})$ and $T^2(\mathbf{v})$. Hence describe the T-cyclic subspace $\langle \mathbf{v} \rangle$ and its dimension.

3. Given $A = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 1 \\ 0 & 0 & 0 & 3 \end{pmatrix}$, find *two distinct* L_A-invariant subspaces $W \le \mathbb{R}^4$ such that $\dim W = 3$.

4. Suppose that W and X are T-invariant subspaces of V. Prove that the sum

$$W + X = \{\mathbf{w} + \mathbf{x} : \mathbf{w} \in W, \mathbf{x} \in X\}$$

is also T-invariant.

5. Prove Lemma 1.18.

6. Give an example of an infinite-dimensional vector space V, a linear map $T \in \mathcal{L}(V)$, and a vector $\mathbf{v}$ such that $\langle \mathbf{v} \rangle = V$.

7. Let $\beta = \{\sin x, \cos x, 2x \sin x, 3x \cos x\}$ and $T = \frac{d}{dx} \in \mathcal{L}(\text{Span}\,\beta)$. Plainly the subspace $W := \text{Span}\{\sin x, \cos x\}$ is T-invariant. Compute the matrices $[T]_\beta$ and $[T_W]_{\beta_W}$ and observe that

$$p(t) = (p_W(t))^2$$

8. Verify explicitly that $A = \begin{pmatrix} 2 & 3 \\ 0 & 2 \end{pmatrix}$ satisfies its characteristic polynomial.

9. Check the details of Example 1.23.3 and evaluate T^4 as a linear combination of I, T and T^2. In particular, check the evaluation of $T^{-1}(f(x))$.

10. Suppose a, b are constants with $a \neq 0$ and define $T(f(x)) = af(x) + bf'(x)$.

   (a) Find an expression for the inverse $T^{-1}(f(x))$ if $T \in \mathcal{L}(P_1(\mathbb{R}))$

   (b) Find an expression for the inverse $T^{-1}(f(x))$ if $T \in \mathcal{L}(P_2(\mathbb{R}))$

   Your answers should be written in terms of f and its derivatives.

11. Let $T(f)(x) = f'(x) + \frac{1}{x}\int_0^x f(t)\, dt$ be a linear map $T \in \mathcal{L}(P_2(\mathbb{R}))$.

   (a) Find the characteristic polynomial of T and identify its eigenspaces. Is T diagonalizable?

   (b) Find $a, b, c \in \mathbb{R}$ such that $T^3 = aT^2 + bT + cI$.

   (c) What are $\dim \mathcal{L}(P_2(\mathbb{R}))$ and $\dim \operatorname{Span}\{T^k : k \in \mathbb{N}_0\}$? Explain.

12. If $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ has non-zero determinant, use the Cayley–Hamilton Theorem to obtain the usual expression for A^{-1}.

13. Recall Examples 1.10.3, 1.17, and 1.20.1 with $A = \left(\begin{smallmatrix} 3 & 1 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 5 \end{smallmatrix}\right)$.

   (a) If $\mathbf{v} = \left(\begin{smallmatrix} x \\ y \\ z \end{smallmatrix}\right)$, show that $\det(\mathbf{v}, A\mathbf{v}, A^2\mathbf{v}) = -4y^2 z$

   (b) Hence determine *all* L_A-cyclic subspaces of $\mathbb{R}^3$.

14. (a) Consider Example 1.20.2 where $T \in \mathcal{L}(P_2(\mathbb{R}))$ is defined by $T(p(x)) = 3p(x) - p''(x)$. Prove that all T-cyclic subspaces have dimension ≤ 2.

   (b) What if we instead consider $S \in \mathcal{L}(P_2(\mathbb{R})$ defined by $S(p(x)) = 3p(x) - p'(x)$?

15. We prove part 1 of Theorem 1.21.

   (a) Explain why the matrix of T_W with respect to the basis β_W is

$$[T_W]_{\beta_W} = \begin{pmatrix} 0 & 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & 0 & & 0 & -a_1 \\ 0 & 1 & 0 & & 0 & -a_2 \\ \vdots & & & \ddots & & \vdots \\ 0 & 0 & 0 & & 0 & -a_{k-2} \\ 0 & 0 & 0 & \cdots & 1 & -a_{k-1} \end{pmatrix} \in M_k(\mathbb{F})$$

   (b) Compute the characteristic polynomial $p_W(t) = \det([T_W]_{\beta_W} - tI_k)$ by expanding the determinant along the first row.

2   Inner Product Spaces, part 1

You should be familiar with the *scalar/dot product* in $\mathbb{R}^2$. For any vectors $\mathbf{x} = \left(\begin{smallmatrix} x_1 \\ x_2 \end{smallmatrix}\right)$, $\mathbf{y} = \left(\begin{smallmatrix} y_1 \\ y_2 \end{smallmatrix}\right)$, define
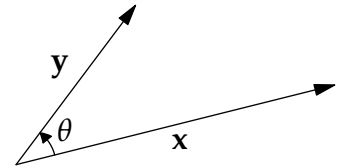
$$\mathbf{x} \cdot \mathbf{y} := x_1 y_1 + x_2 y_2$$

The *norm* or *length* of a vector is $||\mathbf{x}|| = \sqrt{\mathbf{x} \cdot \mathbf{x}}$.

The *angle* θ between vectors satisfies $\mathbf{x} \cdot \mathbf{y} = ||\mathbf{x}|| \, ||\mathbf{y}|| \cos \theta$.

Vectors are *orthogonal* or *perpendicular* precisely when $\mathbf{x} \cdot \mathbf{y} = 0$.

The dot product is what allows us to compute lengths of and angles between vectors in $\mathbb{R}^2$. An *inner product* is an algebraic structure that generalizes this idea to other vector spaces.

2.1   Real and Complex Inner Product Spaces

Unless explicitly stated otherwise, throughout this chapter $\mathbb{F}$ is either the real $\mathbb{R}$ or complex $\mathbb{C}$ field.

Definition 2.1.   An *inner product space* $(V, \langle \, , \, \rangle)$ is a vector space V over $\mathbb{F}$ together with an *inner product*: a function $\langle \, , \, \rangle : V \times V \to \mathbb{F}$ satisfying the following properties $\forall \mathbf{x}, \mathbf{y}, \mathbf{z} \in V$, $\lambda \in \mathbb{F}$,

(a) *Linear*: $\langle \lambda \mathbf{x} + \mathbf{y}, \mathbf{z} \rangle = \lambda \langle \mathbf{x}, \mathbf{z} \rangle + \langle \mathbf{y}, \mathbf{z} \rangle$

(b) *Conjugate-Symmetric*: $\langle \mathbf{y}, \mathbf{x} \rangle = \overline{\langle \mathbf{x}, \mathbf{y} \rangle}$ ⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀⠀(complex conjugate!)

(c) *Positive-definite*: $\mathbf{x} \neq \mathbf{0} \implies \langle \mathbf{x}, \mathbf{x} \rangle > 0$

The *norm* or *length* of $\mathbf{x} \in V$ is $||\mathbf{x}|| := \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$. A *unit vector* has $||\mathbf{x}|| = 1$.

Vectors $\mathbf{x}, \mathbf{y}$ are *perpendicular/orthogonal* if $\langle \mathbf{x}, \mathbf{y} \rangle = 0$ and *orthonormal* if they are additionally unit vectors.

Real inner product spaces

The definition simplifies slightly when $\mathbb{F} = \mathbb{R}$.

- Conjugate-symmetry becomes plain *symmetry*: $\langle \mathbf{y}, \mathbf{x} \rangle = \langle \mathbf{x}, \mathbf{y} \rangle$.

- Linearity $+$ symmetry yields *bilinearity*: a real inner product is also linear in its second slot

$$\langle \mathbf{x}, \lambda \mathbf{y} + \mathbf{z} \rangle = \lambda \langle \mathbf{x}, \mathbf{y} \rangle + \langle \mathbf{x}, \mathbf{z} \rangle$$

A real inner product is often termed a *positive-definite, symmetric, bilinear form.*

The simplest example is the natural generalization of the dot product.

Definition 2.2.   *Euclidean space* means $\mathbb{R}^n$ equipped with the *standard inner (dot) product*,

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x} \cdot \mathbf{y} = \mathbf{y}^T \mathbf{x} = \sum_{j=1}^n x_j y_j = x_1 y_1 + \cdots + x_n y_n, \qquad ||\mathbf{x}|| = \sqrt{\sum_{j=1}^n x_j^2}$$

Unless the inner product is stated explicitly, if we refer to $\mathbb{R}^n$ as an inner product space, we mean Euclidean space. However, there are many other ways to make $\mathbb{R}^n$ into an inner product space...

Example 2.3. Define an alternative inner product on $\mathbb{R}^2$ via

$$\langle \mathbf{x}, \mathbf{y} \rangle = x_1 y_1 + 3 x_2 y_2$$

It is easy to check that this satisfies the required properties:

(a) Linearity: follows from the associative/distributive laws in $\mathbb{R}$,

$$\langle \lambda \mathbf{x} + \mathbf{y}, \mathbf{z} \rangle = (\lambda x_1 + y_1) z_1 + 3(\lambda x_2 + y_2) z_2 = \lambda(x_1 z_1 + 3 x_2 z_2) + (y_1 z_1 + 3 y_2 z_2)$$
$$= \lambda \langle \mathbf{x}, \mathbf{z} \rangle + \langle \mathbf{y}, \mathbf{z} \rangle$$

(b) Symmetry: follows from the commutativity of multiplication in $\mathbb{R}$,

$$\langle \mathbf{y}, \mathbf{x} \rangle = y_1 x_1 + 3 y_2 x_2 = x_1 y_1 + 3 x_2 y_2 = \langle \mathbf{x}, \mathbf{y} \rangle$$

(c) Positive-definiteness: if $\mathbf{x} \neq \mathbf{0}$, then $\langle \mathbf{x}, \mathbf{x} \rangle = x_1^2 + 3 x_2^2 > 0$.

With respect to $\langle \, , \, \rangle$, the concept of *orthogonality* feels strange: e.g., $\{\mathbf{x}, \mathbf{y}\} = \left\{ \frac{1}{2} \left(\begin{smallmatrix} 1 \\ 1 \end{smallmatrix} \right), \frac{1}{2\sqrt{3}} \left(\begin{smallmatrix} 3 \\ -1 \end{smallmatrix} \right) \right\}$ is an orthonormal set!

$$\|\mathbf{x}\|^2 = \frac{1}{4}(1^2 + 3 \cdot 1^2) = 1, \quad \|\mathbf{y}\|^2 = \frac{1}{12}(3^2 + 3 \cdot 1^2) = 1, \quad \langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{4\sqrt{3}}(3 - 3) = 0$$

However, with respect to the standard dot product, these vectors are not special:

$$\mathbf{x} \cdot \mathbf{x} = \frac{1}{2}, \qquad \mathbf{y} \cdot \mathbf{y} = \frac{5}{6}, \qquad \mathbf{x} \cdot \mathbf{y} = \frac{1}{2\sqrt{3}}$$

We have the same *vector space* $\mathbb{R}^2$, but different *inner product spaces*: $(\mathbb{R}^2, \langle \, , \, \rangle) \neq (\mathbb{R}^2, \cdot)$.

The above is an example of a *weighted inner product*: choose *weights* $a_1, \ldots, a_n \in \mathbb{R}^+$ and define

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{j=1}^{n} a_j x_j y_j = a_1 x_1 y_1 + \cdots + a_n x_n y_n$$

It is a simple exercise to check that this defines an inner product on $\mathbb{R}^n$. In particular, $\mathbb{R}^n$ may be equipped with *infinitely many distinct inner products!*

More generally, a symmetric matrix $A \in M_n(\mathbb{R})$ is *positive-definite* if $\mathbf{x}^T A \mathbf{x} > 0$ for all non-zero $\mathbf{x} \in \mathbb{R}^n$. It is straightforward to check that

$$\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{y}^T A \mathbf{x}$$

defines an inner product on $\mathbb{R}^n$. In fact *all* inner products on $\mathbb{R}^n$ arise in this fashion! The weighted inner products correspond to A being diagonal (Euclidean space is $A = I$), but this is not required.

**Example 2.4.** The matrix $A = \left(\begin{smallmatrix} 3 & 1 \\ 1 & 1 \end{smallmatrix} \right)$ is positive-definite and thus defines an inner product

$$\langle \mathbf{x}, \mathbf{y} \rangle = 3 x_1 y_1 + x_1 y_2 + x_2 y_1 + x_2 y_2$$

Lemma 2.5 (Basic properties). *Let V be an inner product space, let $\mathbf{x}, \mathbf{y}, \mathbf{z} \in V$ and $\lambda \in \mathbb{F}$.*

1. $\langle \mathbf{0}, \mathbf{x} \rangle = 0$

2. $||\mathbf{x}|| = 0 \iff \mathbf{x} = \mathbf{0}$

3. $||\lambda \mathbf{x}|| = |\lambda| \, ||\mathbf{x}||$

4. $\langle \mathbf{x}, \mathbf{z} \rangle = \langle \mathbf{y}, \mathbf{z} \rangle$ *for all* $\mathbf{z} \implies \mathbf{x} = \mathbf{y}$

5. *(Cauchy–Schwarz inequality)* $|\langle \mathbf{x}, \mathbf{y} \rangle| \le ||\mathbf{x}|| \, ||\mathbf{y}||$ *with equality if and only if* $\mathbf{x}, \mathbf{y}$ *are parallel*

6. *(Triangle Inequality)* $||\mathbf{x} + \mathbf{y}|| \le ||\mathbf{x}|| + ||\mathbf{y}||$ *with equality if and only if* $\mathbf{x}, \mathbf{y}$ *are parallel and point in the same direction*
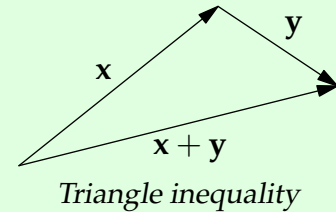
Be careful with notation: $|\lambda|$ means the *absolute value/modulus* (of a *scalar*), while $||\mathbf{x}||$ means the *norm* (of a *vector*).

In the real case, Cauchy–Schwarz allows us to define *angle* via $\cos \theta = \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{||\mathbf{x}|| \, ||\mathbf{y}||}$, since the right-hand side lies in the interval $[-1, 1]$. However, except in Euclidean $\mathbb{R}^2$ and $\mathbb{R}^3$, this notion is of limited use; *orthogonality* ($\langle \mathbf{x}, \mathbf{y} \rangle = 0$) and *orthonormality* are usually all we care about.

Proof. Parts 1–3 are exercises. For simplicity, we prove 5 and 6 only when $\mathbb{F} = \mathbb{R}$.

4. Let $\mathbf{z} = \mathbf{x} - \mathbf{y}$, apply the linearity condition and part 2:

$$\langle \mathbf{x}, \mathbf{z} \rangle = \langle \mathbf{y}, \mathbf{z} \rangle \implies 0 = \langle \mathbf{x} - \mathbf{y}, \mathbf{z} \rangle = \langle \mathbf{x} - \mathbf{y}, \mathbf{x} - \mathbf{y} \rangle = ||\mathbf{x} - \mathbf{y}||^2 \implies \mathbf{x} = \mathbf{y}$$

5. If $\mathbf{y} = \mathbf{0}$, the result is trivial. WLOG (and by part 3) we may assume $||\mathbf{y}|| = 1$; if the inequality holds for this, then it holds for all non-zero $\mathbf{y}$ by parts 1 and 4. Now expand:

$$\begin{aligned}
0 &\le ||\mathbf{x} - \langle \mathbf{x}, \mathbf{y} \rangle \, \mathbf{y}||^2 &&\text{(positive-definiteness)} \\
&= ||\mathbf{x}||^2 + |\langle \mathbf{x}, \mathbf{y} \rangle|^2 \, ||\mathbf{y}||^2 - \langle \mathbf{x}, \mathbf{y} \rangle \langle \mathbf{x}, \mathbf{y} \rangle - \langle \mathbf{x}, \mathbf{y} \rangle \langle \mathbf{y}, \mathbf{x} \rangle &&\text{(bilinearity)} \\
&= ||\mathbf{x}||^2 - |\langle \mathbf{x}, \mathbf{y} \rangle|^2 &&\text{(symmetry)}
\end{aligned}$$

Taking square-roots establishes the inequality.

By part 2, equality holds if and only if $\mathbf{x} = \langle \mathbf{x}, \mathbf{y} \rangle \, \mathbf{y}$, which is precisely when $\mathbf{x}, \mathbf{y}$ are linearly dependent.

6. We establish the squared result.

$$\begin{aligned}
\left(||\mathbf{x}|| + ||\mathbf{y}|| \right)^2 - ||\mathbf{x} + \mathbf{y}||^2 &= ||\mathbf{x}||^2 + 2 \, ||\mathbf{x}|| \, ||\mathbf{y}|| + ||\mathbf{y}||^2 - ||\mathbf{x}||^2 - \langle \mathbf{x}, \mathbf{y} \rangle - \langle \mathbf{y}, \mathbf{x} \rangle - ||\mathbf{y}||^2 \\
&= 2 \left(||\mathbf{x}|| \, ||\mathbf{y}|| - \langle \mathbf{x}, \mathbf{y} \rangle \right) \\
&\ge 2 \left(||\mathbf{x}|| \, ||\mathbf{y}|| - |\langle \mathbf{x}, \mathbf{y} \rangle| \right) \\
&\ge 0 &&\text{(Cauchy–Schwarz)}
\end{aligned}$$

Equality requires both equality in Cauchy–Schwarz ($\mathbf{x}, \mathbf{y}$ parallel) and that $\langle \mathbf{x}, \mathbf{y} \rangle \ge 0$; since $\mathbf{x}, \mathbf{y}$ are already parallel, this means that one is a non-negative multiple of the other. ∎

Complex Inner Product Spaces

Definition 2.1 is already set up nicely when $\mathbb{C} = \mathbb{F}$. One subtle difference comes from how we expand linear combinations in the second slot.

> **Lemma 2.6.** *An inner product is a* positive-definite, conjugate-symmetric, sesquilinear[4]*form: it is* conjugate-linear *(*anti-linear*) in the second slot,*
>
> $$\langle \mathbf{z}, \lambda \mathbf{x} + \mathbf{y} \rangle = \overline{\lambda} \langle \mathbf{z}, \mathbf{x} \rangle + \langle \mathbf{z}, \mathbf{y} \rangle$$

The proof is very easy if you remember your complex conjugates; try it!

Warning! If you dabble in the dark arts of Physics, be aware that their convention[5] is for an inner product to be conjugate-linear in the *first* entry and linear in the second!

> **Definition 2.7.** The *standard (Hermitian) inner product* and *norm* on $\mathbb{C}^n$ are
>
> $$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* \mathbf{x} = \sum_{j=1}^{n} x_j \overline{y_j} = x_1 \overline{y_1} + \cdots + x_n \overline{y_n}, \qquad ||\mathbf{x}|| = \sqrt{\sum_{j=1}^{n} |x_j|^2}$$
>
> where $\mathbf{y}^* = \overline{\mathbf{y}}^T$ is the *conjugate-transpose* of $\mathbf{y}$ and $|x_j|$ is the *modulus*.

Weighted inner products may be defined as on page 16:

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{j=1}^{n} a_j x_j \overline{y_j} = a_1 x_1 \overline{y_1} + \cdots + a_n x_n \overline{y_n}$$

Note that the weights a_j must still be *positive real numbers.* We may similarly define inner products in terms of positive-definite matrices $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* A \mathbf{x}$.

> **Definition 2.8.** A matrix $A \in M_n(\mathbb{C})$ is *Hermitian (self-adjoint)* if $A^* = A$. It is moreover *positive-definite* if $\mathbf{x}^* A \mathbf{x} > 0$ for all non-zero $\mathbf{x} \in \mathbb{C}^n$.

The self-adjoint condition reduces to symmetry ($A^T = A$) if A is a real matrix.

Example 2.9. It can be seen (Exercise 6) that $A = \left(\begin{smallmatrix} 3 & -i \\ i & 3 \end{smallmatrix} \right)$ is positive-definite, whence

$$\langle \mathbf{x}, \mathbf{y} \rangle = \begin{pmatrix} \overline{y_1} & \overline{y_2} \end{pmatrix} \begin{pmatrix} 3 & -i \\ i & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 3x_1 \overline{y_1} + ix_1 \overline{y_2} - ix_2 \overline{y_1} + 3x_2 \overline{y_2}$$

defines an inner product on $\mathbb{C}^2$.

Almost all results in this chapter will be written for general inner product spaces, thus covering the real and complex cases simultaneously. If you don't feel confident with complex numbers, simply let $\mathbb{F} = \mathbb{R}$ and delete all complex conjugates at first read! Very occasionally a different proof will be required depending on the field. For simplicity, examples will more often use real inner products.

---

[4]The prefix *sesqui-* means *one-and-a-half*; for instance a *sesquicentenary* is a 150 year anniversary.
[5]The common Physics notation relates to ours via $\langle \mathbf{x} \,|\, \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$.

Further Examples

As before, the field $\mathbb{F}$ must be either $\mathbb{R}$ or $\mathbb{C}$.

Definition 2.10 (Frobenius inner product).   If $A, B \in M_{m \times n}(\mathbb{F})$, define

$$\langle A, B \rangle = \operatorname{tr}(B^* A)$$

where tr is the *trace* of an $n \times n$ matrix; this makes $M_{m \times n}(\mathbb{F})$ into an inner product space.

This isn't really a new example: if we map $A \to \mathbb{F}^{mn}$ by stacking the columns of A, then the Frobenius inner product is the standard inner product in disguise.

Example 2.11.   In $M_{3 \times 2}(\mathbb{C})$,

$$\left\langle \begin{pmatrix} 1 & i \\ 2-i & 0 \\ 0 & -1 \end{pmatrix}, \begin{pmatrix} 0 & 7 \\ 1 & 2i \\ 3-2i & 4 \end{pmatrix} \right\rangle = \operatorname{tr}\begin{pmatrix} 0 & 1 & 3+2i \\ 7 & -2i & 4 \end{pmatrix} \begin{pmatrix} 1 & i \\ 2-i & 0 \\ 0 & -1 \end{pmatrix} = \operatorname{tr}\begin{pmatrix} 2-i & -3-2i \\ 9-4i & -4+7i \end{pmatrix}$$

$$= 2 - i - 4 + 7i = -2 - 8i$$

$$\left\| \begin{pmatrix} 1 & i \\ 2-i & 0 \\ 0 & -1 \end{pmatrix} \right\|^2 = \operatorname{tr}\begin{pmatrix} 1 & 2+i & 0 \\ -i & 0 & -1 \end{pmatrix} \begin{pmatrix} 1 & i \\ 2-i & 0 \\ 0 & -1 \end{pmatrix} = \operatorname{tr}\begin{pmatrix} 6 & i \\ -i & 2 \end{pmatrix} = 8$$

Definition 2.12 (L^2 inner product).   Given a *real* interval $[a, b]$, the function

$$\langle f, g \rangle := \int_a^b f(t)\overline{g(t)} \, \mathrm{d}t$$

defines an inner product on the space $C[a, b]$ of continuous functions $f : [a, b] \to \mathbb{F}$.

With careful restriction, this works even for infinite intervals and a larger class of functions.[6] Verifying the required properties is straightforward if you know a little analysis; for instance continuity allows us to conclude

$$\|f\|^2 = \int_a^b |f(x)|^2 \, \mathrm{d}x = 0 \iff f(x) \equiv 0$$

This is our first example of an *infinite-dimensional* inner product space.

Example 2.13.   Let $f(x) = x$ and $g(x) = x^2$; these lie in the inner product space $C[-1, 1]$ with respect to the L^2 inner product.

$$\langle f, g \rangle = \int_{-1}^1 x^3 \, \mathrm{d}x = 0, \qquad \|f\|^2 = \int_{-1}^1 x^2 \, \mathrm{d}x = \frac{2}{3}, \qquad \|g\|^2 = \int_{-1}^1 x^4 \, \mathrm{d}x = \frac{2}{5}$$

With some simple scaling, we see that $\left\{ \frac{1}{\|f\|} f, \frac{1}{\|g\|} g \right\} = \left\{ \sqrt{\frac{3}{2}} x, \sqrt{\frac{5}{2}} x^2 \right\}$ forms an orthonormal set.

---

[6]For us, functions will always be continuous (often polynomials) on closed bounded intervals. The *square-integrable functions* and *L^2-spaces* for which the inner product is named are a more complicated business and beyond this course.

Definition 2.14 (ℓ^2 inner product). A sequence (x_n) is *square-summable* if $\sum\limits_{n=1}^{\infty} |x_n|^2 < \infty$. These sequences form a vector space on which we can define an inner product[7]

$$\langle (x_n), (y_n) \rangle = \sum_{n=1}^{\infty} x_n \overline{y_n}$$

In essence we've taken the standard inner product on $\mathbb{F}^n$ and let $n \to \infty$! This example, and its L^2 cousin, are the prototypical *Hilbert spaces,* which have great application to differential equations, signal processing, etc. Since a rigorous discussion requires a significant amount of analysis (convergence of series, completeness, integrability), these objects are generally beyond the course.

Our final example of an inner product is a useful, and hopefully obvious, hack to which we shall repeatedly appeal in examples.

Lemma 2.15. *Let V be a vector space over $\mathbb{R}$ or $\mathbb{C}$. If β is a basis of V, then there exists exactly one inner product on V for which β is an orthonormal set.*

Exercises 2.1 1. Evaluate the inner product of the given vectors.

   (a) $\mathbf{x} = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}$, $\mathbf{y} = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix}$ where $\langle \mathbf{x}, \mathbf{y} \rangle = 2x_1 y_1 + 3x_2 y_2 + x_3 y_3$

   (b) $\mathbf{x} = \begin{pmatrix} 1 \\ 2i \end{pmatrix}$, $\mathbf{y} = \begin{pmatrix} 5i \\ 4 \end{pmatrix}$ where $\langle \mathbf{x}, \mathbf{y} \rangle$ is the standard Hermitian inner product on $\mathbb{C}^2$

   (c) $\mathbf{x} = \begin{pmatrix} 1 \\ 2i \end{pmatrix}$, $\mathbf{y} = \begin{pmatrix} 5i \\ 4 \end{pmatrix}$ where $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* \begin{pmatrix} 2 & -i \\ i & 2 \end{pmatrix} \mathbf{x}$

   (d) $f(x) = x - 1$, $g(x) = x + 1$ where $\langle f, g \rangle$ is the L^2 inner product on $C[0,2]$

2. Suppose $\langle \mathbf{x}, \mathbf{y} \rangle = \langle \mathbf{x}, \mathbf{z} \rangle$ for all $\mathbf{x}$. Prove that $\mathbf{y} = \mathbf{z}$.

3. For each $\mathbf{z} \in V$, The linearity condition says that the map $T_{\mathbf{z}} : V \to \mathbb{F}$ defined by

   $$T_{\mathbf{z}}(\mathbf{x}) = \langle \mathbf{x}, \mathbf{z} \rangle$$

   is linear ($T_{\mathbf{z}}$ is an element of the *dual space* V^*). What, if anything, can you say about the function $U_{\mathbf{z}} : \mathbf{x} \mapsto \langle \mathbf{z}, \mathbf{x} \rangle$?

4. Define $\langle \mathbf{x}, \mathbf{y} \rangle := \sum\limits_{j=1}^{n} x_j y_j$ on $\mathbb{C}^n$. Is this an inner product? Which of the properties (a), (b), (c) from Definition 2.1 does it satisfy?

5. (a) Verify that the matrix in Example 2.4 is positive-definite and that $\langle \mathbf{x}, \mathbf{y} \rangle = 3x_1 y_1 + x_1 y_2 + x_2 y_1 + x_2 y_2$ therefore defines an inner product on $\mathbb{R}^2$.
   (*Hint: Try to write* $||\mathbf{x}||^2$ *as a sum of squares...*)

   (b) Let $\mathbf{x} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. With respect to the inner product in part (a), find a non-zero unit vector $\mathbf{y}$ which is *orthogonal* to $\mathbf{x}$.

6. By multiplying out $|x_1 - ix_2|^2$, show that the matrix $\begin{pmatrix} 3 & -i \\ i & 3 \end{pmatrix}$ is positive-definite.

   (*Hint: recall that* $|a|^2 = a\bar{a}$ *for complex numbers!!!*)

---

[7]Neither of these facts are obvious; in particular, we'd need to see that the sum of two square-summable sequences is also square-summable.

7. Show that every eigenvalue of a positive definite matrix is positive.

8. Prove parts 1, 2 and 3 of Lemma 2.5.

9. Let V be an inner product space, prove Pythagoras' Theorem:

$$\text{If } \langle \mathbf{x}, \mathbf{y} \rangle = 0, \text{ then } ||\mathbf{x} + \mathbf{y}||^2 = ||\mathbf{x}||^2 + ||\mathbf{y}||^2$$

10. Use basic algebra to prove the Cauchy–Schwarz inequality for vectors $\mathbf{x} = \left(\begin{smallmatrix} a \\ b \end{smallmatrix} \right)$ and $\mathbf{y} = \left(\begin{smallmatrix} c \\ d \end{smallmatrix} \right)$ in $\mathbb{R}^2$ with the standard (dot) product.

11. Prove the Cauchy–Schwarz and triangle inequalities for a *complex* inner product space. What has to change compared to the proof of Lemma 2.5?

12. Prove the polarization identities:

    (a) In any real inner product space: $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{4} ||\mathbf{x} + \mathbf{y}||^2 - \frac{1}{4} ||\mathbf{x} - \mathbf{y}||^2$

    (b) In any complex inner product space: $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{4} \sum_{k=1}^{4} i^k ||\mathbf{x} + i^k \mathbf{y}||^2$

    If you know the length of every vector, then you know the inner product!

13. Prove that $\int_0^2 \frac{\sqrt{x}}{x+1} \, dx \leq \frac{2}{\sqrt{3}}$

    (*Hint: use Cauchy–Schwarz*)

14. Let $m \in \mathbb{Z}$ and consider the complex-valued function $f_m(x) = \frac{1}{\sqrt{2\pi}} e^{imx}$. If $\langle \, , \, \rangle$ is the L^2 inner product on $C[-\pi, \pi]$, prove that $\{f_m : m \in \mathbb{Z}\}$ is an orthonormal set.

    This example is central to the study of Fourier series.

    (*Hint: If complex functions are scary, use Euler's formula $e^{imx} = \cos mx + i \sin mx$ and work with the real-valued functions $\cos mx$ and $\sin mx$. The difficulty is that you then need integration by parts...*)

15. Let $\langle \, , \, \rangle$ be an inner product on $\mathbb{F}^n$ (recall that $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$). Define the matrix $A \in M_n(\mathbb{F})$ by $A_{jk} = \langle \mathbf{e}_k, \mathbf{e}_j \rangle$ where $\{\mathbf{e}_1, \ldots, \mathbf{e}_n\}$ is the standard basis. Verify that A is the *matrix of the inner product*:

$$\forall \mathbf{x}, \mathbf{y} \in \mathbb{F}^n, \ \langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* A \mathbf{x}$$

    In particular,

    - A is a *Hermitian/self-adjoint* matrix: $A^* = A$ (if $\mathbb{F} = \mathbb{R}$ this is simply *symmetric*);
    - A is *positive-definite*: for all $\mathbf{x} \in \mathbb{F}^n$, $\mathbf{x}^* A \mathbf{x} > 0$;

    More generally, if $\beta = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ is a basis then $A_{jk} = \langle \mathbf{v}_k, \mathbf{v}_j \rangle$ defines the matrix of the inner product with respect to β:

$$\langle \mathbf{x}, \mathbf{y} \rangle = [\mathbf{y}]_\beta^* A [\mathbf{x}]_\beta$$

2.2 Orthogonal Sets and the Gram–Schmidt Process

We start with a simple definition, relating subspaces to orthogonality.

Definition 2.16. Let U be a subspace of an inner product space V. The *orthogonal complement* $U^\perp$ is the set

$$U^\perp = \{\mathbf{x} \in V : \forall \mathbf{u} \in U, \ \langle \mathbf{x}, \mathbf{u} \rangle = 0\}$$

It is easy to check that $U^\perp$ is itself a subspace of V and that $U \cap U^\perp = \{\mathbf{0}\}$. It can moreover be seen that $U \subseteq (U^\perp)^\perp$, though equality need not hold in infinite dimensions (see Exercise 7).

Example 2.17. $U = \text{Span}\left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \right\} \leq \mathbb{R}^3$ has orthogonal complement $U^\perp = \text{Span}\left\{ \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix} \right\}$.

As in the example, we often have a direct sum decomposition $V = U \oplus U^\perp$: otherwise said,

$$\forall \mathbf{x} \in V, \exists \text{ unique } \mathbf{u} \in U, \mathbf{w} \in U^\perp \text{ such that } \mathbf{x} = \mathbf{u} + \mathbf{w}$$

In such a case, we'll call $\mathbf{u} = \pi_U(\mathbf{x})$ and $\mathbf{w} = \pi_U^\perp(\mathbf{x})$ the *orthogonal projections* of $\mathbf{x}$ onto U and $U^\perp$ respectively. As our first result shows, these are easy to compute whenever U has a finite orthogonal basis.
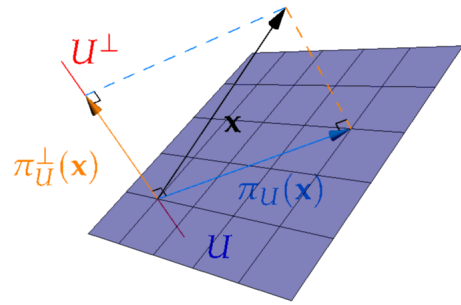
Theorem 2.18. *Let V be an inner product space and let $U = \text{Span}\, \beta$ where $\beta = \{\mathbf{u}_1, \ldots, \mathbf{u}_n\}$ is an orthogonal set of non-zero vectors;*

$$\langle \mathbf{u}_j, \mathbf{u}_k \rangle = \begin{cases} 0 & \text{if } j \neq k \\ ||\mathbf{u}_j||^2 \neq 0 & \text{if } j = k \end{cases}$$

Then:

1. *β is a basis of U and each $\mathbf{x} \in U$ has unique representation*

$$\mathbf{x} = \sum_{j=1}^n \frac{\langle \mathbf{x}, \mathbf{u}_j \rangle}{||\mathbf{u}_j||^2} \mathbf{u}_j \qquad (*)$$

   This simplifies to $\mathbf{x} = \sum \langle \mathbf{x}, \mathbf{u}_j \rangle \mathbf{u}_j$ if β is an orthonormal set.

2. *$V = U \oplus U^\perp$. For any $\mathbf{x} \in V$, we may write $\mathbf{x} = \mathbf{u} + \mathbf{w}$ where*

$$\mathbf{u} = \sum_{j=1}^n \frac{\langle \mathbf{x}, \mathbf{u}_j \rangle}{||\mathbf{u}_j||^2} \mathbf{u}_j \in U \quad and \quad \mathbf{w} = \mathbf{x} - \mathbf{u} \in U^\perp$$

Observe that $(*)$ essentially calculates the co-ordinate vector $[\mathbf{x}]_\beta \in \mathbb{F}^n$. Recalling how unpleasant such calculations have been in the past, often requiring large matrix inversions, we immediately see the power of inner products and orthogonal bases.

Proof. 1. Since β spans U, a given $\mathbf{x} \in U$ may be written

$$\mathbf{x} = \sum_{k=1}^{n} a_k \mathbf{u}_k$$

for some scalars a_k. The orthogonality of β recovers the required expression for a_j:

$$\langle \mathbf{x}, \mathbf{u}_j \rangle = \sum_{k=1}^{n} a_k \langle \mathbf{u}_k, \mathbf{u}_j \rangle = a_j \|\mathbf{u}_j\|^2$$

Finally, let $\mathbf{x} = \mathbf{0}$ to see that β is linearly independent.

2. Clearly $\mathbf{u} \in U$. For each $\mathbf{u}_k$, the orthogonality of β tells us that

$$\langle \mathbf{w}, \mathbf{u}_k \rangle = \langle \mathbf{x} - \mathbf{u}, \mathbf{u}_k \rangle = \langle \mathbf{x}, \mathbf{u}_k \rangle - \sum_{j=1}^{n} \frac{\langle \mathbf{x}, \mathbf{u}_j \rangle}{\|\mathbf{u}_j\|^2} \langle \mathbf{u}_j, \mathbf{u}_k \rangle = \langle \mathbf{x}, \mathbf{u}_k \rangle - \langle \mathbf{x}, \mathbf{u}_k \rangle = 0$$

Since $\mathbf{w}$ is orthogonal to a basis of U it is orthogonal to any element of U; we conclude that $\mathbf{w} \in U^\perp$. Finally, $U \cap U^\perp = \{\mathbf{0}\}$ forces the uniqueness of the decomposition $\mathbf{x} = \mathbf{u} + \mathbf{w}$, whence $V = U \oplus U^\perp$. ∎

Examples 2.19. 1. Consider the standard orthonormal basis $\beta = \{\mathbf{e}_1, \mathbf{e}_2\}$ of $\mathbb{R}^2$. For any $\mathbf{x} = \left(\begin{smallmatrix} x_1 \\ x_2 \end{smallmatrix} \right)$, we easily check that

$$\sum_{j=1}^{2} \langle \mathbf{x}, \mathbf{e}_j \rangle \mathbf{e}_j = x_1 \mathbf{e}_1 + x_2 \mathbf{e}_2 = \mathbf{x}$$

2. In $\mathbb{R}^3$, $\beta = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\} = \left\{ \left(\begin{smallmatrix} 1 \\ 2 \\ 3 \end{smallmatrix} \right), \left(\begin{smallmatrix} 2 \\ -1 \\ 0 \end{smallmatrix} \right), \left(\begin{smallmatrix} 3 \\ 6 \\ -5 \end{smallmatrix} \right) \right\}$ is an orthogonal set and thus a basis. We compute the co-ordinates of $\mathbf{x} = \left(\begin{smallmatrix} 7 \\ 4 \\ 2 \end{smallmatrix} \right)$ with respect to β:

$$\mathbf{x} = \sum_{j=1}^{3} \frac{\langle \mathbf{x}, \mathbf{u}_j \rangle}{\|\mathbf{u}_j\|^2} \mathbf{u}_j = \frac{7+8+6}{1+4+9} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + \frac{14-4+0}{4+1+0} \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix} + \frac{21+24-10}{9+36+25} \begin{pmatrix} 3 \\ 6 \\ -5 \end{pmatrix}$$

$$= \frac{3}{2} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + 2 \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 3 \\ 6 \\ -5 \end{pmatrix} \implies [\mathbf{x}]_\beta = \begin{pmatrix} 3/2 \\ 2 \\ 1/2 \end{pmatrix}$$

Compare this with the painfully slow augmented matrix method for finding co-ordinates!

3. Revisiting Example 2.17, let $\mathbf{x} = \left(\begin{smallmatrix} 1 \\ 1 \\ 1 \end{smallmatrix} \right)$. Since $\beta = \left\{ \left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix} \right), \left(\begin{smallmatrix} 0 \\ -1 \\ 3 \end{smallmatrix} \right) \right\}$ is an orthogonal basis of U, we observe that

$$\pi_U(\mathbf{x}) = \frac{1}{1} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} + \frac{2}{10} \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} = \frac{1}{5} \begin{pmatrix} 5 \\ -1 \\ 3 \end{pmatrix}, \qquad \pi_U^\perp(\mathbf{x}) = \mathbf{x} - \pi_U(\mathbf{x}) = \frac{2}{5} \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix}$$

are the orthogonal projections corresponding to $\mathbb{R}^3 = U \oplus U^\perp$.

The Gram–Schmidt Process

Theorem 2.18 tells us how to compute the orthogonal projections corresponding to $V = U \oplus U^\perp$, *provided U has a finite, orthogonal basis.* Given how useful such bases are, our next goal is to see that such exist for any finite-dimensional subspace. Helpfully there exists a constructive algorithm.

> **Theorem 2.20 (Gram–Schmidt).** *Suppose $S = \{\mathbf{s}_1, \ldots, \mathbf{s}_n\}$ is a linearly independent subset of an inner product space V. Construct a sequence of vectors $\mathbf{u}_1, \ldots, \mathbf{u}_n$ inductively:*
>
> - *Choose $\mathbf{u}_1 = a_1 \mathbf{s}_1$ where $a_1 \neq 0$*
>
> - *For each $k \geq 2$, choose $\mathbf{u}_k = a_k \left(\mathbf{s}_k - \sum_{j=1}^{k-1} \dfrac{\langle \mathbf{s}_k, \mathbf{u}_j \rangle}{||\mathbf{u}_j||^2} \mathbf{u}_j \right)$ where $a_k \neq 0$* $\hspace{1cm}$ (†)
>
> *Then $\beta := \{\mathbf{u}_1, \ldots, \mathbf{u}_n\}$ is an orthogonal basis of $\operatorname{Span} S$.*

The purpose of the scalars a_k is to give you some freedom; choose them to avoid unpleasant fractions! If you want a set of *orthonormal* vectors, it is easier to scale everything after the algorithm is complete. Indeed, by taking S to be a basis of V and normalizing the resulting β, we conclude:

> **Corollary 2.21.** *Every finite-dimensional inner product space has an orthonormal basis.*

Example 2.22. $S = \{\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3\} = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}$ is a linearly independent subset of $\mathbb{R}^3$.

1. Choose $\mathbf{u}_1 = \mathbf{s}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \implies ||\mathbf{u}_1||^2 = 1$

2. $\mathbf{s}_2 - \dfrac{\langle \mathbf{s}_2, \mathbf{u}_1 \rangle}{||\mathbf{u}_1||^2} \mathbf{u}_1 = \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix} - \frac{2}{1} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix}$: choose $\mathbf{u}_2 = \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \implies ||\mathbf{u}_2||^2 = 10$

3. $\mathbf{s}_3 - \dfrac{\langle \mathbf{s}_3, \mathbf{u}_1 \rangle}{||\mathbf{u}_1||^2} \mathbf{u}_1 - \dfrac{\langle \mathbf{s}_3, \mathbf{u}_2 \rangle}{||\mathbf{u}_2||^2} \mathbf{u}_2 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} - \frac{1}{1} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} - \frac{2}{10} \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} = \frac{2}{5} \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix}$: choose $\mathbf{u}_3 = \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix}$

The orthogonality of $\beta = \{\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3\}$ is clear. It is now trivial to observe that $\left\{ \mathbf{u}_1, \frac{1}{\sqrt{10}} \mathbf{u}_2, \frac{1}{\sqrt{10}} \mathbf{u}_3 \right\}$ is an orthonormal basis of $\mathbb{R}^3$.

Proof of Theorem 2.20. For each $k \leq n$, define $S_k = \{\mathbf{s}_1, \ldots, \mathbf{s}_k\}$ and $\beta_k = \{\mathbf{u}_1, \ldots, \mathbf{u}_k\}$. We prove by induction that each β_k is an orthogonal set of non-zero vectors and that $\operatorname{Span} \beta_k = \operatorname{Span} S_k$. The Theorem is then the terminal case $k = n$.

(Base case $k = 1$) Certainly $\beta_1 = \{a_1 \mathbf{s}_1\}$ is an orthogonal set and $\operatorname{Span} \beta_1 = \operatorname{Span} S_1$.

(Induction step) Fix $k \geq 2$, assume β_{k-1} is an orthogonal non-zero set and that $\operatorname{Span} \beta_{k-1} = \operatorname{Span} S_{k-1}$. By Theorem 2.18, $\mathbf{u}_k \in (\operatorname{Span} \beta_{k-1})^\perp$. We also see that $\mathbf{u}_k \neq \mathbf{0}$, for if not,

$$(†) \implies \mathbf{s}_k \in \operatorname{Span} \beta_{k-1} = \operatorname{Span} S_{k-1}$$

and S would be linearly dependent. It follows that β_k is an orthogonal set of non-zero vectors. Moreover, $\mathbf{s}_k \in \operatorname{Span} \beta_k \implies \operatorname{Span} S_k \leq \operatorname{Span} \beta_k$. Since these spaces have the same (finite) dimension k, we conclude that $\operatorname{Span} \beta_k = \operatorname{Span} S_k$.

By induction, β is an orthogonal, non-zero spanning set for $\operatorname{Span} S$; by Theorem 2.18, it is a basis. $\blacksquare$

Example 2.23. This time we work in the space of polynomials $P(\mathbb{R})$ equipped with the L^2 inner product $\langle f, g \rangle = \int_0^1 f(x)g(x)\,\mathrm{d}x$ on the interval $[0, 1]$. Let $S = \{1, x, x^2\}$ and apply the algorithm:

1. Choose $f_1(x) = 1 \implies ||f_1||^2 = \int_0^1 1\,\mathrm{d}x = 1$

2. $x - \frac{\langle x, f_1 \rangle}{||f_1||^2} f_1 = x - \int_0^1 x\,\mathrm{d}x = x - \frac{1}{2}$

   We choose $f_2(x) = 2x - 1$, with $||f_2||^2 = \int_0^1 (2x - 1)^2\,\mathrm{d}x = \frac{1}{3}$

3. $x^2 - \frac{\langle x^2, f_1 \rangle}{||f_1||^2} f_1 - \frac{\langle x^2, f_2 \rangle}{||f_2||^2} = x^2 - \int_0^1 x^2\,\mathrm{d}x - \frac{\int_0^1 x^2(2x-1)\,\mathrm{d}x}{1/3}(2x - 1) = x^2 - x + \frac{1}{6}$

   We choose $f_3(x) = 6x^2 - 6x + 1$ with $||f_3||^2 = \int_0^1 \left(6x^2 - 6x + 1\right)^2\,\mathrm{d}x = \frac{1}{5}$

It follows that Span S has an orthonormal basis

$$\beta = \left\{ 1, \sqrt{3}(2x - 1), \sqrt{5}(6x^2 - 6x + 1) \right\}$$

This example can be extended to arbitrary degree since the countable set $\{1, x, x^2, \ldots\}$ is basis of $P(\mathbb{R})$. Indeed this shows that $(P(\mathbb{R}), \langle\,,\,\rangle)$ has an orthonormal basis.

Gram–Schmidt also shows that our earlier discussion of orthogonal projections is generic.

> **Corollary 2.24.** *If U is a finite-dimensional subspace of an inner product space V, then $V = U \oplus U^\perp$ and the orthogonal projections may be computed as in Theorem 2.18.*

If U is an infinite-dimensional subspace of V, then we need not have $V = U \oplus U^\perp$ and the orthogonal projections might not be well-defined (see, for example, Exercises 7 and 8). Instead, if β is an orthonormal basis of U, it is common to describe the coefficients $\langle \mathbf{x}, \mathbf{u} \rangle$ for each $\mathbf{u} \in \beta$ as the *Fourier coefficients*, and the infinite sum

$$\sum_{\mathbf{u} \in \beta} \langle \mathbf{x}, \mathbf{u} \rangle \mathbf{u}$$

as the *Fourier series* of $\mathbf{x}$, provided the sum converges.

Exercises 2.2 1. Apply Gram–Schmidt to obtain an orthogonal basis β for Span S. Then obtain the co-ordinate representation (Fourier coefficients) of the given vector with respect to β.

(a) $S = \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\} \subseteq \mathbb{R}^3$, $\mathbf{x} = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$

(b) $S = \left\{ \begin{pmatrix} 3 & 5 \\ -1 & 1 \end{pmatrix}, \begin{pmatrix} -1 & 9 \\ 5 & -1 \end{pmatrix}, \begin{pmatrix} 7 & -17 \\ 2 & -6 \end{pmatrix} \right\} \subseteq M_2(\mathbb{R})$, $X = \begin{pmatrix} -1 & 27 \\ -4 & 8 \end{pmatrix}$ (use the Frobenius product)

(c) $S = \left\{ \begin{pmatrix} 1 \\ i \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ i \end{pmatrix} \right\} \subseteq \mathbb{C}^3$ with $\mathbf{x} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$

(d) $S = \{1, x, x^2\}$ with $\langle f, g \rangle = \int_{-1}^1 f(x)g(x)\,\mathrm{d}x$ and $f(x) = x^2$.

Important! You'll likely need *much* more practice than this to get comfortable with Gram–Schmidt; make up your own problems!

25

2. Let $S = \{\mathbf{s}_1, \mathbf{s}_2\} = \left\{ \begin{pmatrix} 1 \\ 0 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix} \right\}$ and $U = \mathrm{Span}\, S \leq \mathbb{R}^3$. Find $\pi_U(\mathbf{x})$ if $\mathbf{x} = \begin{pmatrix} 3 \\ -1 \\ -2 \end{pmatrix}$.

   (*Hint: First apply Gram–Schmidt*)

3. Find the orthogonal complement to $U = \mathrm{Span}\{x^2\} \leq P_2(\mathbb{R})$ with respect to the inner product
   $\langle f, g \rangle = \int_0^1 f(t)g(t)\, \mathrm{d}t$

4. Let $\mathrm{T} \in \mathcal{L}(V, W)$ where V, W are inner product spaces with orthonormal bases $\beta = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ and $\gamma = \{\mathbf{w}_1, \ldots, \mathbf{w}_m\}$ respectively. Prove that the matrix $A = [\mathrm{T}]_\beta^\gamma \in M_{m \times n}(\mathbb{F})$ of T with respect to these bases has jk^{th} entry

   $$A_{jk} = \langle \mathrm{T}(\mathbf{v}_k), \mathbf{w}_j \rangle$$

5. Suppose that β is an orthonormal basis of an n-dimensional inner product space V. Prove that,

   $$\forall \mathbf{x}, \mathbf{y} \in V, \ \langle \mathbf{x}, \mathbf{y} \rangle = [\mathbf{y}]_\beta^* [\mathbf{x}]_\beta$$

   Otherwise said, the co-ordinate isomorphism $\phi_\beta : V \to \mathbb{F}^n$ defined by $\phi_\beta(\mathbf{x}) = [\mathbf{x}]_\beta$ is an isomorphism of inner product spaces where we use the standard inner product on $\mathbb{F}^n$

6. Let U be a subspace of an inner product space V. Prove the following:

   (a) $U^\perp$ is a subspace of V.
   (b) $U \cap U^\perp = \{\mathbf{0}\}$
   (c) $U \subseteq (U^\perp)^\perp$
   (d) If $V = U \oplus U^\perp$, then $U = (U^\perp)^\perp$ (this is always the case when $\dim U < \infty$)

7. Let ℓ^2 be the set of square-summable sequences of real numbers (Definition 2.14). Consider the sequences $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \ldots$, where $\mathbf{u}_j$ is the zero sequence except for a single 1 in the j^{th} entry. For instance,

   $$\mathbf{u}_4 = (0, 0, 0, 1, 0, 0, 0, 0, \ldots)$$

   (a) Let $U = \mathrm{Span}\{\mathbf{u}_j : j \in \mathbb{N}\}$. Prove that $U^\perp$ contains only the zero sequence.
   (b) Show that the sequence $\mathbf{y} = \left(\frac{1}{n}\right)$ lies in ℓ^2, but does not lie in U.
   U is therefore a proper subset of $(U^\perp)^\perp = \ell^2$ and $\ell^2 \neq U \oplus U^\perp$.

8. Recall Exercise 2.1.14 where we saw that the set $\beta = \{\frac{1}{\sqrt{2\pi}} e^{imx} : m \in \mathbb{Z}\}$ is orthonormal with respect to $\langle f, g \rangle = \int_{-\pi}^\pi f(t)\overline{g(t)}\, \mathrm{d}t$.

   (a) Show that the *Fourier series* of $f(x) = x$ is

   $$\mathcal{F}(x) := \sum_{m=-\infty}^\infty \left\langle x, \frac{1}{\sqrt{2\pi}} e^{imx} \right\rangle \frac{1}{\sqrt{2\pi}} e^{imx} = \sum_{m=1}^\infty \frac{2(-1)^{m+1}}{m} \sin mx$$

   (b) Briefly explain why the Fourier series is *not* an element of $\mathrm{Span}\,\beta$.
   (c) Sketch a few of the Fourier approximations (sum up to $m = 5$ or $7\ldots$) and observe, when extended to $\mathbb{R}$, how they approximate a discontinuous periodic function.

9. (Hard)   Much of Theorem 2.18 remains true, with suitable modifications, even if β is an *infinite* set. Restate and prove as much as you can, and identify the false part(s).

2.3   The Adjoint of a Linear Operator

Recall how the standard inner product on $\mathbb{F}^n$ may be written in terms of the conjugate-transpose

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* \mathbf{x} = \overline{\mathbf{y}}^T \mathbf{x}$$

We start by inserting a matrix into this expression and interpreting in two different ways. Suppose $A \in M_{m \times n}(\mathbb{F})$, $\mathbf{v} \in \mathbb{F}^n$ and $\mathbf{w} \in \mathbb{F}^m$, then

$$\underbrace{\langle A^* \mathbf{w}, \mathbf{v} \rangle}_{\text{in } \mathbb{F}^n} = \mathbf{v}^*(A^*\mathbf{w}) = (\mathbf{v}^* A^*)\mathbf{w} = (A\mathbf{v})^*\mathbf{w} = \underbrace{\langle \mathbf{w}, A\mathbf{v} \rangle}_{\text{in } \mathbb{F}^m} \tag{†}$$

Example 2.25.   As a sanity check, let $A = \left(\begin{smallmatrix} 1 & 2 \\ 0 & 3 \end{smallmatrix} \right) \in M_2(\mathbb{R})$, $\mathbf{w} = \left(\begin{smallmatrix} x \\ y \end{smallmatrix} \right)$ and $\mathbf{v} = \left(\begin{smallmatrix} p \\ q \end{smallmatrix} \right)$. Then,

$$\left\langle A^T \begin{pmatrix} x \\ y \end{pmatrix}, \begin{pmatrix} p \\ q \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} 1 & 0 \\ 2 & 3 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}, \begin{pmatrix} p \\ q \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} x \\ 2x + 3y \end{pmatrix}, \begin{pmatrix} p \\ q \end{pmatrix} \right\rangle = xp + (2x + 3y)q$$

$$\left\langle \begin{pmatrix} x \\ y \end{pmatrix}, A \begin{pmatrix} p \\ q \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} x \\ y \end{pmatrix}, \begin{pmatrix} 1 & 2 \\ 0 & 3 \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix} \right\rangle = \left\langle \begin{pmatrix} x \\ y \end{pmatrix}, \begin{pmatrix} p + 2q \\ 3q \end{pmatrix} \right\rangle = x(p + 2q) + 3yq$$

Note how the inner products are evaluated on *different spaces.* At the level of linear maps this is a relationship between $L_A \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$ and $L_{A^*} \in \mathcal{L}(\mathbb{F}^m, \mathbb{F}^n)$, one that is easily generalizable.

> **Definition 2.26.**   Let $T \in \mathcal{L}(V, W)$ where V, W are inner product spaces over the same field $\mathbb{F}$. The *adjoint* of T is a function $T^* : W \to V$ (read 'T-star') satisfying
>
> $$\forall \mathbf{v} \in V, \mathbf{w} \in W, \quad \langle T^*(\mathbf{w}), \mathbf{v} \rangle = \langle \mathbf{w}, T(\mathbf{v}) \rangle$$
>
> Note that the first inner product is computed within V and the second within W.

The adjoint effectively extends the conjugate-transpose to linear maps. We now use the same notation for three objects, so be careful!

- If A is a real or complex matrix, then $A^* = \overline{A}^T$ is its *conjugate-transpose.*

- If T is a linear map, then T^* is its *adjoint.*

- If V is a vector space, then $V^* = \mathcal{L}(V, \mathbb{F})$ is its *dual space.*

Thankfully the two notations line up nicely, as part 3 of our first result shows.

> **Theorem 2.27 (Basic Properties).**   1. *If an adjoint exists,[8] then it is unique and linear.*
>
> 2. *If* T *and* S *have adjoints, then*
>
> $$(T^*)^* = T, \quad (TS)^* = S^*T^*, \quad (\lambda T + S)^* = \overline{\lambda} T^* + S^*$$
>
> 3. *Suppose* V, W *are finite-dimensional with orthonormal bases* β, γ *respectively. Then the matrix of the adjoint of* $T \in \mathcal{L}(V, W)$ *is the conjugate-transpose of the original:* $[T^*]_\gamma^\beta = ([T]_\beta^\gamma)^*$.

---

[8]Existence of adjoints is trickier, so we postpone this a little: see Corollary 2.34 and Exercise 12.

Proof.   1. (*Uniqueness*)   Suppose T^* and S^* are adjoints of T. Then

$$\langle T^*(\mathbf{x}), \mathbf{y} \rangle = \langle \mathbf{x}, T(\mathbf{y}) \rangle = \langle S^*(\mathbf{x}), \mathbf{y} \rangle$$

Since this holds for all $\mathbf{y}$, Lemma 2.5 part 4 says that $\forall \mathbf{x}$, $T^*(\mathbf{x}) = S^*(\mathbf{x})$, whence $T^* = S^*$.

(*Linearity*)   Simply translate across, use the linearity of T, and again appeal to Lemma 2.5:

$$\forall \mathbf{z}, \ \langle T^*(\lambda\mathbf{x} + \mathbf{y}), \mathbf{z} \rangle = \langle \lambda\mathbf{x} + \mathbf{y}, T(\mathbf{z}) \rangle = \lambda \langle \mathbf{x}, T(\mathbf{z}) \rangle + \langle \mathbf{y}, T(\mathbf{z}) \rangle = \lambda \langle T^*(\mathbf{x}), \mathbf{z} \rangle + \langle T^*(\mathbf{y}), \mathbf{z} \rangle$$
$$= \langle \lambda T^*(\mathbf{x}) + T^*(\mathbf{y}), \mathbf{z} \rangle$$
$$\implies T^*(\lambda\mathbf{x} + \mathbf{y}) = \lambda T^*(\mathbf{x}) + T^*(\mathbf{y})$$

2. These may be proved similarly to part 1 and are left as an exercise.

3. By Exercise 2.2.4, the jk^{th} entry of $[T^*]_\gamma^\beta$ is

$$\langle T^*(\mathbf{w}_k), \mathbf{v}_j \rangle = \langle \mathbf{w}_k, T(\mathbf{v}_j) \rangle = \overline{\langle T(\mathbf{v}_j), \mathbf{w}_k \rangle} = \overline{A_{kj}}$$

∎

We revisit our motivating set-up (†) in the language of part 3. Suppose:

- $V = \mathbb{F}^n$ and $W = \mathbb{F}^m$ have standard orthonormal bases $\beta = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ and $\gamma = \{\mathbf{e}_1, \dots, \mathbf{e}_m\}$.

- $T = L_A \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$.

Since the matrix of T with respect to the standard bases is simply A itself, the theorem confirms our earlier observation that the adjoint of L_A is left multiplication by the conjugate-transpose A^*:

$$[T^*]_\gamma^\beta = ([T]_\beta^\gamma)^* = A^* = [L_{A^*}]_\gamma^\beta \implies T^* = (L_A)^* = L_{A^*}$$

Here is another straightforward example, this time using the standard inner products on $\mathbb{C}^2$ and $\mathbb{C}^3$.

Example 2.28.   Let $T = L_A \in \mathcal{L}(\mathbb{C}^3, \mathbb{C}^2)$ where $A = \left(\begin{smallmatrix} -i & 1 & -3 \\ 2 & 1-i & 4+2i \end{smallmatrix} \right)$.
Plainly $A = [T]_\beta^\gamma$ with respect to $\beta = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ and $\gamma = \{\mathbf{e}_1, \mathbf{e}_2\}$. We conclude that $T^* = L_{A^*}$:

$$[T^*]_\gamma^\beta = A^* = \begin{pmatrix} i & 2 \\ 1 & 1+i \\ -3 & 4-2i \end{pmatrix}$$

As a sanity check, multiply out a few examples of $\langle A^*\mathbf{w}, \mathbf{v} \rangle = \langle \mathbf{w}, A\mathbf{v} \rangle$; make sure you're comfortable with the fact that the left inner product is on $\mathbb{C}^2$ and the right on $\mathbb{C}^3$!

The Theorem tells us that every linear map $T \in \mathcal{L}(V, W)$ between finite-dimensional spaces has an adjoint and moreover how to compute it:

1. Choose orthonormal bases (exist by Corollary 2.21) and find the matrix $[T]_\beta^\gamma$.

2. Take the conjugate-transpose $\left([T]_\beta^\gamma\right)^*$ and translate back to find $T^* \in \mathcal{L}(W, V)$.

The prospect of twice applying Gram–Schmidt and translating between linear maps and their matrices is unappealing; calculating this way can quickly become an enormous mess! In practice, it is often better to try a modified approach; see for instance part 2(b) of the next Example.

28

Examples 2.29. Let $T = \frac{d}{dx} \in \mathcal{L}(P_1(\mathbb{R}))$ be the derivative operator; $T(a + bx) = b$. We treat $P_1(\mathbb{R})$ as an inner product space in two ways.

1. Equip the inner product for which the standard basis $\epsilon = \{1, x\}$ is orthonormal. Then

$$[T]_\epsilon = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \implies [T^*]_\epsilon = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \implies T^*(a + bx) = ax$$

2. Equip the L^2 inner product $\langle f, g \rangle = \int_0^1 f(x)g(x)\,dx$. As we saw in Example 2.23, the basis $\beta = \{f_1, f_2\} = \{1, 2x - 1\}$ is orthogonal with $||f_1|| = 1$ and $||f_2|| = \frac{1}{\sqrt{3}}$. We compute the adjoint of $T = \frac{d}{dx}$ in two different ways.

   (a) The basis $\gamma = \{g_1, g_2\} = \left\{f_1, \sqrt{3}f_2\right\} = \left\{1, \sqrt{3}(2x - 1)\right\}$ is orthonormal. Observe that

   $$T(g_1) = 0, \quad T(g_2) = 2\sqrt{3} \implies [T]_\gamma = \begin{pmatrix} 0 & 2\sqrt{3} \\ 0 & 0 \end{pmatrix} \implies [T^*]_\gamma = \begin{pmatrix} 0 & 0 \\ 2\sqrt{3} & 0 \end{pmatrix}$$

   $$\implies T^*(a + bx) = T^* \left(a + \frac{b}{2} + \frac{b}{2\sqrt{3}}\sqrt{3}(2x - 1) \right) = \left(a + \frac{b}{2} \right) \cdot 2\sqrt{3}g_2$$

   $$= 3(2a + b)(2x - 1)$$

   (b) Use the orthogonal basis β and the projection formula (Theorem 2.18). With $p(x) = a + bx$,

   $$T^*(p) = \frac{\langle T^*(p), f_1 \rangle}{||f_1||^2} f_1 + \frac{\langle T^*(p), f_2 \rangle}{||f_2||^2} f_2 = \langle p, T(1) \rangle + \langle p, T(2x - 1) \rangle \cdot 3(2x - 1)$$

   $$= \langle p, 0 \rangle + 3 \langle p, 2 \rangle (2x - 1) = 3 \left(\int_0^1 2(a + bx)\,dx \right)(2x - 1)$$

   $$= 3(2a + b)(2x - 1)$$

   Note the advantages here: no square roots and no need to change basis at the end!

The calculations for the second example were much nastier, even though we were already in possession of an orthogonal basis. The crucial point is that the two examples produce *different maps* T^*: the adjoint depends on the inner product!

Why should we care about adjoints?

Adjoints might seem merely to be an abstraction of something simple (transposes) for its own sake. A convincing explanation of why adjoints are useful takes a lot of work; here is a short version.

Given a linear map $T \in \mathcal{L}(V)$ on an inner product space, we now have *two desirable types of basis.*

1. Eigenbasis: diagonalizes T.

2. Orthonormal basis: recall (Theorem 2.18) how these simplify computations.

The capstone result of this course is the famous *spectral theorem* (Theorem 2.37) which says, in short, that *self-adjoint operators* ($T^* = T$) have an *orthonormal eigenbasis*, the holy grail of easy computation! Such operators are important both theoretically and in applications such as quantum mechanics.

The Fundamental Subspaces Theorem

To every linear map are associated its range and nullspace. These interact nicely with the adjoint...

> **Theorem 2.30.** *If* $T \in \mathcal{L}(V, W)$ *has adjoint* T^*, *then,*
>
> 1. $\mathcal{R}(T^*)^\perp = \mathcal{N}(T)$
>
> 2. *If* V *is finite dimensional, then* $\mathcal{R}(T^*) = \mathcal{N}(T)^\perp$
>
> *The corresponding results hold if we swap* $V \leftrightarrow W$ *and* $T \leftrightarrow T^*$.

The proof is left to Exercise 6. You've likely observed this with transposes of small matrices.

Example 2.31. Let $A = \begin{pmatrix} 1 & 2 & -1 \\ 0 & 3 & -2 \end{pmatrix}$. Viewed as a linear map between Euclidean spaces, $T = L_A$ has adjoint $T^* = L_{A^T}$. It is easy to compute the relevant subspaces:

$$\mathcal{R}(A) = \mathbb{R}^2, \quad \mathcal{N}(A^T) = \{\mathbf{0}\}, \qquad \mathcal{R}(A^T) = \text{Span}\left\{ \begin{pmatrix} 1 \\ 2 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ 3 \\ -2 \end{pmatrix} \right\}, \quad \mathcal{N}(A) = \text{Span}\begin{pmatrix} -1 \\ 2 \\ 3 \end{pmatrix}$$

The Riesz Representation Theorem

This powerful result demonstrates a natural relation between an inner product space and its *dual-space* $V^* = \mathcal{L}(V, \mathbb{F})$.

> **Theorem 2.32.** *If* V *is finite-dimensional and* $g : V \to \mathbb{F}$ *is linear, then there exists a unique* $\mathbf{y} \in V$ *such that* $g(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$.

Example 2.33. $g(p) := \int_0^1 p(x)\, dx$ is a linear map $g : P_2(\mathbb{R}) \to \mathbb{R}$. Equip $P_2(\mathbb{R})$ with the inner product for which the standard basis $\{1, x, x^2\}$ is orthonormal. Then

$$g(a + bx + cx^2) = a + \tfrac{1}{2}b + \tfrac{1}{3}c = \left\langle a + bx + cx^2, 1 + \tfrac{1}{2}x + \tfrac{1}{3}x^2 \right\rangle$$

We conclude that $g(p) = \langle p, q \rangle$, where $q(x) = 1 + \tfrac{1}{2}x + \tfrac{1}{3}x^2$.

The idea of the proof is very simple: if $g(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$ then the nullspace of g must equal $\text{Span}\{\mathbf{y}\}^\perp$...

Proof. If g is the zero map, take $\mathbf{y} = \mathbf{0}$. Otherwise, rank $g = 1$ and

$$V = \mathcal{N}(g) \oplus \mathcal{N}(g)^\perp \quad \text{where} \quad \dim \mathcal{N}(g)^\perp = 1 \qquad \text{(rank–nullity theorem and Exercise 2.2.6)}$$

Let $\mathbf{u} \in \mathcal{N}(g)^\perp$ be either of the two unit vectors and define, independently of $\mathbf{u}$,

$$\mathbf{y} := \overline{g(\mathbf{u})}\mathbf{u} \in V$$

Following the decomposition, write $\mathbf{x} = \mathbf{n} + \alpha\mathbf{u}$ where $\mathbf{n} \in \mathcal{N}(g)$ and observe that

$$\langle \mathbf{x}, \mathbf{y} \rangle = \left\langle \mathbf{n} + \alpha\mathbf{u}, \overline{g(\mathbf{u})}\mathbf{u} \right\rangle = g(\mathbf{u})\alpha = 0 + g(\alpha\mathbf{u}) = g(\mathbf{n} + \alpha\mathbf{u}) = g(\mathbf{x})$$

The uniqueness of $\mathbf{y}$ follows from the cancellation property (Lemma 2.5, part 4). ∎

Due to the tight correspondence, the map is often decorated as g_y. Riesz's theorem indeed says that $\mathbf{y} \mapsto g_\mathbf{y}$ is an isomorphism $V \cong V^*$. While there are infinitely many isomorphisms between these spaces, the inner product structure identifies a *canonical* or preferred choice.

> **Corollary 2.34.** *Every linear map on a finite-dimensional inner product space has an adjoint.*

Note how only the domain is required to be finite-dimensional! Riesz's Theorem and the Corollary also apply to continuous linear operators on (infinite-dimensional) Hilbert spaces, though the proof is a little trickier.

Proof. Let $T \in \mathcal{L}(V, W)$ where $\dim V < \infty$, and suppose $\mathbf{z} \in W$ is given. Simply *define* $T^*(\mathbf{z}) := \mathbf{y}$ where $\mathbf{y} \in V$ is the unique vector in Riesz's Theorem arising from the linear map

$$g : V \to \mathbb{F}, \quad g(\mathbf{x}) = \langle T(\mathbf{x}), \mathbf{z} \rangle$$

and check the required property:

$$\langle T^*(\mathbf{z}), \mathbf{x} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle = \overline{g(\mathbf{x})} = \langle \mathbf{z}, T(\mathbf{x}) \rangle$$

■

Exercises 2.3 1. For each inner product space V and linear operator $T \in \mathcal{L}(V)$, evaluate T^* on the given vector.

    (a) $V = \mathbb{R}^2$ with the standard inner product, $T\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right) = \left(\begin{smallmatrix} 2x+y \\ x-3y \end{smallmatrix}\right)$ and $\mathbf{x} = \left(\begin{smallmatrix} 3 \\ 5 \end{smallmatrix}\right)$

    (b) $V = \mathbb{C}^2$ with the standard inner product, $T\left(\begin{smallmatrix} z \\ w \end{smallmatrix}\right) = \left(\begin{smallmatrix} 2z+iw \\ (1-i)z \end{smallmatrix}\right)$ and $\mathbf{x} = \left(\begin{smallmatrix} 3-i \\ 1+2i \end{smallmatrix}\right)$

    (c) $V = P_1(\mathbb{R})$ with $\langle f, g \rangle = \int_0^1 f(t)g(t)\,dt$, $T(f) = f' + 3f$ and $f(t) = 4 - 2t$

2. Suppose $A = \left(\begin{smallmatrix} 1 & 1 \\ 4 & 3 \end{smallmatrix}\right)$ and consider the linear map $T = L_A \in \mathcal{L}(\mathbb{R}^2)$ where $\mathbb{R}^2$ is equipped with the weighted inner product

$$\langle \mathbf{x}, \mathbf{y} \rangle = 4x_1 y_1 + x_2 y_2$$

    (a) Find the matrix of T with respect to the orthonormal basis $\beta = \{\mathbf{v}_1, \mathbf{v}_2\} = \{\frac{1}{2}\mathbf{e}_1, \mathbf{e}_2\}$.

    (b) Find the adjoint T^* and its matrix with respect to the standard basis $\epsilon = \{\mathbf{e}_1, \mathbf{e}_2\}$.

        (*Hint: the answer isn't A^T!*)

3. Extending Examples 2.29, find the adjoint of $T = \frac{d}{dx} \in \mathcal{L}(P_2(\mathbb{R}))$ with respect to:

    (a) The inner product where the standard basis $\epsilon = \{1, x, x^2\}$ is orthonormal.

    (b) (Hard!) The L^2 inner product $\int_0^1 f(x)g(x)\,dx$.

4. Let $T(f) = f''$ be a linear transformation of $P_2(\mathbb{R})$ and let $\epsilon = \{1, x, x^2\}$ be the standard basis. Find $T^*(a + bx + cx^2)$:

    (a) With respect to the inner product where ϵ is orthonormal;

    (b) With respect to the L^2 inner product $\langle f, g \rangle = \int_{-1}^1 f(t)g(t)\,dt$.

        (*Hint: $\{1, x, 3x^2 - 1\}$ is orthogonal*)

5. Prove part 2 of Theorem 2.27.

6. Prove the Fundamental Subspaces Theorem 2.30.

7. For each inner product space V and linear transformation $g : V \to \mathbb{F}$, find a vector $\mathbf{y} \in V$ such that $g(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$ for all $\mathbf{x} \in V$.

   (a) $V = \mathbb{R}^3$ with the standard inner product, and $g\left(\begin{smallmatrix} x \\ y \\ z \end{smallmatrix}\right) = x - 2y + 4z$

   (b) $V = \mathbb{C}^2$ with the standard inner product, and $g\left(\begin{smallmatrix} z \\ w \end{smallmatrix}\right) = iz - 2w$

   (c) $V = P_2(\mathbb{R})$ with the L^2 inner product $\langle f, h \rangle = \int_0^1 f(x)h(x)\, dx$, and $g(f) = f'(1)$

8. (a) In the proof of Theorem 2.32, explain why $\mathbf{y}$ depends only on g (not $\mathbf{u}$).

   (b) In the proof of Corollary 2.34, check that $g(\mathbf{x}) := \langle T(\mathbf{x}), \mathbf{z} \rangle$ is linear.

9. Let $\mathbf{y}, \mathbf{z} \in V$ be fixed vectors and define $T \in \mathcal{L}(V)$ by $T(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle \mathbf{z}$. Show that T^* exists and find an explicit expression.

10. Suppose $A \in M_{m \times n}(\mathbb{F})$. Prove that A^*A is diagonal if and only if the columns of A are orthogonal. What additionally would it mean if $A^*A = I$?

11. Suppose $T \in \mathcal{L}(V)$ where V is a finite-dimensional inner product space.

    (a) Prove that the eigenvalues of T^* are the complex conjugates of those of T.

    (*Hint: relate the characteristic polynomial $p^*(t) = \det(T^* - tI)$ to that of T*)

    (b) Prove that T^* is diagonalizable if and only if T is.

12. (Hard)   We present two linear maps which do not have an adjoint!

    (a) Since $\epsilon = \{1, x, x^2, \ldots\}$ is a basis of $P(\mathbb{R})$, we may define a linear map $T \in \mathcal{L}(P(\mathbb{R}))$ via $T(x^n) = 1$ for all n; for instance

    $$T(4 + 3x + 2x^5) = 9$$

    Let $\langle\ ,\ \rangle$ be the inner product for which ϵ is orthonormal. If T^* existed, show that

    $$T^*(1) = \sum_{n=0}^{\infty} x^n$$

    would be an *infinite series*: T^* therefore does not exist.

    (b) For a related challenge, recall the space ℓ^2 of square-summable real sequences. For any sequence $(x_n)_{n=1}^{\infty} \in \ell^2$), define $T \in \mathcal{L}(\ell^2)$ via

    $$T(x_n) = \left(\sum_{n=1}^{\infty} \frac{1}{n} x_n, 0, 0, 0, 0, 0, \ldots \right)$$

    Find the adjoint T^*. If $V \leq \ell^2$ is the subspace whose elements have only finitely many non-zero terms, show that the restriction T_V does not have an adjoint.

2.4 Normal & Self-Adjoint Operators and the Spectral Theorem

We now come to the fundamental question of this chapter: for which linear operators $T \in \mathcal{L}(V)$ can we find an orthonormal eigenbasis? Many linear maps are, of course, not even diagonalizable, so in general this is far too much to hope for! Let's see what happens if such a basis exists...

If β is an orthonormal basis of eigenvectors of T, then

$$[T]_\beta = \mathrm{diag}(\lambda_1, \cdots, \lambda_n) \implies [T^*]_\beta = \mathrm{diag}(\overline{\lambda_1}, \cdots, \overline{\lambda_n})$$

If V is a *real* inner product space, then these matrices are identical and so $T^* = T$. In the complex case, we instead observe that

$$[TT^*]_\beta = \mathrm{diag}(|\lambda_1|^2, \cdots, |\lambda_n|^2) = [T^*T]_\beta \implies TT^* = T^*T$$

Definition 2.35. Suppose T is a linear operator on an inner product space V and assume T has an adjoint. We say that T is,

> *Normal* if $TT^* = T^*T$,
>
> *Self-adjoint* if $T^* = T$.

The definitions for square matrices over $\mathbb{R}$ and $\mathbb{C}$ are identical, where * now denotes the conjugate-transpose.

A real self-adjoint matrix $A \in M_n(\mathbb{R})$ is plainly *symmetric*: $A^T = A$. A complex self-adjoint matrix is also called *Hermitian*: $A^* = A$.

If T is self-adjoint then it is certainly normal, but the converse is false as the next example shows.

Example 2.36. The (non-symmetric) real matrix $A = \begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}$ is normal but not self-adjoint:

$$AA^T = \begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix}\begin{pmatrix} 2 & 1 \\ -1 & 2 \end{pmatrix} = \begin{pmatrix} 5 & 0 \\ 0 & 5 \end{pmatrix} = \begin{pmatrix} 2 & 1 \\ -1 & 2 \end{pmatrix}\begin{pmatrix} 2 & -1 \\ 1 & 2 \end{pmatrix} = A^T A$$

More generally, every non-zero *skew-hermitian* matrix $A^* = -A$ is normal but *not* self-adjoint:

$$A^* = -A \implies AA^* = -A^2 = A^*A$$

We saw above that linear maps with an orthonormal eigenbasis are either self-adjoint or normal depending whether the inner product space is real or complex. Amazingly, this provides a complete characterisation of such maps!

Theorem 2.37 (Spectral Theorem, version 1). *Let* T *be a linear operator on a finite-dimensional inner product space* V.

Complex case: V *has an orthonormal basis of eigenvectors of* T *if and only if* T *is normal.*

Real case: V *has an orthonormal basis of eigenvectors of* T *if and only if* T *is self-adjoint.*

The theorem gets is name from the *spectrum* (set of eigenvalues) of T.

Examples 2.38. 1. We diagonalize the self-adjoint linear map $T = L_A \in \mathcal{L}(\mathbb{R}^2)$ where $A = \left(\begin{smallmatrix} 6 & 3 \\ 3 & -2 \end{smallmatrix}\right)$.

Characteristic polynomial $p(t) = (6-t)(-2-t) - 9 = t^2 - 4t - 21 = (t-7)(t+3)$

Eigenvalues $\lambda_1 = 7, \quad \lambda_2 = -3$

Eigenvectors (normalized) $\mathbf{w}_1 = \frac{1}{\sqrt{10}} \left(\begin{smallmatrix} 3 \\ 1 \end{smallmatrix}\right), \quad \mathbf{w}_2 = \frac{1}{\sqrt{10}} \left(\begin{smallmatrix} -1 \\ 3 \end{smallmatrix}\right)$

The basis $\beta = \{\mathbf{w}_1, \mathbf{w}_2\}$ is orthonormal, with respect to which $[T]_\beta = \left(\begin{smallmatrix} 7 & 0 \\ 0 & -3 \end{smallmatrix}\right)$ is diagonal.

2. The map $T = L_A \in \mathcal{L}(\mathbb{R}^2)$ where $A = \left(\begin{smallmatrix} 1 & 3 \\ 0 & -2 \end{smallmatrix}\right)$ is neither self-adjoint nor normal:

$$AA^* = \begin{pmatrix} 1 & 3 \\ 0 & -2 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 3 & -2 \end{pmatrix} = \begin{pmatrix} 10 & -6 \\ -6 & 4 \end{pmatrix} \neq \begin{pmatrix} 1 & 3 \\ 3 & 13 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 3 & -2 \end{pmatrix} \begin{pmatrix} 1 & 3 \\ 0 & -2 \end{pmatrix} = A^*A$$

It is diagonalizable, indeed

$$\gamma = \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ -1 \end{pmatrix} \right\} \rightsquigarrow [T]_\gamma = \begin{pmatrix} 1 & 0 \\ 0 & -2 \end{pmatrix}$$

In accordance with the spectral theorem, γ is not orthogonal.

3. Let $A = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ and consider $T = L_A$ acting on both $\mathbb{C}^2$ and $\mathbb{R}^2$. Since T is normal but not self-adjoint, we'll see how the field really matters in the spectral theorem.

First the complex case: $T \in \mathcal{L}(\mathbb{C}^2)$ is normal and thus diagonalizable with respect to an orthonormal basis of eigenvectors. Here are the details.

Characteristic polynomial $p(t) = t^2 + 1 = (t-i)(t+i)$

Eigenvalues $\lambda_1 = i, \quad \lambda_2 = -i$

Eigenvectors (normalized) $\mathbf{w}_1 = \frac{1}{\sqrt{2}} \left(\begin{smallmatrix} i \\ 1 \end{smallmatrix}\right), \quad \mathbf{w}_2 = \frac{1}{\sqrt{2}} \left(\begin{smallmatrix} -i \\ 1 \end{smallmatrix}\right)$

Certainly $\langle \mathbf{w}_1, \mathbf{w}_2 \rangle = 0$, $\beta = \{\mathbf{w}_1, \mathbf{w}_2\}$ is orthonormal, and $[T]_\beta = \left(\begin{smallmatrix} i & 0 \\ 0 & -i \end{smallmatrix}\right)$ is diagonal.

Now for the real case: $T \in \mathcal{L}(\mathbb{R}^2)$ is not self-adjoint and thus should not be diagonalizable with respect to an orthonormal basis of eigenvectors. Indeed this is trivial; the characteristic polynomial has no roots in $\mathbb{R}$ and so there are no real eigenvalues! It is also clear geometrically: T is rotation by 90° counter-clockwise around the origin, so it has no eigenvectors.

4. Let $A = \left(\begin{smallmatrix} 3 & i \\ -i & 3 \end{smallmatrix}\right)$ and consider the self-adjoint operator $T = L_A \in \mathcal{L}(\mathbb{C}^2)$.

Characteristic polynomial $p(t) = t^2 - 6t + 9 - 1 = (t-2)(t-4)$

Eigenvalues $\lambda_1 = 2, \quad \lambda_2 = 4$

Eigenvectors (normalized) $\mathbf{w}_1 = \frac{1}{\sqrt{2}} \left(\begin{smallmatrix} 1 \\ i \end{smallmatrix}\right), \quad \mathbf{w}_2 = \frac{1}{\sqrt{2}} \left(\begin{smallmatrix} 1 \\ -i \end{smallmatrix}\right)$

With respect to the orthonormal basis $\beta = \{\mathbf{w}_1, \mathbf{w}_2\}$, we have $[T]_\beta = \left(\begin{smallmatrix} 2 & 0 \\ 0 & 4 \end{smallmatrix}\right)$.

Proving the Spectral Theorem for Self-Adjoint Operators

> **Lemma 2.39 (Basic properties of self-adjoint operators).**   *Let* $T \in \mathcal{L}(V)$ *be self-adjoint.*
>
> 1. *If $W \leq V$ is T-invariant then the restriction T_W is self-adjoint;*
>
> 2. *Every eigenvalue of T is real;*
>
> 3. *If $\dim V$ is finite then T has an eigenvalue.*

It is irrelevant whether V is real or complex. The previous example demonstrates part 2; even when $V = \mathbb{C}^2$ is a complex inner product space, the eigenvalues of a self-adjoint matrix are real.

Proof.   1. Let $\mathbf{w}_1, \mathbf{w}_2 \in W$. Then

$$\langle T_W(\mathbf{w}_1), \mathbf{w}_2 \rangle = \langle T(\mathbf{w}_1), \mathbf{w}_2 \rangle = \langle \mathbf{w}_1, T(\mathbf{w}_2) \rangle = \langle \mathbf{w}_1, T_W(\mathbf{w}_2) \rangle$$

2. Suppose $(\lambda, \mathbf{x})$ is an eigenpair. Then

$$\lambda \, ||\mathbf{x}||^2 = \langle T(\mathbf{x}), \mathbf{x} \rangle = \langle \mathbf{x}, T(\mathbf{x}) \rangle = \overline{\lambda} \, ||\mathbf{x}||^2 \implies \lambda \in \mathbb{R}$$

3. This is trivial if V is complex since every characteristic polynomial splits over $\mathbb{C}$. We therefore assume V is real. Choose any orthonormal basis γ of V, let $A = [T]_\gamma \in M_n(\mathbb{R})$, and define $S := L_A \in \mathcal{L}(\mathbb{C}^n)$. Then;

   - The characteristic polynomial of S splits over $\mathbb{C}$, whence there exists an eigenvalue $\lambda \in \mathbb{C}$.
   - The characteristic polynomials of S and T are *identical* (to that of A).
   - S is self-adjoint and thus (part 2) $\lambda \in \mathbb{R}$.

   It follows that T has the same *real* eigenvalue λ. ∎

We are now able to prove the spectral theorem for self-adjoint operators on a finite-dimensional inner product space V. The argument applies regardless of whether V is real or complex.

Proof of the Spectral Theorem (self-adjoint case). We prove by induction on $\dim V$.

(Base case)   If $\dim V = 1$, then $V = \mathrm{Span}\{\mathbf{x}\}$ and $T(\mathbf{x}) = \lambda \mathbf{x}$ for some unit vector $\mathbf{x}$ and scalar $\lambda \in \mathbb{R}$; plainly $\{\mathbf{x}\}$ is an orthonormal eigenbasis for T.

(Induction step)   Fix $n \in \mathbb{N}$ and assume that *every* self-adjoint operator on *every* inner product space of dimension n satisfies the spectral theorem. Let $\dim V = n + 1$ and $T \in \mathcal{L}(V)$ be self-adjoint. By part 3 of the Lemma, T has an eigenpair $(\lambda, \mathbf{x})$ where we may assume $\mathbf{x}$ has unit length. Let $W = \mathrm{Span}\{\mathbf{x}\}^{\perp}$. If $\mathbf{w} \in W$, then

$$\langle \mathbf{x}, T(\mathbf{w}) \rangle = \langle T(\mathbf{x}), \mathbf{w} \rangle = \lambda \langle \mathbf{x}, \mathbf{w} \rangle = 0 \qquad (*)$$

whence W is T-invariant. Plainly $\dim W = n$. By part 1 of the Lemma, T_W is self-adjoint. By the induction hypothesis, T_W is diagonalized by some orthonormal basis γ of W. But then T is diagonalized by the orthonormal basis $\beta = \gamma \cup \{\mathbf{x}\}$ of V. ∎

Proving the Spectral Theorem for Normal Operators

What changes for normal operators on complex inner product spaces? Not much! Indeed the proof is almost identical when T is merely normal.

- We don't need parts 2 and 3 of Lemma 2.39: every linear operator on a finite-dimensional complex inner product space has an eigenvalue and we no longer care whether eigenvalues are real.

- Two parts of the induction step need fixed:
  - *W* being T-invariant: This isn't quite as simple as (∗), but thankfully part 3 of the next result provides the needed correction.
  - T_W being normal: We need a replacement for part 1 of Lemma 2.39; this is a little more involved.

Rather than write out all the details, we leave this to Exercises 6 and 7.

For completeness, and as an analogue/extension of Lemma 2.39, we summarize some of the basic properties of normal operators. These also apply to self-adjoint operators as a special case.

Lemma 2.40 (Basic properties of normal operators).   *Let* T *be normal on* V. *Then:*

1. $\forall \mathbf{x} \in V, ||T(\mathbf{x})|| = ||T^*(\mathbf{x})||$.

2. $T - tI$ *is normal for any scalar* t.

3. $T(\mathbf{x}) = \lambda\mathbf{x} \iff T^*(\mathbf{x}) = \bar{\lambda}\mathbf{x}$ *so that* T *and* T^* *have the same eigenvectors and conjugate eigenvalues. This recovers the previously established fact that* $\lambda \in \mathbb{R}$ *if* T *is self-adjoint.*

4. *Distinct eigenvalues of* T *have orthogonal eigenvectors.*

Proof.   1. $||T(\mathbf{x})||^2 = \langle T(\mathbf{x}), T(\mathbf{x})\rangle = \langle T^*T(\mathbf{x}), \mathbf{x}\rangle = \langle TT^*(\mathbf{x}), \mathbf{x}\rangle = \langle T^*(\mathbf{x}), T^*(\mathbf{x})\rangle = ||T^*(\mathbf{x})||^2$.

2. $\langle \mathbf{x}, (T - tI)(\mathbf{y})\rangle = \langle \mathbf{x}, T(\mathbf{y})\rangle - \bar{t}\langle \mathbf{x}, \mathbf{y}\rangle = \langle T^*(\mathbf{x}), \mathbf{y}\rangle - \langle \bar{t}\mathbf{x}, \mathbf{y}\rangle = \langle (T^* - \bar{t}I)(\mathbf{x}), \mathbf{y}\rangle$ shows that $T - tI$ has adjoint $T^* - \bar{t}I$. It is trivial to check that these commute.

3. $T(\mathbf{x}) = \lambda\mathbf{x} \iff ||(T - \lambda I)(\mathbf{x})|| = 0 \overset{\text{parts}}{\underset{1\&2}{\iff}} ||(T^* - \bar{\lambda}I)(\mathbf{x})|| = 0 \iff T^*(\mathbf{x}) = \bar{\lambda}\mathbf{x}$.

4. In part this follows from the spectral theorem, but we can also prove more straightforwardly. Suppose $T(\mathbf{x}) = \lambda\mathbf{x}$ and $T(\mathbf{y}) = \mu\mathbf{y}$ where $\lambda \neq \mu$. By part 3,

$$\lambda\langle \mathbf{x}, \mathbf{y}\rangle = \langle \lambda\mathbf{x}, \mathbf{y}\rangle = \langle T(\mathbf{x}), \mathbf{y}\rangle = \langle \mathbf{x}, T^*(\mathbf{y})\rangle = \langle \mathbf{x}, \bar{\mu}\mathbf{y}\rangle$$
$$= \mu\langle \mathbf{x}, \mathbf{y}\rangle$$

This is a contradiction unless $\langle \mathbf{x}, \mathbf{y}\rangle = 0$. ∎

Schur's Lemma

It is reasonable to ask how useful an orthonormal basis can be in general. Here is one answer.

> **Lemma 2.41 (Schur).** *Suppose* T *is a linear operator on a finite-dimensional inner product space* V. *If the characteristic polynomial of* T *splits, then there exists an orthonormal basis* β *of* V *such that* $[T]_\beta$ *is upper-triangular.*

The spectral theorem is a special case; since the proof is similar, we leave it to the exercises.

The conclusion of Schur's lemma is weaker than the spectral theorem, though it applies to more operators: indeed if V is complex, it applies to *any* T! Every example of the spectral theorem is also an example of Schur's lemma. Example 2.38.2 provides another, since the matrix A is already upper triangular with respect to the standard orthonormal basis. Here is a another example.

Example 2.42. Consider $T(f) = 2f'(x) + xf(1)$ as a linear map $T \in \mathcal{L}(P_1(\mathbb{R}))$ with respect to the L^2 inner product $\langle f, g \rangle = \int_0^1 f(t)g(t)\,dt$. We have

$$T(a + bx) = 2b + (a + b)x$$

If $[T]_\beta$ is to be upper-triangular, the first vector in β must be an eigenvector of T. It is easily checked that $f_1 = 1 + x$ is such with eigenvalue 2. To find a basis satisfying Schur's lemma, we need only find f_2 orthogonal to this and then normalize. This can be done by brute force since the problem is small, but for the sake of practice we apply Gram–Schmidt to the polynomial 1:

$$1 - \frac{\langle 1, 1+x \rangle}{||1+x||^2}(1+x) = 1 - \frac{1 + \frac{1}{2}}{1 + 1 + \frac{1}{3}}(1+x) = \frac{1}{14}(5 - 9x) \implies f_2 = 5 - 9x$$

Indeed we obtain an upper-triangular matrix for T:

$$T(f_2) = -18 - 4x = -13(1 + x) - (5 - 9x) = -13f_1 - f_2 \implies [T]_{\{f_1, f_2\}} = \begin{pmatrix} 2 & -13 \\ 0 & -1 \end{pmatrix}$$

We can also work with the corresponding orthonormal basis as posited in the theorem, though the matrix is messier:

$$\beta = \{g_1, g_2\} = \left\{ \sqrt{\frac{3}{7}}(1+x), \frac{1}{\sqrt{7}}(5 - 9x) \right\} \implies [T]_\beta = \begin{pmatrix} 2 & -\frac{13}{\sqrt{3}} \\ 0 & -1 \end{pmatrix}$$

Alternatively, we could have started with the other eigenvector $h_1 = 2 - x$: an orthogonal vector to this is $h_2 = 4 - 9x$, with respect to which

$$[T]_{\{h_1, h_2\}} = \begin{pmatrix} -1 & -13 \\ 0 & 2 \end{pmatrix}$$

In both cases the eigenvalues are down the diagonal, as must be for an upper-triangular matrix.

In general, it is difficult to quickly find a suitable basis satisfying Schur's lemma. After trying the proof in the exercises, you should be able to describe a method, though it is impractically slow!

Exercises 2.4   1. For each linear operator T on an inner product space V, decide whether T is normal, self-adjoint, or neither. If the spectral theorem permits, find an orthonormal eigenbasis.

(a) $V = \mathbb{R}^2$ and $\mathrm{T}\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right) = \left(\begin{smallmatrix} 2a-b \\ -2a+5b \end{smallmatrix}\right)$

(b) $V = \mathbb{R}^3$ and $\mathrm{T}\left(\begin{smallmatrix} x \\ y \\ z \end{smallmatrix}\right) = \left(\begin{smallmatrix} -a-b \\ 5b \\ 4a-2b+5c \end{smallmatrix}\right)$

(c) $V = \mathbb{C}^2$ and $\mathrm{T}\left(\begin{smallmatrix} z \\ w \end{smallmatrix}\right) = \left(\begin{smallmatrix} 2z+iw \\ z+2w \end{smallmatrix}\right)$

(d) $V = \mathbb{R}^4$ with $\mathrm{T} : (\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4) \mapsto (\mathbf{e}_3, \mathbf{e}_4, \mathbf{e}_1, \mathbf{e}_2)$

(e) $V = P_2(\mathbb{R})$ with $\langle f, g \rangle = \int_0^1 f(t)g(t)\, dt$ and $\mathrm{T}(f) = f'$

(*Hint: Don't compute* T*! *Instead assume* T *is normal and aim for a contradiction…*)

2. Let $\mathrm{T}(f(x)) = f'(x) + 4xf(0)$ where $\mathrm{T} \in \mathcal{L}(P_1(\mathbb{R}))$ and $\langle f, g \rangle = \int_{-1}^1 f(t)g(t)\, dt$. Find an orthonormal basis of $P_1(\mathbb{R})$ with respect to which the matrix of T is upper-triangular.

3. Suppose S, T are self-adjoint operators on an inner product space V. Prove that ST is self-adjoint if and only if $\mathrm{ST} = \mathrm{TS}$.

(*Hint: recall Theorem 2.27*)

4. Let T be normal on a finite-dimensional inner product space V. Prove that $\mathcal{N}(\mathrm{T}^*) = \mathcal{N}(\mathrm{T})$ and that $\mathcal{R}(\mathrm{T}^*) = \mathcal{R}(\mathrm{T})$.

(*Hint: Use Lemma 2.40 and the Fundamental Subspaces Theorem 2.30*)

5. Let T be self-adjoint on a finite-dimensional inner product space V. Prove that

$$\forall \mathbf{x} \in V, \quad ||\mathrm{T}(\mathbf{x}) \pm i\mathbf{x}||^2 = ||\mathrm{T}(\mathbf{x})||^2 + ||\mathbf{x}||^2$$

Hence prove that $\mathrm{T} - i\mathrm{I}$ is invertible and that $[(\mathrm{T} - i\mathrm{I})^{-1}]^* = (\mathrm{T} + i\mathrm{I})^{-1}$.

6. Let W be a T-invariant subspace of an inner product space V and let $\mathrm{T}_W \in \mathcal{L}(W)$ be the restriction of T to W. Prove:

(a) $W^\perp$ is T*-invariant.

(b) If W is both T- and T*-invariant, then $(\mathrm{T}_W)^* = (\mathrm{T}^*)_W$.

(c) If W is both T- and T*-invariant and T is normal, then T_W is normal.

7. Use the previous question to complete the proof of the spectral theorem for a normal operator on a finite-dimensional complex inner product space.

8. (a) Suppose S is a normal operator on a finite-dimensional *complex* inner product space, all of whose eigenvalues are *real.* Prove that S is self-adjoint.

(b) Let T be a normal operator on a finite-dimensional *real* inner product space V whose characteristic polynomial splits. Prove that T is self-adjoint and that there exists an orthonormal basis of V of eigenvectors of T.

(*Hint: Mimic the proof of Lemma 2.39 part 3 and use part (a)*)

9. Prove Schur's lemma by induction, similarly to the proof of the spectral theorem.

(*Hint:* T* *has an eigenvector* $\mathbf{x}$*; why? Now show that* $W = \mathrm{Span}\{\mathbf{x}\}^\perp$ *is* T-*invariant…*)

2.5 Unitary and Orthogonal Operators and their Matrices

In this section we focus on length-preserving transformations of an inner product space.

> **Definition 2.43.** A *linear*[9]*isometry* of an inner product space V is a linear map T satisfying
>
> $$\forall \mathbf{x} \in V, \ ||T(\mathbf{x})|| = ||\mathbf{x}||$$

Every eigenvalue of an isometry must have modulus 1: if $T(\mathbf{w}) = \lambda\mathbf{w}$, then

$$||\mathbf{w}||^2 = ||T(\mathbf{w})||^2 = ||\lambda\mathbf{w}||^2 = |\lambda|^2 \, ||\mathbf{w}||^2$$

Example 2.44. Let $T = L_A \in \mathcal{L}(\mathbb{R}^2)$, where $A = \frac{1}{5}\left(\begin{smallmatrix} 4 & -3 \\ 3 & 4 \end{smallmatrix}\right)$. Then

$$\left|\left|T\begin{pmatrix} x \\ y \end{pmatrix}\right|\right|^2 = \left|\left|\frac{1}{5}\begin{pmatrix} 4x - 3y \\ 3x + 4y \end{pmatrix}\right|\right|^2 = \frac{1}{25}\left((4x - 3y)^2 + (3x + 4y)^2\right) = x^2 + y^2 = \left|\left|\begin{pmatrix} x \\ y \end{pmatrix}\right|\right|^2$$

This matrix is very special in that its inverse equals its transpose:

$$A^{-1} = \frac{1}{\frac{16}{25} + \frac{9}{25}}\begin{pmatrix} 4 & 3 \\ -3 & 4 \end{pmatrix} = \frac{1}{5}\begin{pmatrix} 4 & 3 \\ -3 & 4 \end{pmatrix} = A^T$$

We call such matrices *orthogonal.* The simple version of what follows is that every linear isometry on $\mathbb{R}^n$ is multiplication by an orthogonal matrix.

> **Definition 2.45.** A *unitary operator* T on an inner product space V is an invertible linear map satisfying $T^*T = I = TT^*$. A *unitary matrix* is a (real or complex) matrix satisfying $A^*A = I$.
>
> If V is real, we usually call these *orthogonal operators/matrices*; this isn't necessary, since *unitary* encompasses both real and complex spaces. An orthogonal matrix satisfies $A^TA = I$.

Example 2.46. The matrix $A = \frac{1}{3}\left(\begin{smallmatrix} i & 2+2i \\ 2-2i & i \end{smallmatrix}\right)$ is unitary:

$$A^*A = \frac{1}{9}\begin{pmatrix} -i & 2+2i \\ 2-2i & -i \end{pmatrix}\begin{pmatrix} i & 2+2i \\ 2-2i & i \end{pmatrix} = \frac{1}{9}\begin{pmatrix} -i^2 + 4 + 4 & (-i+i)(2+2i) \\ (i-i)(2-2i) & 4 + 4 - i^2 \end{pmatrix} = I$$

If V is finite-dimensional, the operator/matrix notions correspond straightforwardly. By Theorem 2.27, if we choose any orthonormal basis β of V, then

$$T \in \mathcal{L}(V) \text{ is unitary/orthogonal} \iff [T]_\beta \text{ is unitary/orthogonal}$$

Moreover, we need only assume $T^*T = I$ (or $TT^* = I$) if V is finite-dimensional: if β is an orthonormal basis, then

$$T^*T = I \iff [T^*]_\beta[T]_\beta = I \iff [T]_\beta[T^*]_\beta = I \iff TT^* = I$$

In infinite dimensions, we need T^* to be both the left- and right-inverse of T. This isn't an empty requirement (see Exercise 13).

---

[9]There also exist *non-linear* isometries: for instance *translations* ($T(\mathbf{x}) = \mathbf{x} + \mathbf{a}$ for any constant $\mathbf{a}$) and *complex conjugation* ($T(\mathbf{x}) = \overline{\mathbf{x}}$) on $\mathbb{C}^n$. Together with linear isometries, these essentially comprise all isometries in finite dimensions.

We now tackle the correspondence between unitary operators and isometries.

Theorem 2.47. *Let* T *be a linear operator on an inner product space* V.

1. *If* T *is a unitary/orthogonal operator, then it is a linear isometry.*

2. *If* T *is a linear isometry and* V *is finite-dimensional, then* T *is unitary/orthogonal.*

Proof.   1. If T is unitary, then

$$\forall \mathbf{x}, \mathbf{y} \in V, \quad \langle \mathbf{x}, \mathbf{y} \rangle = \langle T^*T(\mathbf{x}), \mathbf{y} \rangle = \langle T(\mathbf{x}), T(\mathbf{y}) \rangle \tag{†}$$

In particular taking $\mathbf{x} = \mathbf{y}$ shows that T is an isometry.

2. $(I - T^*T)^* = I^* - (T^*T)^* = I - T^*T$ is self-adjoint. By the spectral theorem, there exists an orthonormal basis of V of eigenvectors of $I - T^*T$. For any such $\mathbf{x}$ with (real) eigenvalue λ,

$$0 = ||\mathbf{x}||^2 - ||T(\mathbf{x})||^2 = \langle \mathbf{x}, \mathbf{x} \rangle - \langle T(\mathbf{x}), T(\mathbf{x}) \rangle = \langle \mathbf{x}, (I - T^*T)\mathbf{x} \rangle = \lambda ||\mathbf{x}||^2$$
$$\implies \lambda = 0$$

Since $I - T^*T = 0$ on a basis, $T^*T = I$. Since V is finite-dimensional, we also have $TT^* = I$ whence T is unitary. ∎

The finite-dimensional restriction is important in part 2: we use the existence of adjoints, the spectral theorem, and that a left-inverse is also a right-inverse. See Exercise 13 for an example of a *non-unitary* isometry in infinite dimensions.

The proof shows a little more:

Corollary 2.48. *On a finite-dimensional space, being unitary is equivalent to each of the following:*

(a) *Preservation of the inner product*[10] *(†).*

(b) *The existence of an orthonormal basis* $\beta = \{\mathbf{w}_1, \ldots, \mathbf{w}_n\}$ *such that* $T(\beta) = \{T(\mathbf{w}_1), \ldots, T(\mathbf{w}_n)\}$ *is also orthonormal.*

(c) *That* every *orthonormal basis* β *of* V *is mapped to an orthonormal basis* $T(\beta)$.

While (a) is simply (†), claims (b) and (c) are also worth proving explicitly: see Exercise 9. If β is the standard orthonormal basis of $\mathbb{F}^n$ and $T = L_A$, then the columns of A form the orthonormal set $T(\beta)$. This makes identifying unitary/orthogonal matrices easy:

Corollary 2.49. *A matrix* $A \in M_n(\mathbb{R})$ *is orthogonal if and only if its columns form an orthonormal basis of* $\mathbb{R}^n$ *with respect to the standard (dot) inner product.*

A matrix $A \in M_n(\mathbb{C})$ *is unitary if and only if its columns form an orthonormal basis of* $\mathbb{C}^n$ *with respect to the standard (Hermitian) inner product.*

---

[10]In particular, in a real inner product space isometries also preserve the *angle* θ between vectors since $\cos\theta = \frac{\langle \mathbf{x},\mathbf{y} \rangle}{||\mathbf{x}||||\mathbf{y}||}$.

Examples 2.50. 1. The matrix $A_\theta = \left(\begin{smallmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{smallmatrix}\right) \in M_2(\mathbb{R})$ is orthogonal for any θ. Example 2.44 is this with $\theta = \tan^{-1}\frac{3}{4} = \sin^{-1}\frac{3}{5} = \cos^{-1}\frac{4}{5}$. More generally (Exercise 6), it can be seen that every real orthogonal 2×2 matrix has the form A_θ or

$$B_\theta = \begin{pmatrix} \cos\theta & \sin\theta \\ \sin\theta & -\cos\theta \end{pmatrix}$$

for some angle θ. The effect of the L_{A_θ} is to rotate counter-clockwise by θ, while that of L_{B_θ} is to reflect across the line making angle $\frac{1}{2}\theta$ with the positive x-axis.

2. $A = \frac{1}{\sqrt{6}}\left(\begin{smallmatrix} \sqrt{2} & \sqrt{3} & 1 \\ \sqrt{2} & 0 & -2 \\ -\sqrt{2} & \sqrt{3} & -1 \end{smallmatrix}\right) \in M_3(\mathbb{R})$ is orthogonal: check the columns!.

3. $A = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix} 1 & i \\ i & 1 \end{smallmatrix}\right) \in M_2(\mathbb{C})$ is unitary: indeed it maps the standard basis to the orthonormal basis

$$T(\beta) = \left\{ \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ i \end{pmatrix}, \frac{1}{\sqrt{2}}\begin{pmatrix} i \\ 1 \end{pmatrix} \right\}$$

It is also easy to check that the characteristic polynomial is

$$p(t) = \det\begin{pmatrix} \frac{1}{\sqrt{2}} - t & \frac{i}{\sqrt{2}} \\ \frac{i}{\sqrt{2}} & \frac{1}{\sqrt{2}} - t \end{pmatrix} = \left(t - \frac{1}{\sqrt{2}}\right)^2 + \frac{1}{2} \implies t = \frac{1}{\sqrt{2}}(1 \pm i) = e^{\pm\pi i/4}$$

whence the eigenvalues of T both have modulus 1.

4. Here is an example of an infinite-dimensional unitary operator. On the space $C[-\pi, \pi]$, the function $T(f(x)) = e^{ix}f(x)$ is linear. Moreover

$$\left\langle e^{ix}f(x), g(x) \right\rangle = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{ix}f(x)\overline{g(x)}\,dx = \frac{1}{2\pi}\int_{-\pi}^{\pi} f(x)\overline{e^{-ix}g(x)}\,dx = \left\langle f(x), e^{-ix}g(x) \right\rangle$$

whence $T^*(f(x)) = e^{-ix}f(x)$. Indeed $T^* = T^{-1}$ and so T is a unitary operator.

Since $C[-\pi, \pi]$ is infinite-dimensional, we don't expect all parts of the Corollary to hold:

(a) T does preserve the inner product.

(b), (c) $C[-\pi, \pi]$ doesn't have an orthonormal basis; there is no orthonormal set $\beta = \{f_k\}$ so that every continuous function is a *finite* linear combination.[11] We cannot therefore claim that T maps orthonormal bases to orthonormal bases!

---

[11] An *infinite* orthonormal set $\beta = \{f_k : k \in \mathbb{Z}\}$ can be found so that every function f 'equals' an *infinite series* in the sense that $\|f - \sum a_k f_k\| = 0$. Since these are not finite sums, β isn't strictly a basis, though it isn't uncommon for it to be so described. Moreover, given that the norm is defined by an integral, this also isn't a claim that f and $\sum a_k f_k$ are equal as functions. Indeed the infinite series need not be continuous! For these reasons, when working with Fourier series, one tends to consider a broader class than the continuous functions.

Unitary and Orthogonal Equivalence

Suppose $A \in M_n(\mathbb{R})$ is symmetric (self-adjoint) $A^T = A$. By the spectral theorem, A has an orthonormal eigenbasis $\beta = \{\mathbf{w}_1, \ldots, \mathbf{w}_n\}$: $A\mathbf{w}_j = \lambda_j \mathbf{w}_j$. Arranging the eigenbasis as the columns of a matrix, we see that the columns of $U = (\mathbf{w}_1 \cdots \mathbf{w}_n)$ are orthonormal and so U is an orthogonal matrix. We can therefore write

$$A = UDU^{-1} = U \begin{pmatrix} \lambda_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_n \end{pmatrix} U^T$$

A similar approach works if $A \in M_n(\mathbb{C})$ is normal: we now have $A = UDU^*$ where U is unitary.

Example 2.51. The matrix $A = \left(\begin{smallmatrix} 1+i & 1+i \\ -1-i & 1+i \end{smallmatrix} \right)$ is normal as can easily be checked. Its characteristic polynomial is

$$p(t) = t^2 - 2(1+i)t + 4i = (t - 2i)(t - 2)$$

with corresponding orthonormal eigenvectors

$$\mathbf{w}_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix}, \qquad \mathbf{w}_{2i} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix}$$

We conclude that

$$A = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix} \begin{pmatrix} 2 & 0 \\ 0 & 2i \end{pmatrix} \left[\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix} \right]^{-1} = \begin{pmatrix} 1 & 1 \\ -i & i \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix} \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix}$$

This is an example of unitary equivalence.

Definition 2.52. Square matrices A, B are *unitarily equivalent* if there exists a unitary matrix U such that $B = U^* A U$. *Orthogonal equivalence* is similar: $B = U^T A U$.

The above discussion proves half the following:

Theorem 2.53. $A \in M_n(\mathbb{C})$ *is normal if and only if it is unitarily equivalent to a diagonal matrix (the matrix of its eigenvalues).*

$A \in M_n(\mathbb{R})$ *is self-adjoint (symmetric) if and only if it is orthogonally equivalent to a diagonal matrix.*

Proof. We've already observed the ($\Rightarrow$) direction.

For the converse, let D be diagonal, U unitary, and $A = U^* D U$. Then

$$A^* A = (U^* D U)^* U^* D U = U^* D^* U U^* D U = U^* \overline{D} D U = U^* D \overline{D} U = U^* D U U^* \overline{D} U$$
$$= U^* D U (U^* D U)^* = A A^*$$

since $U^* = U^{-1}$ and because diagonal matrices commute: $\overline{D} D = D \overline{D}$.

In the special case where A is real and U is orthogonal, then A is symmetric:

$$A^T = (U^T D U)^T = U^T D^T U = U^T D U = A$$

$\blacksquare$

Exercises 2.5 1. For each matrix A find an orthogonal or unitary U and a diagonal $D = U^*AU$.

(a) $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ (b) $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ (c) $\begin{pmatrix} 2 & 3-3i \\ 3+3i & 5 \end{pmatrix}$ (d) $\begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}$

2. Which of the following pairs are unitarily/orthogonally equivalent? Explain your answers.

(a) $A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ and $B = \begin{pmatrix} 0 & 2 \\ 2 & 0 \end{pmatrix}$ (b) $A = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}$

(c) $A = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & i & 0 \\ 0 & 0 & -i \end{pmatrix}$

3. Let $a, b \in \mathbb{C}$ be such that $|a|^2 + |b|^2 = 1$. Prove that every 2×2 matrix of the form $\begin{pmatrix} a & -e^{i\theta}\bar{b} \\ b & e^{i\theta}\bar{a} \end{pmatrix}$ is unitary. Are these all the unitary 2×2 matrices? Prove or disprove.

4. If A, B are orthogonal/unitary, prove that AB and A^{-1} are also orthogonal/unitary.
   (*This proves that orthogonal/unitary matrices are groups under matrix multiplication*)

5. Check that $A = \frac{1}{3} \begin{pmatrix} 5 & -4i \\ 4i & 5 \end{pmatrix} \in M_2(\mathbb{C})$ satisfies $A^T A = I$ (it is a *complex orthogonal* matrix).
   (*These don't have the same nice relationship with inner products, and are thus less useful to us*)

6. Supply the details of Exercise 2.50.1.
   (*Hints:* $\beta = \{\mathbf{i}, \mathbf{j}\}$ *is orthonormal, whence* $\{A\mathbf{i}, A\mathbf{j}\}$ *must be orthonormal. Now draw pictures to compute the result of rotating and reflecting the vectors* $\mathbf{i}$ *and* $\mathbf{j}$.)

7. Show that the linear map in Example 2.50.4 has no eigenvectors.

8. Prove that $A \in M_n(\mathbb{C})$ has an orthonormal basis of eigenvectors whose eigenvalues have modulus 1, if and only if A is unitary.

9. Prove parts (b) and (c) of Corollary 2.48 for a finite-dimensional inner product space:

(a) If β is an orthonormal basis such that $T(\beta)$ is orthonormal, then T is unitary.

(b) If T is unitary, and η is an orthonormal basis, then $T(\eta)$ is an orthonormal basis.

10. Let T be a linear operator on a finite-dimensional inner product space V. If $\|T(\mathbf{x})\| = \|\mathbf{x}\|$ for all $\mathbf{x}$ in some orthonormal basis of V, must T be unitary? Prove or disprove.

11. Let T be a unitary operator on an inner product space V and let W be a finite-dimensional T-invariant subspace of V. Prove:

(a) $T(W) = W$ (*Hint: show that* T_W *is injective*);

(b) $W^\perp$ is T-invariant.

12. Let W a subspace of an inner product space V such that $V = W \oplus W^\perp$. Define $T \in \mathcal{L}(V)$ by $T(\mathbf{u} + \mathbf{w}) = \mathbf{u} - \mathbf{w}$ where $\mathbf{u} \in W$ and $\mathbf{w} \in W^\perp$. Prove that T is unitary and self-adjoint.

13. In the inner product space ℓ^2 of square-summable sequences, consider the linear operator $T(x_1, x_2, \ldots) = (0, x_1, x_2, \ldots)$. Prove that T is an isometry and compute its adjoint. Check that T is non-invertible and non-unitary.

14. Prove Schur's Lemma for matrices. Every $A \in M_n(\mathbb{R})$ is orthogonally equivalent and every $A \in M_n(\mathbb{C})$ is unitarily equivalent to an upper triangular matrix.

2.6 Orthogonal Projections

Recall the discussion of the Gram-Schmidt process, where we saw that any finite-dimensional subspace W of an inner product space V has an orthonormal basis $\beta_W = \{\mathbf{w}_1, \ldots, \mathbf{w}_n\}$. In such a situation, we can define the *orthogonal projections* onto W and $W^\perp$ via

$$\pi_W : V \to W : \mathbf{x} \mapsto \sum_{j=1}^{n} \langle \mathbf{x}, \mathbf{w}_j \rangle \, \mathbf{w}_j, \qquad \pi_W^\perp : V \to W^\perp : \mathbf{x} \mapsto \mathbf{x} - \pi_W(\mathbf{x})$$

Our previous goal was to use orthonormal bases to ease *computation.* In this section we develop projections more generally. First recall the notion of a direct sum within a vector space V:

$$V = X \oplus Y \iff \forall \mathbf{v} \in V, \, \exists \text{ unique } \mathbf{x} \in X, \, \mathbf{y} \in Y \text{ such that } \mathbf{v} = \mathbf{x} + \mathbf{y}$$
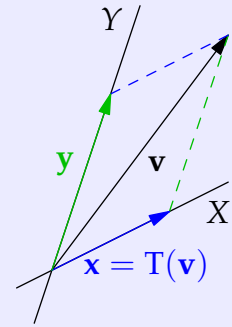
Definition 2.54. A linear map $T \in \mathcal{L}(V)$ is a *projection* if:

$$V = \mathcal{R}(T) \oplus \mathcal{N}(T) \quad \text{and} \quad T_{\mathcal{R}(T)} = I_{\mathcal{R}(T)}$$

Otherwise said, $T(\mathbf{r} + \mathbf{n}) = \mathbf{r}$ whenever $\mathbf{r} \in \mathcal{R}(T)$ and $\mathbf{n} \in \mathcal{N}(T)$.

Alternatively, given $V = X \oplus Y$, the *projection along Y onto X* is the map $\mathbf{v} = \mathbf{x} + \mathbf{y} \mapsto \mathbf{x}$.

We call a $A \in M_n(\mathbb{F})$ a *projection matrix* if $L_A \in \mathcal{L}(\mathbb{F}^n)$ is a projection.

Example 2.55. $A = \frac{1}{5} \begin{pmatrix} 6 & -2 \\ 3 & -1 \end{pmatrix}$ is a projection matrix with $\mathcal{R}(A) = \text{Span}\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\mathcal{N}(A) = \text{Span}\begin{pmatrix} 1 \\ 3 \end{pmatrix}$. Indeed, it is straightforward to describe all projection matrices in $M_2(\mathbb{R})$. There are three cases:

1. $A = I$ is the identity matrix: $\mathcal{R}(A) = \mathbb{R}^2$ and $\mathcal{N}(A) = \{\mathbf{0}\}$;

2. $A = 0$ is the zero matrix: $\mathcal{R}(A) = \{\mathbf{0}\}$ and $\mathcal{N}(A) = \mathbb{R}^2$;

3. Choose *distinct* subspaces $\mathcal{R}(A) = \text{Span}\begin{pmatrix} a \\ b \end{pmatrix}$ and $\mathcal{N}(A) = \text{Span}\begin{pmatrix} c \\ d \end{pmatrix}$, then

$$A = \frac{1}{ad - bc} \begin{pmatrix} a \\ b \end{pmatrix} (d \quad -c) = \frac{1}{ad - bc} \begin{pmatrix} ad & -ac \\ bd & -bc \end{pmatrix}$$

Think about why this last does what we claim.

It should be clear that every projection T has (at most) two eigenspaces:

- $\mathcal{R}(T)$ is an eigenspace with eigenvalue 1

- $\mathcal{N}(T)$ is an eigenspace with eigenvalue 0

If V is finite-dimensional and ρ, η are bases of $\mathcal{R}(T)$, $\mathcal{N}(T)$ respectively, then the matrix of T with respect to $\rho \cup \eta$ has block form

$$[T]_{\rho \cup \eta} = \left(\begin{array}{c|c} I & 0 \\ \hline 0 & 0 \end{array} \right)$$

where $\text{rank } I = \text{rank } T$. In particular, every finite-dimensional projection is diagonalizable.

Lemma 2.56. $T \in \mathcal{L}(V)$ *is a projection if and only if* $T^2 = T$.

Proof. Throughout, assume $\mathbf{r} \in \mathcal{R}(T)$ and $\mathbf{n} \in \mathcal{N}(T)$.

($\Rightarrow$) Since every vector in V has a unique representation $\mathbf{v} = \mathbf{r} + \mathbf{n}$, simply compute

$$T^2(\mathbf{v}) = T(T(\mathbf{r}+\mathbf{n})) = T(\mathbf{r}) = \mathbf{r} = T(\mathbf{v})$$

($\Leftarrow$) Suppose $T^2 = T$. Note first that if $\mathbf{r} \in \mathcal{R}(T)$, then $\mathbf{r} = T(\mathbf{v})$ for some $\mathbf{v} \in V$, whence

$$T(\mathbf{r}) = T^2(\mathbf{v}) = T(\mathbf{v}) = \mathbf{r} \tag{†}$$

Thus T is the identity on $\mathcal{R}(T)$. Moreover, if $\mathbf{x} \in \mathcal{R}(T) \cap \mathcal{N}(T)$, (†) says that $\mathbf{x} = T(\mathbf{x}) = \mathbf{0}$, whence

$$\mathcal{R}(T) \cap \mathcal{N}(T) = \{\mathbf{0}\}$$

and so $\mathcal{R}(T) \oplus \mathcal{N}(T)$ is a well-defined *subspace* of V.[12]To finish things off, let $\mathbf{v} \in V$ and observe that

$$T(\mathbf{v} - T(\mathbf{v})) = T(\mathbf{v}) - T^2(\mathbf{v}) = \mathbf{0} \implies \mathbf{v} - T(\mathbf{v}) \in \mathcal{N}(T)$$

so that $\mathbf{v} = T(\mathbf{v}) + (\mathbf{v} - T(\mathbf{v}))$ is a decomposition into $\mathcal{R}(T)$- and $\mathcal{N}(T)$-parts. We conclude that $V = \mathcal{R}(T) \oplus \mathcal{N}(T)$ and that T is a projection. ∎

Thus far the discussion hasn't had anything to do with inner products...

Definition 2.57. An *orthogonal projection* is a projection $T \in \mathcal{L}(V)$ on an inner product space for which we additionally have

$$\mathcal{N}(T) = \mathcal{R}(T)^\perp \quad \text{and} \quad \mathcal{R}(T) = \mathcal{N}(T)^\perp$$

Alternatively, given $V = W \oplus W^\perp$, the *orthogonal projection* π_W is the projection along $W^\perp$ onto W: that is

$$\mathcal{R}(\pi_W) = W \quad \text{and} \quad \mathcal{N}(\pi_W) = W^\perp$$

The complementary orthogonal projection $\pi_W^\perp = I - \pi_W$ has $\mathcal{R}(\pi_W^\perp) = W^\perp$ and $\mathcal{N}(\pi_W^\perp) = W$.

Example (2.55 continued). The identity and zero matrices are both 2×2 orthogonal projection matrices, while those of type 3 are orthogonal if $\left(\begin{smallmatrix} a \\ b \end{smallmatrix}\right) \cdot \left(\begin{smallmatrix} c \\ d \end{smallmatrix}\right) = 0$: we obtain

$$A = \frac{1}{a^2 + b^2} \begin{pmatrix} a \\ b \end{pmatrix} (a \ b) = \frac{1}{a^2 + b^2} \begin{pmatrix} a^2 & ab \\ ab & b^2 \end{pmatrix}$$

More generally, if $W \leq \mathbb{F}^n$ has orthonormal basis $\{\mathbf{w}_1, \dots, \mathbf{w}_k\}$, then the matrix of π_W is $\sum_{j=1}^k \mathbf{w}_j \mathbf{w}_j^*$.

---

[12]In finite dimensions, the rank–nullity theorem and dimension counting finishes the proof here without having to proceed further:

$$\dim(\mathcal{R}(T) \oplus \mathcal{N}(T)) = \operatorname{rank} T + \operatorname{null} T = \dim V \implies \mathcal{R}(T) \oplus \mathcal{N}(T) = V$$

Theorem 2.58. *A projection* $T \in \mathcal{L}(V)$ *is orthogonal if and only if it is self-adjoint* $T = T^*$.

Proof. ($\Rightarrow$) By assumption, $\mathcal{R}(T)$ and $\mathcal{N}(T)$ are orthogonal subspaces. Letting $\mathbf{x}, \mathbf{y} \in V$ and using subscripts to denote $\mathcal{R}(T)$- and $\mathcal{N}(T)$-parts, we see that

$$\langle \mathbf{x}, T(\mathbf{y}) \rangle = \langle \mathbf{x}_r + \mathbf{x}_n, \mathbf{y}_r \rangle = \langle \mathbf{x}_r, \mathbf{y}_r + \mathbf{y}_n \rangle = \langle T(\mathbf{x}), \mathbf{y} \rangle \implies T^* = T$$

($\Leftarrow$) Suppose T is a self-adjoint projection. By the fundamental subspaces theorem,

$$\mathcal{N}(T) = \mathcal{N}(T^*) = \mathcal{R}(T)^\perp$$

Since T is a projection already, we have $V = \mathcal{R}(T) \oplus \mathcal{N}(T) = \mathcal{R}(T) \oplus \mathcal{R}(T)^\perp$, from which

$$\mathcal{R}(T) = \left(\mathcal{R}(T)^\perp \right)^\perp = \mathcal{N}(T)^\perp \qquad \blacksquare$$

The language of projections allows us to rephrase the Spectral Theorem.

Theorem 2.59 (Spectral Theorem, mk. II). *Let V be a finite-dimensional complex/real inner product space and $T \in \mathcal{L}(V)$ be normal/self-adjoint with spectrum $\{\lambda_1, \ldots, \lambda_k\}$ and corresponding eigenspaces $E_1, \ldots, E_k$. Let $\pi_j \in \mathcal{L}(V)$ be the orthogonal projection onto E_j. Then:*

1. *$V = E_1 \oplus \cdots \oplus E_k$ is a direct sum of orthogonal subspaces, in particular, $E_j^\perp$ is the direct sum of the remaining eigenspaces.*

2. *$\pi_i \pi_j = 0$ if $i \neq j$.*

3. *(Resolution of the identity)* *$I_V = \pi_1 + \cdots + \pi_k$*

4. *(Spectral decomposition)* *$T = \lambda_1 \pi_1 + \cdots + \lambda_k \pi_k$*

Proof. 1. T is diagonalizable and so V is the direct sum of the eigenspaces of T. Since T is normal, the eigenvectors corresponding to distinct eigenvalues are orthogonal, whence the eigenspaces are mutually orthogonal. In particular, this says that

$$\hat{E}_j := \bigoplus_{i \neq j} E_i \leq E_j^\perp$$

Since V is finite-dimensional, we have $V = E_j \oplus E_j^\perp$, whence

$$\dim \hat{E}_j = \sum_{i \neq j} \dim E_i = \dim V - \dim E_j = \dim E_j^\perp \implies \hat{E}_j = E_j^\perp$$

2. This is clear by part 1, since $\mathcal{N}(\pi_j) = E_j^\perp = \hat{E}_j$.

3. Write $\mathbf{x} = \sum_{j=1}^k \mathbf{x}_j$ where each $\mathbf{x}_j \in E_j$. Then $\pi_j(\mathbf{x}) = \mathbf{x}_j$: now add…

4. $T(\mathbf{x}) = \sum_{j=1}^k T(\mathbf{x}_j) = \sum_{j=1}^k \lambda_j \mathbf{x}_j = \sum_{j=1}^k \lambda_j \pi_j(\mathbf{x})$. $\qquad \blacksquare$

Examples 2.60. We verify the resolution of the identity and the spectral decomposition; for clarity, we index projections and eigenspaces by eigenvalue rather than the natural numbers.

1. The symmetric matrix $A = \left(\begin{smallmatrix} 10 & 2 \\ 2 & 7 \end{smallmatrix}\right)$ has spectrum $\{6, 11\}$ and orthonormal eigenvectors

$$\mathbf{w}_6 = \frac{1}{\sqrt{5}}\begin{pmatrix} 1 \\ -2 \end{pmatrix}, \qquad \mathbf{w}_{11} = \frac{1}{\sqrt{5}}\begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

The corresponding projections therefore have matrices

$$\pi_6 = \mathbf{w}_6\mathbf{w}_6^T = \frac{1}{5}\begin{pmatrix} 1 \\ -2 \end{pmatrix}(1 \ \ -2) = \frac{1}{5}\begin{pmatrix} 1 & -2 \\ -2 & 4 \end{pmatrix}, \qquad \pi_{11} = \mathbf{w}_{11}\mathbf{w}_{11}^T = \frac{1}{5}\begin{pmatrix} 4 & 2 \\ 2 & 1 \end{pmatrix}$$

from which the resolution of the identity and the spectral decomposition are readily verified:

$$\pi_6 + \pi_{11} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad 6\pi_6 + 11\pi_{11} = \frac{1}{5}\begin{pmatrix} 6+44 & -12+22 \\ -12+22 & 24+11 \end{pmatrix} = A$$

2. The normal matrix $B = \left(\begin{smallmatrix} 1+i & 1+i \\ -1-i & 1+i \end{smallmatrix}\right) \in M_2(\mathbb{C})$ has spectrum $\{2, 2i\}$ and corresponding orthonormal eigenvectors

$$\mathbf{w}_2 = \frac{1}{\sqrt{2}}\begin{pmatrix} i \\ 1 \end{pmatrix}, \qquad \mathbf{w}_{2i} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ i \end{pmatrix}$$

The orthogonal projection matrices are therefore

$$\pi_2 = \mathbf{w}_2\mathbf{w}_2^* = \frac{1}{2}\begin{pmatrix} i \\ 1 \end{pmatrix}(-i \ \ 1) = \frac{1}{2}\begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix} \qquad \pi_{2i} = \mathbf{w}_{2i}\mathbf{w}_{2i}^* = \frac{1}{2}\begin{pmatrix} 1 & -i \\ i & 1 \end{pmatrix}$$

from which

$$\pi_2 + \pi_{2i} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad 2\pi_2 + 2i\pi_{2i} = \begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix} + \begin{pmatrix} i & 1 \\ -1 & i \end{pmatrix} = B$$

3. The matrix $C = \left(\begin{smallmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{smallmatrix}\right)$ has spectrum $\{-1, 2\}$, an orthonormal eigenbasis

$$\{\mathbf{u}, \mathbf{v}, \mathbf{w}\} = \left\{ \frac{1}{\sqrt{3}}\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \frac{1}{\sqrt{6}}\begin{pmatrix} 1 \\ 1 \\ -2 \end{pmatrix} \right\}$$

and eigenspaces $E_2 = \text{Span}\{\mathbf{u}\}$ and $E_{-1} = \text{Span}\{\mathbf{v}, \mathbf{w}\}$. The orthogonal projections have matrices

$$\pi_2 = \mathbf{u}\mathbf{u}^T = \frac{1}{3}\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}(1 \ 1 \ 1) = \frac{1}{3}\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

$$\pi_{-1} = \mathbf{v}\mathbf{v}^T + \mathbf{w}\mathbf{w}^T = \frac{1}{2}\begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}(1 \ \ -1 \ 0) + \frac{1}{6}\begin{pmatrix} 1 \\ 1 \\ -2 \end{pmatrix}(1 \ 1 \ \ -2)$$

$$= \frac{1}{2}\begin{pmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} + \frac{1}{6}\begin{pmatrix} 1 & 1 & -2 \\ 1 & 1 & -2 \\ -2 & -2 & 4 \end{pmatrix} = \frac{1}{3}\begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}$$

It is now easy to check the resolution of the identity and the spectral decomposition:

$$\pi_2 + \pi_{-1} = I \quad \text{and} \quad 2\pi_2 - \pi_{-1} = C$$

Orthogonal Projections and Minimization Problems

We finish this section with an important observation that drives much of the application of inner product spaces to other parts of mathematics and beyond. Throughout this discussion, X and Y denote inner product spaces.
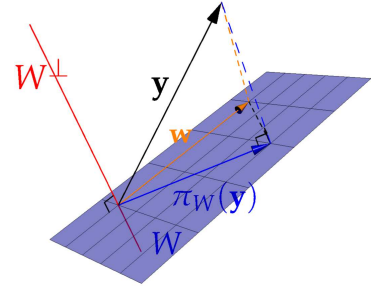
> **Theorem 2.61.** *Suppose $Y = W \oplus W^\perp$. For any $\mathbf{y} \in Y$, the orthogonal projection $\pi_W(\mathbf{y})$ is the unique element of W which minimizes the distance to $\mathbf{y}$:*
>
> $$\forall \mathbf{w} \in W, \ ||\mathbf{y} - \pi_W(\mathbf{y})|| \le ||\mathbf{y} - \mathbf{w}||$$

Proof. Apply Pythagoras' Theorem: since $\pi_W^\perp(\mathbf{y}) = \mathbf{y} - \pi_W(\mathbf{y}) \in W^\perp$,

$$||\mathbf{y} - \mathbf{w}||^2 = ||\mathbf{y} - \pi_W(\mathbf{y}) + \pi_W(\mathbf{y}) - \mathbf{w}||^2$$
$$= ||\mathbf{y} - \pi_W(\mathbf{y})||^2 + ||\pi_W(\mathbf{y}) - \mathbf{w}||^2$$
$$\ge ||\mathbf{y} - \pi_W(\mathbf{y})||^2$$

with equality if and only if $\mathbf{w} = \pi_W(\mathbf{y})$. ∎

This set up can be used to compute accurate approximations in many contexts.

Examples 2.62.   1. To obtain a quadratic polynomial approximation $p(x) = a + bx + cx^2$ to e^x on the interval $[-1, 1]$ we choose to minimize the integral $\int_{-1}^{1} |e^x - p(x)|^2 \, dx$, namely the squared L^2-norm $||e^x - p(x)||^2$ on $C[-1, 1]$. If we let $W = \mathrm{Span}\{1, x, x^2\}$, then the finite-dimensionality of W means that $C[-1, 1] = W \oplus W^\perp$. By the Theorem, the solution is $p(x) = \pi_W(e^x)$.

To compute this, recall that we have an orthonormal basis for W, namely

$$\left\{ \frac{1}{\sqrt{2}}, \sqrt{\frac{3}{2}}x, \sqrt{\frac{5}{8}}(3x^2 - 1) \right\}$$

from which

$$p(x) = \frac{1}{2}\langle 1, e^x \rangle + \frac{3}{2}\langle x, e^x \rangle x + \frac{5}{8}(3x^2 - 1)\langle 3x^2 - 1, e^x \rangle$$
$$= \frac{1}{2}\int_{-1}^{1} e^x \, dx + \frac{3}{2}x \int_{-1}^{1} xe^x \, dx$$
$$\quad + \frac{5}{8}(3x^2 - 1)\int_{-1}^{1}(3x^2 - 1)e^x \, dx$$
$$= \frac{1}{2}(e - e^{-1}) + 3e^{-1}x + \frac{5}{4}(e - 7e^{-1})(3x^2 - 1)$$
$$\approx 1.18 + 1.10x + 0.179(3x^2 - 1)$$
$$\approx 1 + 1.1x + 0.537x^2$$
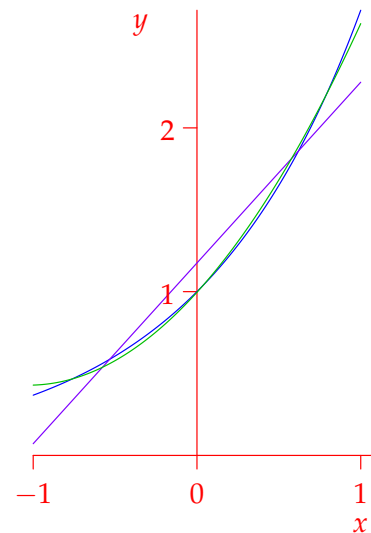
The linear and quadratic approximations to $y = e^x$ are drawn. Compare this with the Maclaurin polynomial $e^x \approx 1 + x + \frac{1}{2}x^2$ from calculus.

2. The n^{th} Fourier approximation of a function $f(x)$ is its orthogonal projection onto the finite-dimensional subspace

$$W_n = \text{Span}\{1, e^{ix}, e^{-ix}, \ldots, e^{inx}, e^{-inx}\}$$
$$= \text{Span}\{1, \cos x, \sin x, \cos 2x, \sin 2x, \ldots, \cos nx, \sin nx\}$$

within $L^2[-\pi, \pi]$, namely

$$\mathcal{F}_n(x) = \frac{1}{2\pi} \sum_{k=-n}^{n} \left\langle f(x), e^{ikx} \right\rangle e^{ikx}$$
$$= \frac{1}{2\pi} \langle f(x), 1 \rangle + \frac{1}{\pi} \sum_{k=1}^{n} \langle f(x), \cos kx \rangle \cos kx + \langle f(x), \sin kx \rangle \sin kx$$

According to the Theorem, this is the unique function in $\mathcal{F}_n(x) \in W_n$ minimizing the integral

$$||f(x) - \mathcal{F}_n(x)||^2 = \int_{-\pi}^{\pi} |f(x) - \mathcal{F}_n(x)|^2 \, dx$$

For example, if $f(x) = \begin{cases} 1 & \text{if } 0 < x \leq \pi \\ -1 & \text{if } -\pi < x \leq 0 \end{cases}$ is extended periodically, then

$$\mathcal{F}_{2n-1}(x) = \frac{4}{\pi} \sum_{j=1}^{n} \frac{\sin(2j-1)x}{2j-1} = \frac{4}{\pi}\left(\sin x + \frac{\sin 3x}{3} + \frac{\sin 5x}{5} + \cdots + \frac{\sin(2n-1)x}{2n-1}\right)$$

$y = f(x)$ and its eleventh Fourier approximation $y = \mathcal{F}_{11}(x)$
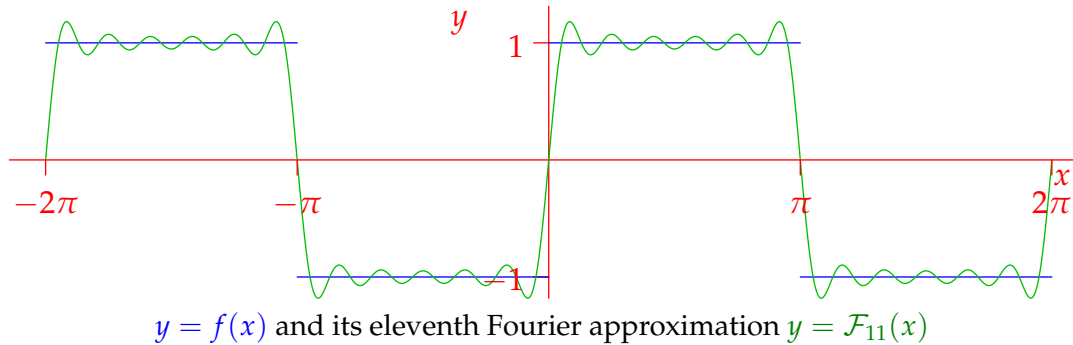
We'll return to related minimization problems in the next section.

Exercises 2.6   1. Compute the matrices of the orthogonal projections onto W viewed as subspaces of the standard inner product spaces $\mathbb{R}^n$ or $\mathbb{C}^n$.

(a) $W = \text{Span}\begin{pmatrix} 4 \\ -1 \end{pmatrix}$ (b) $W = \text{Span}\left\{\begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}\right\}$ (c) $W = \text{Span}\left\{\begin{pmatrix} i \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ i \\ 1 \end{pmatrix}\right\}$

(d) $W = \text{Span}\left\{\begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}\right\}$ *(watch out, these vectors aren't orthogonal!)*

2. For each of the following matrices, compute the projections onto each eigenspace, verify the resolution of the identity and the spectral decomposition.

(a) $\begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}$ (b) $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ (c) $\begin{pmatrix} 2 & 3-3i \\ 3+3i & 5 \end{pmatrix}$ (d) $\begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}$

   (You should have orthonormal eigenbases from the previous section)

3. If W be a finite-dimensional subspace of an inner product space V. If $T = \pi_W$ is the orthogonal projection onto W, prove that $I - T$ is the orthogonal projection onto $W^\perp$.

4. Let $T \in \mathcal{L}(V)$ where V is finite-dimensional.

   (a) If T is an orthogonal projection, prove that $||T(\mathbf{x})|| \leq ||\mathbf{x}||$ for all $\mathbf{x} \in V$.

   (b) Give an example of a projection for which the inequality in (a) is *false.*

   (c) If T is a projection for which $||T(\mathbf{x})|| = ||\mathbf{x}||$ for all $\mathbf{x} \in V$, what is T?

   (d) If T is a projection for which $||T(\mathbf{x})|| \leq ||\mathbf{x}||$ for all $\mathbf{x} \in V$, prove that T is an *orthogonal* projection.

5. Let T be a normal operator on a finite-dimensional inner product space. If T is a projection, prove that it must be an orthogonal projection.

6. Let T be a normal operator on a finite-dimensional complex inner product space V. Use the spectral decomposition $T = \lambda_1 \pi_1 + \cdots + \lambda_k \pi_k$ to prove:

   (a) If T^n is the zero map for some $n \in \mathbb{N}$, then T is the zero map.

   (b) $U \in \mathcal{L}(V)$ commutes with T if and only if U commutes with each π_j.

   (c) There exists a normal $U \in \mathcal{L}(V)$ such that $U^2 = T$.

   (d) T is invertible if and only if $\lambda_j \neq 0$ for all j.

   (e) T is a projection if and only if every $\lambda_j = 0$ or 1.

   (f) $T = -T^*$ if and only if every λ_j is imaginary.

7. Find a linear approximation to $f(x) = e^x$ on $[0, 1]$ using the L^2 inner product.

8. Consider the L^2 inner product on $C[-\pi, \pi]$ inner product.

   (a) Explain why $\langle \sin x, x^{2n} \rangle = 0$ for all n.

   (b) Find linear and *cubic* approximations to $f(x) = \sin x$.

   (Feel free to use a computer algebra package to evaluate the integrals!)

9. Revisit Example 2.62.2

   (a) Verify that the general complex (e^{ikx}) and real $(\cos kx, \sin kx)$ expressions for the Fourier approximation are correct.

   (Hint: use Euler's formula $e^{ikx} = \cos kx + i \sin kx$)

   (b) Verify the explicit expression for $\mathcal{F}_{2n-1}(x)$ when $f(x)$ is the given step-function. What is $\mathcal{F}_{2n}(x)$ in this case?

2.7 The Singular Value Decomposition and the Pseudoinverse

Given $T \in \mathcal{L}(V, W)$ between finite-dimensional inner product spaces, the overarching concern of this chapter is the existence and computation of bases β, γ of V, W with two properties:

- That β, γ be *orthonormal,* thus facilitating easy calculation within V, W;

- That the matrix $[T]_\beta^\gamma$ be as simple as possible.

We have already addressed two special cases:

Spectral Theorem When $V = W$ and T is normal/self-adjoint, $\exists \beta = \gamma$ such that $[T]_\beta$ is diagonal.

Schur's Lemma When $V = W$ and $p(t)$ splits, $\exists \beta = \gamma$ such that $[T]_\beta$ is upper triangular.

In this section we allow $V \neq W$ and $\beta \neq \gamma$, and obtain a result that applies to any linear map between finite-dimensional inner product spaces.

Example 2.63. Let $A = \begin{pmatrix} 3 & 1 \\ 2 & -2 \\ 1 & 3 \end{pmatrix}$ and consider orthonormal bases $\beta = \{\mathbf{v}_1, \mathbf{v}_2\}$ of $\mathbb{R}^2$ and $\gamma = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3\}$ of $\mathbb{R}^3$ respectively:

$$\beta = \left\{ \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix} \right\}, \quad \gamma = \left\{ \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{6}} \begin{pmatrix} -1 \\ -2 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ -1 \\ -1 \end{pmatrix} \right\}$$

Since $A\mathbf{v}_1 = 4\mathbf{w}_1$ and $A\mathbf{v}_2 = 2\sqrt{3}\mathbf{w}_2$, whence $[L_A]_\beta^\gamma = \begin{pmatrix} 4 & 0 \\ 0 & 2\sqrt{3} \\ 0 & 0 \end{pmatrix}$ is almost diagonal.

Our main result says that such bases always exist.

Theorem 2.64 (Singular Value Decomposition). *Suppose V, W are finite-dimensional inner product spaces and that $T \in \mathcal{L}(V, W)$ has rank r. Then:*

1. *There exist orthonormal bases $\beta = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ of V and $\gamma = \{\mathbf{w}_1, \ldots, \mathbf{w}_m\}$ of W, and positive scalars $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r$ such that*

$$T(\mathbf{v}_j) = \begin{cases} \sigma_j \mathbf{w}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases} \quad \text{equivalently} \quad [T]_\beta^\gamma = \begin{pmatrix} \text{diag}(\sigma_1, \ldots, \sigma_r) & O \\ O & O \end{pmatrix}$$

2. *Any such β is an eigenbasis of T^*T, whence the scalars σ_j are uniquely determined by T: indeed*

$$T^*T(\mathbf{v}_j) = \begin{cases} \sigma_j^2 \mathbf{v}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases} \quad \text{and} \quad T^*(\mathbf{w}_j) = \begin{cases} \sigma_j \mathbf{v}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases}$$

3. *If $A \in M_{m \times n}(\mathbb{F})$ with rank $A = r$, then $A = P\Sigma Q^*$ where*

$$\Sigma = [L_A]_\beta^\gamma = \begin{pmatrix} \text{diag}(\sigma_1, \ldots, \sigma_r) & O \\ O & O \end{pmatrix}, \quad P = (\mathbf{w}_1, \ldots, \mathbf{w}_m), \quad Q = (\mathbf{v}_1, \ldots, \mathbf{v}_n)$$

Since the columns of P, Q are orthonormal, these matrices are unitary.

> **Definition 2.65.** The numbers $\sigma_1, \ldots, \sigma_r$ are the *singular values* of T. If T is not maximum rank, we have additional zero singular values $\sigma_{r+1} = \cdots = \sigma_{\min(m,n)} = 0$.

Freedom of Choice   While the singular values are uniquely determined, there is often significant freedom regarding the bases β and γ, particularly if any eigenspace of T^*T has dimension ≥ 2.

Special Case (Spectral Theorem)   If $V = W$ and T is normal/self-adjoint, we may *choose* β to be an eigenbasis of T, then σ_j is the *modulus* of the corresponding eigenvalue (see Exercise 7).

Rank-one decomposition   If we write $g_j : V \to \mathbb{F}$ for the linear map $g_j : \mathbf{v} \mapsto \langle \mathbf{v}, \mathbf{v}_j \rangle$ (recall Riesz's Theorem), then the singular value decomposition says

$$T = \sum_{j=1}^{r} \sigma_j \mathbf{w}_j g_j \quad \text{that is} \quad T(\mathbf{v}) = \sum_{j=1}^{r} \sigma_j \langle \mathbf{v}, \mathbf{v}_j \rangle \mathbf{w}_j$$

thus rewriting T as a linear combination of rank-one maps ($\mathbf{w}_j g_j : V \to W$). For matrices, g_j is simply multiplication by the row vector $\mathbf{v}_j^*$ and we may write[13]

$$A = \sum_{j=1}^{r} \sigma_j \mathbf{w}_j \mathbf{v}_j^*$$

> **Example (2.63 cont).**   We apply the method in the Theorem to $A = \begin{pmatrix} 3 & 1 \\ 2 & -2 \\ 1 & 3 \end{pmatrix}$.

The symmetric(!) matrix $A^T A = \begin{pmatrix} 14 & 2 \\ 2 & 14 \end{pmatrix}$ has eigenvalues $\sigma_1^2 = 16, \sigma_2^2 = 12$ and orthonormal eigenvectors $\mathbf{v}_1 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \mathbf{v}_2 = \frac{1}{\sqrt{2}} \begin{pmatrix} -1 \\ 1 \end{pmatrix}$. The singular values are therefore $\sigma_1 = 4, \sigma_2 = 2\sqrt{3}$. Now compute

$$\mathbf{w}_1 = \frac{1}{\sigma_1} A\mathbf{v}_1 = \frac{1}{4\sqrt{2}} A \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad \mathbf{w}_2 = \frac{1}{\sigma_2} A\mathbf{v}_2 = \frac{1}{2\sqrt{6}} A \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{6}} \begin{pmatrix} -1 \\ -2 \\ 1 \end{pmatrix}$$

and observe that these are orthonormal. Finally choose $\mathbf{w}_3 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ -1 \\ -1 \end{pmatrix}$ to complete the orthonormal basis γ of $\mathbb{R}^3$. A singular value decomposition is therefore

$$A = \begin{pmatrix} 3 & 1 \\ 2 & -2 \\ 1 & 3 \end{pmatrix} = P\Sigma Q^* = \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{-1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \\ 0 & \frac{-2}{\sqrt{6}} & \frac{-1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & \frac{-1}{\sqrt{3}} \end{pmatrix} \begin{pmatrix} 4 & 0 \\ 0 & 2\sqrt{3} \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix}$$

By expanding the decomposition, A is expressed as the sum of rank-one matrices:

$$\sigma_1 \mathbf{w}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{w}_2 \mathbf{v}_2^T = 4 \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} + 2\sqrt{3} \begin{pmatrix} \frac{-1}{\sqrt{6}} \\ \frac{-2}{\sqrt{6}} \\ \frac{1}{\sqrt{6}} \end{pmatrix} \begin{pmatrix} \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} 2 & 2 \\ 0 & 0 \\ 2 & 2 \end{pmatrix} + \begin{pmatrix} 1 & -1 \\ 2 & -2 \\ -1 & 1 \end{pmatrix}$$

---

[13]Since β is orthonormal, it is common to write $\mathbf{v}_j^*$ for the map $g_j = \langle \ , \mathbf{v}_j \rangle$ in general contexts. To those familiar with the *dual space* $V^* = \mathcal{L}(V, \mathbb{F})$, the set $\{g_1, \ldots, g_n\} = \{\mathbf{v}_1^*, \ldots, \mathbf{v}_n^*\}$ is the *dual basis* to β. In this course $\mathbf{v}_j^*$ will only ever mean the *conjugate-transpose* of a *column vector* in $\mathbb{F}^n$. This discussion is part of why physicists write inner products differently!

Proof.   1. Since T^*T is self-adjoint, the spectral theorem says it has an orthonormal basis of eigenvectors $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$. If $T^*T(\mathbf{v}_j) = \lambda_j \mathbf{v}_j$, then

$$\langle T(\mathbf{v}_j), T(\mathbf{v}_k) \rangle = \langle T^*T(\mathbf{v}_j), \mathbf{v}_k \rangle = \lambda_j \langle \mathbf{v}_j, \mathbf{v}_k \rangle = \begin{cases} \lambda_j & \text{if } j = k \\ 0 & \text{if } j \neq k \end{cases} \tag{$*$}$$

whence every eigenvalue is a non-negative real number: $\lambda_j = ||T(\mathbf{v}_j)||^2 \geq 0$.
Since $\operatorname{rank} T^*T = \operatorname{rank} T = r$ (Exercise 8), exactly r eigenvalues are non-zero; by reordering basis vectors if necessary, we may assume

$$\lambda_1 \geq \cdots \geq \lambda_r > 0$$

If $j \leq r$, define $\sigma_j := \sqrt{\lambda_j} > 0$ and $\mathbf{w}_j := \frac{1}{\sigma_j} T(\mathbf{v}_j)$, then the set $\{\mathbf{w}_1, \dots, \mathbf{w}_r\}$ is orthonormal $(*)$. If necessary, extend this to an orthonormal basis γ of W.

2. If orthonormal bases β and γ exist such that $[T]_\beta^\gamma = \begin{pmatrix} \operatorname{diag}(\sigma_1,\dots,\sigma_r) & O \\ O & O \end{pmatrix}$, then $[T^*]_\gamma^\beta$ is *essentially the same matrix* just that its dimensions have been reversed:

$$[T^*]_\gamma^\beta = \begin{pmatrix} \operatorname{diag}(\sigma_1,\dots,\sigma_r) & O \\ O & O \end{pmatrix} \implies T^*T = \begin{pmatrix} \operatorname{diag}(\sigma_1^2,\dots,\sigma_r^2) & O \\ O & O \end{pmatrix}$$

whence T^* and T^*T are as claimed.

3. This is merely part 1 in the context of $T = L_A \in \mathcal{L}(\mathbb{F}^n, \mathbb{F}^m)$. The orthonormal bases β, γ consist of *column vectors* and so the (change of co-ordinate) matrices P, Q are unitary. ∎

Examples 2.66.   1. The matrix $A = \begin{pmatrix} 2 & 3 \\ 0 & 2 \end{pmatrix}$ has $A^T A = \begin{pmatrix} 4 & 6 \\ 6 & 13 \end{pmatrix}$ with eigenvalues $\sigma_1^2 = 16$ and $\sigma_2^2 = 1$ and orthonormal eigenbasis

$$\beta = \left\{ \frac{1}{\sqrt{5}} \begin{pmatrix} 1 \\ 2 \end{pmatrix}, \frac{1}{\sqrt{5}} \begin{pmatrix} -2 \\ 1 \end{pmatrix} \right\}$$

The singular values are therefore $\sigma_1 = 4$ and $\sigma_2 = 1$, from which we obtain

$$\gamma = \left\{ \frac{1}{\sigma_1} A\mathbf{v}_1, \frac{1}{\sigma_2} A\mathbf{v}_2 \right\} = \left\{ \frac{1}{\sqrt{5}} \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{5}} \begin{pmatrix} -1 \\ 2 \end{pmatrix} \right\}$$

and the decomposition

$$A = P\Sigma Q^* = \begin{pmatrix} \frac{2}{\sqrt{5}} & \frac{-1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \end{pmatrix} \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{-2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{pmatrix}$$

Multiplying out, we may write A as a sum of rank-one matrices

$$A = \frac{4}{5} \begin{pmatrix} 2 \\ 1 \end{pmatrix} \begin{pmatrix} 1 & 2 \end{pmatrix} + \frac{1}{5} \begin{pmatrix} -1 \\ 2 \end{pmatrix} \begin{pmatrix} -2 & 1 \end{pmatrix} = \frac{4}{5} \begin{pmatrix} 2 & 4 \\ 1 & 2 \end{pmatrix} + \frac{1}{5} \begin{pmatrix} 2 & -1 \\ -4 & 2 \end{pmatrix}$$

2. The decomposition can be very messy to find in non-matrix situations. Here is a classic example where we simply observe the structure directly.

The L^2 inner product $\langle f, g \rangle = \int_0^1 f(x)g(x)\,\mathrm{d}x$ on $P_2(\mathbb{R})$ and $P_1(\mathbb{R})$ admits orthonormal bases

$$\beta = \left\{ \sqrt{5}(6x^2 - 6x + 1), \sqrt{3}(2x - 1), 1 \right\}, \qquad \gamma = \left\{ \sqrt{3}(2x - 1), 1 \right\}$$

Let $T = \frac{\mathrm{d}}{\mathrm{d}x}$ be the derivative operator. The matrix of T is already in the required form!

$$[T]_\beta^\gamma = \begin{pmatrix} 2\sqrt{15} & 0 & 0 \\ 0 & 2\sqrt{3} & 0 \end{pmatrix}$$

thus β, γ are suitable bases and the singular values of T are $\sigma_1 = 2\sqrt{15}$ and $\sigma_2 = 2\sqrt{3}$.

Since β, γ are orthonormal, we could have used the adjoint method to evaluate this directly,

$$[T^*T]_\beta = ([T]_\beta^\gamma)^T [T]_\beta^\gamma = \begin{pmatrix} 60 & 0 & 0 \\ 0 & 12 & 0 \\ 0 & 0 & 0 \end{pmatrix} \implies \sigma_1^2 = 60, \ \sigma_2^2 = 12$$

Up to sign, $\{[\mathbf{v}_1]_\beta, [\mathbf{v}_2]_\beta, [\mathbf{v}_3]_\beta\}$ is therefore forced to be the standard ordered basis of $\mathbb{R}^3$, confirming that β was the correct basis of $P_2(\mathbb{R})$ all along!

The Pseudoinverse

The singular value decomposition of a map $T \in \mathcal{L}(V, W)$ gives rise to a natural map from W back to V. This map behaves somewhat like an inverse even when the operator is non-invertible!

Definition 2.67. Given the singular value decomposition of a rank r map $T \in \mathcal{L}(V, W)$, the *pseudoinverse* of T is the linear map $T^\dagger \in \mathcal{L}(W, V)$ defined by

$$T^\dagger(\mathbf{w}_j) = \begin{cases} \frac{1}{\sigma_j}\mathbf{v}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases}$$

Restricted to $\mathrm{Span}\{\mathbf{v}_1, \ldots, \mathbf{v}_r\} \to \mathrm{Span}\{\mathbf{w}_1, \ldots, \mathbf{w}_r\}$, the pseudoinverse really does invert T:

$$T^\dagger T(\mathbf{v}_j) = \begin{cases} \mathbf{v}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases}$$

$$V = \underbrace{\mathrm{Span}\{\mathbf{v}_1, \ldots, \mathbf{v}_r\}}_{\mathcal{N}(T)^\perp = \mathcal{R}(T^\dagger)} \oplus \underbrace{\mathrm{Span}\{\mathbf{v}_{r+1}, \ldots, \mathbf{v}_n\}}_{\mathcal{N}(T) = \mathcal{R}(T^\dagger)^\perp} \quad \{\mathbf{0}_V\}$$

$$T \Big\updownarrow T^\dagger \qquad \text{bijection}$$

$$TT^\dagger(\mathbf{w}_j) = \begin{cases} \mathbf{w}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases}$$

$$W = \underbrace{\mathrm{Span}\{\mathbf{w}_1, \ldots, \mathbf{w}_r\}}_{\mathcal{R}(T) = \mathcal{N}(T^\dagger)^\perp} \oplus \underbrace{\mathrm{Span}\{\mathbf{w}_{r+1}, \ldots, \mathbf{w}_m\}}_{\mathcal{R}(T)^\perp = \mathcal{N}(T^\dagger)} \quad \{\mathbf{0}_W\}$$
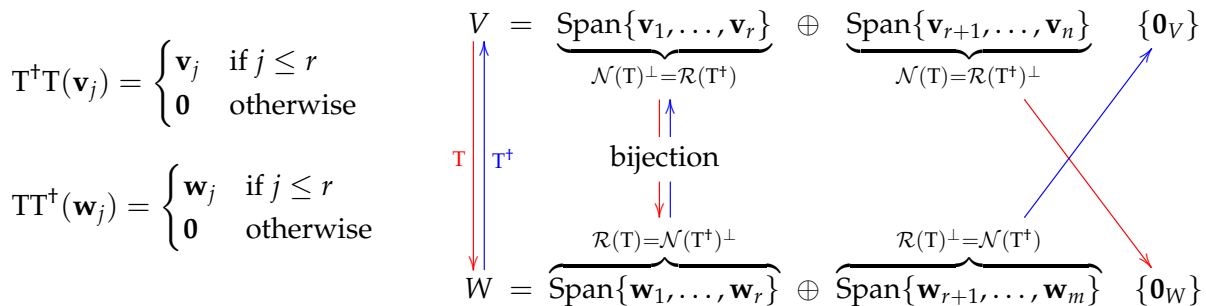
Otherwise said, the combinations are *orthogonal projections*: $T^\dagger T = \pi_{\mathcal{N}(T)}^\perp$ and $TT^\dagger = \pi_{\mathcal{R}(T)}$

Given the singular value decomposition of a matrix $A = P\Sigma Q^*$, its pseudoinverse is that of matrix of $(L_A)^\dagger$, namely

$$A^\dagger = \sum_{j=1}^{r} \frac{1}{\sigma_j} \mathbf{v}_j \mathbf{w}_j^* = Q\Sigma^\dagger P^* \quad \text{where} \quad \Sigma^\dagger = \begin{pmatrix} \text{diag}(\sigma_1^{-1}, \ldots, \sigma_r^{-1}) & O \\ O & O \end{pmatrix}$$

Examples 2.68. 1. Again continuing Example 2.63, $A = \begin{pmatrix} 3 & 1 \\ 2 & -2 \\ 1 & 3 \end{pmatrix}$ has pseudoinverse

$$A^\dagger = \frac{1}{\sigma_1} \mathbf{v}_1 \mathbf{w}_1^* + \frac{1}{\sigma_2} \mathbf{v}_2 \mathbf{w}_2^*$$

$$= \frac{1}{4\sqrt{2}\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} (1\ 0\ 1) + \frac{1}{2\sqrt{3}\sqrt{2}\sqrt{6}} \begin{pmatrix} -1 \\ 1 \end{pmatrix} (-1\ -2\ 1)$$

$$= \frac{1}{8} \begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix} + \frac{1}{12} \begin{pmatrix} 1 & 2 & -1 \\ -1 & -2 & 1 \end{pmatrix} = \frac{1}{24} \begin{pmatrix} 5 & 4 & 1 \\ 1 & -4 & 5 \end{pmatrix}$$

which is exactly what we would have found by computing $A^\dagger = Q\Sigma^\dagger P^*$. Observe that

$$A^\dagger A = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad AA^\dagger = \frac{1}{3} \begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}$$

are the orthogonal projection matrices onto the spaces $\mathcal{N}(A)^\perp = \text{Span}\{\mathbf{v}_1, \mathbf{v}_2\} = \mathbb{R}^2$ and $\mathcal{R}(A) = \text{Span}\{\mathbf{w}_1, \mathbf{w}_2\} \leq \mathbb{R}^3$ respectively. Both spaces have dimension 2, since rank $A = 2$.

2. The pseudoinverse of $T = \frac{d}{dx} : P_2(\mathbb{R}) \to P_1(\mathbb{R})$, as seen in Example 2.66.2, maps

$$T^\dagger(\sqrt{3}(2x-1)) = \frac{1}{2\sqrt{15}}\sqrt{5}(6x^2 - 6x + 1) = \frac{1}{2\sqrt{3}}(6x^2 - 6x + 1)$$

$$T^\dagger(1) = \frac{1}{2\sqrt{3}}\sqrt{3}(2x-1) = x - \frac{1}{2}$$

$$\implies T^\dagger(a+bx) = T^\dagger\left(a + \frac{b}{2} + \frac{b}{2\sqrt{3}}\sqrt{3}(2x-1)\right)$$

$$= \left(a + \frac{b}{2}\right)\left(x - \frac{1}{2}\right) + \frac{b}{12}(6x^2 - 6x + 1)$$

$$= \frac{b}{2}x^2 + ax - \frac{a}{2} - \frac{b}{6}$$

The pseudoinverse of 'differentiation' therefore returns a particular choice of anti-derivative, namely the unique anti-derivative of $a + bx$ lying in $\text{Span}\{\sqrt{5}(6x^2 - 6x + 1), \sqrt{3}(2x - 1)\}$.

Exercises 2.7 1. Find the ingredients β, γ and the singular values for each of the following:

(a) $T \in \mathcal{L}(\mathbb{R}^2, \mathbb{R}^3)$ where $T\left(\begin{smallmatrix} x \\ y \end{smallmatrix}\right) = \left(\begin{smallmatrix} x \\ x+y \\ x-y \end{smallmatrix}\right)$

(b) $T : P_2(\mathbb{R}) \to P_1(\mathbb{R})$ and $T(f) = f''$ where $\langle f, g \rangle := \int_0^1 f(x)g(x)\,dx$

(c) $V = W = \text{Span}\{1, \sin x, \cos x\}$ and $\langle f, g \rangle = \int_0^{2\pi} f(x)g(x)\,dx$, with $T(f) = f' + 2f$

2. Find a singular value decomposition of each of the matrices:

(a) $\begin{pmatrix} 1 & 1 \\ 1 & 1 \\ -1 & -1 \end{pmatrix}$ (b) $\begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & -1 \end{pmatrix}$ (c) $\begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & 0 \\ 1 & 0 & -1 \end{pmatrix}$

3. Find an explicit formula for $T^\dagger$ for each map in Exercise 1.

4. Find the pseudoinverse of each of the matrix in Exercise 2.

5. Suppose $A = P\Sigma Q^*$ is a singular value decomposition.

   (a) Describe a singular value decomposition of A^*.
   (b) Explain why $A^\dagger = Q\Sigma^\dagger P^*$ *isn't* a singular value decomposition of $A^\dagger$; what would be the correct decomposition? (*Hint: what is wrong with $\Sigma^\dagger$?*)

6. Suppose $T : V \to W$ is written according to the singular value theorem. Prove that γ is a basis of eigenvectors of TT^* with the *same* non-zero eigenvalues as T^*T, including repetitions.

7. (a) Suppose $T = \mathcal{L}(V)$ is normal. Prove that each v_j in the singular value theorem may be chosen to be an eigenvector of T and that σ_j is the *modulus* of the corresponding eigenvalue.
   (b) Let $A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$. Show that *any* orthonormal basis β of $\mathbb{R}^2$ satisfies the singular value theorem. What is γ here? What is it about the *eigenvalues* of A that make this possible?

   (*Even when T is self-adjoint, the vectors in β need not also be eigenvectors of T!*)

8. In the proof of the singular value theorem we claimed that $\text{rank } T^*T = \text{rank } T$. Verify this by checking explicitly that $\mathcal{N}(T^*T) = \mathcal{N}(T)$.

   (*This is circular logic if you use the decomposition, so you must do without!*)

9. Let V, W be finite-dimensional inner product spaces and $T \in \mathcal{L}(V, W)$. Prove:

   (a) $T^*TT^\dagger = T^\dagger TT^* = T^*$.
   (*Hint: evaluate on the basis $\gamma = \{w_1, \ldots, w_m\}$ in the singular value theorem*)
   (b) If T is injective, then T^*T is invertible and $T^\dagger = (T^*T)^{-1}T^*$.
   (c) If T is surjective, then TT^* is invertible and $T^\dagger = T^*(TT^*)^{-1}$.

10. Consider the equation $T(x) = b$, where T is a linear map between finite-dimensional inner product spaces. A *least-squares solution* is a vector x which minimizes $||T(x) - b||$.

    (a) Prove that $x_0 = T^\dagger(b)$ is a least-squares solution and that any other has the form $x_0 + n$ for some $n \in \mathcal{N}(T)$.
    (*Hint: Theorem 2.61 says that x_0 is a least-squares solution if and only if $T(x_0) = \pi_{\mathcal{R}(T)}(b)$*)
    (b) Prove that $x_0 = T^\dagger(b)$ has smaller norm than any other least-squares solution.
    (c) If T is injective, prove that $x_0 = T^\dagger(b)$ is the unique least-squares solution.

11. Find the minimal norm solution to the first system, and the least-squares solution to the second:

$$\begin{cases} 3x + 2y + z = 9 \\ x - 2y + 3z = 3 \end{cases} \qquad \begin{cases} 3x + y = 1 \\ 2x - 2y = 0 \\ x + 3y = 0 \end{cases}$$

Linear Regression (non-examinable)

Given a data set $\{(t_j, y_j) : 1 \le j \le m\}$, we may employ the least-squares method to find a *best-fitting line* $y = c_0 + c_1 t$; often used to predict y given a value of t.

The trick is to minimize the sum of the squares of the vertical deviations of the line from the data set.

$$\sum_{j=1}^{m} |y_j - c_0 - c_1 t_j|^2 = ||\mathbf{y} - A\mathbf{x}||^2 \quad \text{where} \quad A = \begin{pmatrix} t_1 & 1 \\ \vdots & \vdots \\ t_m & 1 \end{pmatrix} \quad \mathbf{x} = \begin{pmatrix} c_1 \\ c_0 \end{pmatrix} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_m \end{pmatrix}$$

With the indicated notation, we recognize this as a least-squares problem. Indeed if there are at least two distinct t-values in the data set, then rank $A = 2$ is maximal and we have a unique best-fitting line with coefficients given by

$$\begin{pmatrix} c_1 \\ c_0 \end{pmatrix} = A^\dagger \mathbf{y} = (A^T A)^{-1} A^T \mathbf{y}$$

Example 2.69. Given the data set $\{(0,1), (1,1), (2,0), (3,2), (4,2)\}$, we compute

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \\ 2 & 1 \\ 3 & 1 \\ 4 & 1 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 2 \\ 2 \end{pmatrix} \implies \mathbf{x}_0 = \begin{pmatrix} c_1 \\ c_0 \end{pmatrix} = (A^T A)^{-1} A^T \mathbf{y} = \begin{pmatrix} 30 & 10 \\ 10 & 5 \end{pmatrix}^{-1} \begin{pmatrix} 15 \\ 6 \end{pmatrix} = \frac{3}{10} \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

The regression line therefore has equation $y = \frac{3}{10}(t + 2)$.

The process can be applied more generally to approximate using other functions. To find the best-fitting quadratic polynomial $y = c_0 + c_1 t + c_2 t^2$, we'd instead work with

$$A = \begin{pmatrix} t_1^2 & t_1 & 1 \\ \vdots & \vdots & \vdots \\ t_m^2 & t_m & 1 \end{pmatrix} \quad \mathbf{x} = \begin{pmatrix} c_2 \\ c_1 \\ c_0 \end{pmatrix} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_m \end{pmatrix} \implies \sum_{j=1}^{m} \left| y_j - c_0 - c_1 t - c_2 t_j^2 \right|^2 = ||\mathbf{y} - A\mathbf{x}||^2$$

Provided we have at least three distinct values t_1, t_2, t_3, the matrix A is guaranteed to have rank 3 and there will be a best-fitting least-squares quadratic: in this case

$$y = \frac{1}{70}(15t^2 - 39t + 72)$$

This curve and the best-fitting straight line are shown below.
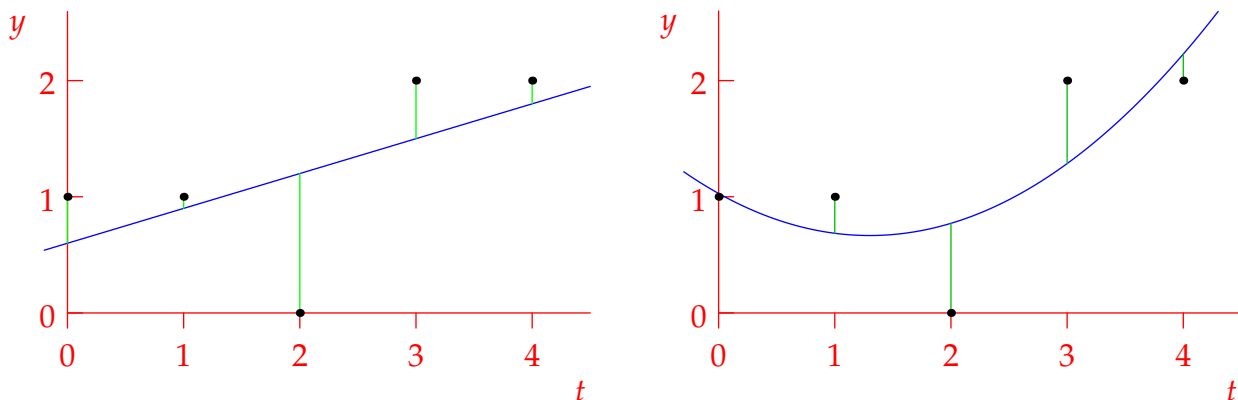
Optional Problems   Use a computer to invert any 3×3 matrices!

1. Check the calculation for the best-fitting least-squares quadratic in Example 2.69.

2. Find the best-fitting least-squares linear and quadratic approximations to the data set

$$\{(1,2), (3,4), (5,7), (7,9), (9,12)\}$$

3. Suppose a data set $\{(t_j, y_j) : 1 \leq j \leq m\}$ has unique regression line $y = ct + d$.

   (a) Show that the equations $A^T A \mathbf{x}_0 = A^T \mathbf{y}$ can be written in matrix form

$$\begin{pmatrix} \sum t_j^2 & \sum t_j \\ \sum t_j & m \end{pmatrix} \begin{pmatrix} c \\ d \end{pmatrix} = \begin{pmatrix} \sum t_j y_j \\ \sum y_j \end{pmatrix}$$

   (b) Recover the standard expressions from statistics:

$$c = \frac{\mathrm{Cov}(t,y)}{\sigma_t^2} \text{ and } d = \bar{y} - c\bar{t}$$

   where

   - $\bar{t} = \dfrac{1}{m} \sum_{j=1}^{m} t_j$ and $\bar{y} = \dfrac{1}{m} \sum_{j=1}^{m} y_j$ are the *means* (averages),

   - $\sigma_t^2 = \dfrac{1}{m} \sum_{j=1}^{m} (t_j - \bar{t})^2$ is the *variance*,

   - $\mathrm{Cov}(t,y) = \dfrac{1}{m} \sum_{j=1}^{m} (t_j - \bar{t})(y_j - \bar{y})$ is the *covariance*.

2.8 Bilinear and Quadratic Forms

In this section we slightly generalize the idea of an inner product. Throughout, V is a vector space over a field $\mathbb{F}$; it need not be an inner product space and $\mathbb{F}$ can be any field (not just $\mathbb{R}$ or $\mathbb{C}$).

Definition 2.70. A *bilinear form* $B : V \times V \to \mathbb{F}$ is linear in each entry: $\forall \mathbf{v}, \mathbf{x}, \mathbf{y} \in V, \lambda \in \mathbb{F}$,

$$B(\lambda\mathbf{x} + \mathbf{y}, \mathbf{v}) = \lambda B(\mathbf{x}, \mathbf{v}) + B(\mathbf{y}, \mathbf{v}), \qquad B(\mathbf{v}, \lambda\mathbf{x} + \mathbf{y}) = \lambda B(\mathbf{v}, \mathbf{x}) + B(\mathbf{v}, \mathbf{y})$$

Additionally, B is *symmetric* if $\forall \mathbf{x}, \mathbf{y} \in V, B(\mathbf{x}, \mathbf{y}) = B(\mathbf{y}, \mathbf{x})$.

Examples 2.71. 1. If V is a *real* inner product space, then the inner product $\langle \, , \, \rangle$ is a symmetric bilinear form. Note that a *complex* inner product is *not bilinear!*

2. If $A \in M_n(\mathbb{F})$, then $B(\mathbf{x}, \mathbf{y}) := \mathbf{x}^T A \mathbf{y}$ is a bilinear form on $\mathbb{F}^n$. For instance, on $\mathbb{R}^2$,

$$B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 1 & 2 \\ 2 & 0 \end{pmatrix} \mathbf{y} = x_1 y_1 + 2x_1 y_2 + 2x_2 y_1$$

defines a symmetric bilinear form, though not an inner product since it isn't positive definite; for example $B(\mathbf{j}, \mathbf{j}) = 0$.

As seen above, we often make use of a matrix.

Definition 2.72. Let B be a bilinear form on a finite-dimensional space with basis $\epsilon = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$. The *matrix of B with respect to ϵ* is the matrix $[B]_\epsilon = A \in M_n(\mathbb{F})$ with ij^{th} entry

$$A_{ij} = B(\mathbf{v}_i, \mathbf{v}_j)$$

Given $\mathbf{x}, \mathbf{y} \in V$, compute their co-ordinate vectors $[\mathbf{x}]_\epsilon, [\mathbf{y}]_\epsilon$ with respect to ϵ, then

$$B(\mathbf{x}, \mathbf{y}) = [\mathbf{x}]_\epsilon^T A [\mathbf{y}]_\epsilon$$

The set of bilinear forms on V is therefore in bijective correspondence with $M_n(\mathbb{F})$. Moreover,

$$B(\mathbf{y}, \mathbf{x}) = [\mathbf{y}]_\epsilon^T A [\mathbf{x}]_\epsilon = \left([\mathbf{y}]_\epsilon^T A [\mathbf{x}]_\epsilon \right)^T = [\mathbf{x}]_\epsilon^T A^T [\mathbf{y}]_\epsilon$$

Finally, if β is another basis of V, then an appeal to the change of co-ordinate matrix Q_ϵ^β yields

$$B(\mathbf{x}, \mathbf{y}) = [\mathbf{x}]_\epsilon^T A [\mathbf{y}]_\epsilon = (Q_\beta^\epsilon [\mathbf{x}]_\beta)^T A (Q_\beta^\epsilon [\mathbf{y}]_\epsilon) = [\mathbf{x}]_\beta^T (Q_\beta^\epsilon)^T A Q_\beta^\epsilon [\mathbf{y}]_\beta \implies [B]_\beta = (Q_\beta^\epsilon)^T [B]_\epsilon Q_\beta^\epsilon$$

To summarize:

Lemma 2.73. *Let B be a bilinear form on a finite-dimensional vector space.*

1. *If A is the matrix of B with respect to some basis, then every other matrix of B has the form $Q^T A Q$ for some invertible Q.*

2. *B is symmetric if and only if its matrix with respect to any (and all) bases is symmetric.*

Naturally, the simplest situation is when the matrix of B is diagonal...

Examples 2.74. 1. Example 2.71.2 can be written

$$B(\mathbf{x},\mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 1 & 2 \\ 2 & 0 \end{pmatrix} \mathbf{y} = x_1 y_1 + 2x_1 y_2 + 2x_2 y_1 = (x_1 + 2x_2)(y_1 + 2y_2) - 4x_2 y_2$$

$$= \begin{pmatrix} x_1 + 2x_2 & x_2 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix} \begin{pmatrix} y_1 + 2y_2 \\ y_2 \end{pmatrix} = \mathbf{x}^T \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}^T \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \mathbf{y}$$

If ϵ is the standard basis, then $[B]_\beta = \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix}$ where $Q_\epsilon^\beta = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$. It follows that $Q_\beta^\epsilon = \begin{pmatrix} 1 & -2 \\ 0 & 1 \end{pmatrix}$ from which $\beta = \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 1 \end{pmatrix} \right\}$ is a diagonalizing basis.

2. In general, we may perform a sequence of *simultaneous* row and column operations to diagonalize *any* symmetric B; we require only elementary matrices $E_{ij}^{(\lambda)}$ of type III.[14] For instance:

$$B(\mathbf{x},\mathbf{y}) = \mathbf{x}^T A \mathbf{y} = \mathbf{x}^T \begin{pmatrix} 1 & -2 & 3 \\ -2 & 0 & 4 \\ 3 & 4 & -1 \end{pmatrix} \mathbf{y} \qquad (A = [B]_\epsilon \text{ with respect to the standard basis})$$

$$E_{21}^{(2)} A E_{12}^{(2)} = \begin{pmatrix} 1 & 0 & 3 \\ 0 & -4 & 10 \\ 3 & 10 & -1 \end{pmatrix} \qquad \text{(add twice row 1 to row 2, columns similarly)}$$

$$E_{31}^{(-3)} E_{21}^{(2)} A E_{12}^{(2)} E_{13}^{(-3)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -4 & 10 \\ 0 & 10 & -10 \end{pmatrix} \qquad \text{(subtract 3 times row 1 from row 3, etc.)}$$

$$E_{23}^{(1)} E_{31}^{(-3)} E_{21}^{(2)} A E_{12}^{(2)} E_{13}^{(-3)} E_{32}^{(1)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & -10 \end{pmatrix} \qquad \text{(add row 3 to row 2, etc.)}$$

If β is the diagonalizing basis, the the change of co-ordinate matrix is

$$Q_\beta^\epsilon = E_{12}^{(2)} E_{13}^{(-3)} E_{32}^{(1)} = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & -3 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & -1 & -3 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}$$

from which $\beta = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} -3 \\ 0 \\ 1 \end{pmatrix} \right\}$. If you're having trouble believing this, invert the change of co-ordinate matrix and check that

$$B(\mathbf{x},\mathbf{y}) = (x_1 - 2x_2 + 3x_3)(y_1 - 2y_2 + 3y_3) + 6x_2 y_2 - 10(-x_2 + x_3)(-y_2 + y_3)$$

Warning! If $\mathbb{F} = \mathbb{R}$ then every symmetric B may be diagonalized by an orthonormal basis (for the usual dot product on $\mathbb{R}^n$). It is *very* unlikely that our algorithm will produce such! The algorithm has two main advantages over the spectral theorem: it is typically faster and it applies to vector spaces over any field. As a disadvantage, it is highly non-unique.

---

[14] Recall that $E_{ij}^{(\lambda)}$ is the identity matrix with an additional λ in the ij^{th} entry.
- As a *column* operation (right-multiplication), $A \mapsto A E_{ij}^{(\lambda)}$ adds λ times the i^{th} column to the j^{th}.
- As a *row* operation (left-multiplication), $A \mapsto E_{ji}^{(\lambda)} A = (E_{ij}^{(\lambda)})^T A$ adds λ times the i^{th} row to the j^{th}.

Example 2.75. We diagonalize $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \left(\begin{smallmatrix} 1 & 6 \\ 6 & 3 \end{smallmatrix}\right) \mathbf{y} = x_1 y_1 + 6 x_1 y_2 + 6 x_2 y_1 + 3 x_2 y_2$ in three ways.

- $\left(\begin{smallmatrix} 1 & 0 \\ -6 & 1 \end{smallmatrix}\right) \left(\begin{smallmatrix} 1 & 6 \\ 6 & 3 \end{smallmatrix}\right) \left(\begin{smallmatrix} 1 & -6 \\ 0 & 1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -33 \end{smallmatrix}\right) = [B]_\beta$ where $\beta = \left\{ \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} -6 \\ 1 \end{smallmatrix}\right) \right\}$. This corresponds to

  $$B(\mathbf{x}, \mathbf{y}) = (x_1 + 6x_2)(y_1 + 6y_2) - 33 x_2 y_2$$

- $\left(\begin{smallmatrix} 1 & -2 \\ 0 & 1 \end{smallmatrix}\right) \left(\begin{smallmatrix} 1 & 6 \\ 6 & 3 \end{smallmatrix}\right) \left(\begin{smallmatrix} 1 & 0 \\ -2 & 1 \end{smallmatrix}\right) = \left(\begin{smallmatrix} -11 & 0 \\ 0 & 3 \end{smallmatrix}\right) = [B]_\gamma$ where $\gamma = \left\{ \left(\begin{smallmatrix} 1 \\ -2 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right) \right\}$. This corresponds to

  $$B(\mathbf{x}, \mathbf{y}) = -11 x_1 y_1 + 3(2x_1 + x_2)(2y_1 + y_2)$$

- If $\mathbb{F} = \mathbb{R}$, we may apply the spectral theorem to see that $[B]_\eta = \left(\begin{smallmatrix} 2+\sqrt{37} & 0 \\ 0 & 2-\sqrt{37} \end{smallmatrix}\right)$ is diagonal with respect to $\eta = \left\{ \left(\begin{smallmatrix} 6 \\ 1+\sqrt{37} \end{smallmatrix}\right), \left(\begin{smallmatrix} -1-\sqrt{37} \\ 6 \end{smallmatrix}\right) \right\}$. The expression for $B(\mathbf{x}, \mathbf{y})$ in these co-ordinates is disgusting, so we omit it; that B can be diagonalized orthogonally doesn't mean it should be!

Theorem 2.76. *Suppose B is a bilinear form of a finite-dimensional space V over $\mathbb{F}$.*

1. *If B is diagonalizable, then it is symmetric.*

2. *If B is symmetric and $\mathbb{F}$ does not have characteristic two (see aside), then B is diagonalizable.*

Proof. 1. If B is diagonalizable, $\exists \beta$ such that $[B]_\beta$ is diagonal and thus symmetric.

2. Suppose B is non-zero (otherwise the result is trivial). We prove by induction on $n = \dim V$.

If $n = 1$, the result is trivial: $B(x, y) = axy$ for some $a \in \mathbb{F}$ is clearly symmetric.

Fix $n \in \mathbb{N}$ and assume that every non-zero symmetric bilinear form on a dimension n vector space over $\mathbb{F}$ is diagonalizable. Let $\dim V = n + 1$. By the discussion below, $\exists \mathbf{x} \in V$ such that $B(\mathbf{x}, \mathbf{x}) \neq 0$. Consider the linear map

$$T : V \to \mathbb{F} : \mathbf{v} \mapsto B(\mathbf{x}, \mathbf{v})$$

Clearly $\operatorname{rank} T = 1 \implies \dim \mathcal{N}(T) = n$. Moreover, B is symmetric when restricted to $\mathcal{N}(T)$; by the induction hypothesis there exists a basis β of $\mathcal{N}(T)$ such that $[B_{\mathcal{N}(T)}]_\beta$ is diagonal. But then B is diagonal with respect to the basis $\beta \cup \{\mathbf{x}\}$. ∎

Aside: Characteristic two fields This means $1 + 1 = 0$ in $\mathbb{F}$; this holds, for instance, in the field $\mathbb{Z}_2 = \{0, 1\}$ of remainders modulo 2. We now see the importance of $\operatorname{char} \mathbb{F} \neq 2$ to the above result.

- The proof uses the existence of $\mathbf{x} \in V$ such that $B(\mathbf{x}, \mathbf{x}) \neq 0$. If B is non-zero, $\exists \mathbf{u}, \mathbf{v}$ such that $B(\mathbf{u}, \mathbf{v}) \neq 0$. If both $B(\mathbf{u}, \mathbf{u}) = 0 = B(\mathbf{v}, \mathbf{v})$, then $\mathbf{x} = \mathbf{u} + \mathbf{v}$ does the job whenever $\operatorname{char} \mathbb{F} \neq 2$:

  $$B(\mathbf{x}, \mathbf{x}) = B(\mathbf{u}, \mathbf{v}) + B(\mathbf{v}, \mathbf{u}) = 2B(\mathbf{u}, \mathbf{v}) \neq 0$$

- To see that the requirement isn't idle, consider $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right) \mathbf{y}$ on the finite vector space $\mathbb{Z}_2^2 = \left\{ \left(\begin{smallmatrix} 0 \\ 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 1 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 \\ 1 \end{smallmatrix}\right) \right\}$ over $\mathbb{Z}_2$. Every element of this space satisfies $B(\mathbf{x}, \mathbf{x}) = 0$! Perhaps surprisingly, the matrix of B is identical with respect to *any* basis of $\mathbb{Z}_2^2$, whence B is symmetric but non-diagonalizable.

In Example 2.75 notice how the three diagonal matrix representations have something in common: one each of the diagonal entries are positive and negative. This is a general phenomenon:

> **Theorem 2.77 (Sylvester's Law of Inertia).** *Suppose B is a symmetric bilinear form on a real vector space V with diagonal matrix representation $\mathrm{diag}(\lambda_1,\dots,\lambda_n)$. Then the number of entries λ_j which are positive/negative/zero is independent of the diagonal representation.*

> **Definition 2.78.** The *signature* of a symmetric bilinear form B is the triple (n_+, n_-, n_0) representing how many positive, negative and zero terms are in any diagonal representation. Sylvester's Law says that the signature is an *invariant* of a symmetric bilinear form.

Positive-definiteness says that a real inner product on an n-dimensional space has signature $(n,0,0)$. Practitioners of relativity often work in *Minkowski spacetime*: $\mathbb{R}^4$ equipped with a signature $(1,3,0)$ bilinear form, typically

$$B(\mathbf{x},\mathbf{y}) = \mathbf{x}^T \begin{pmatrix} c^2 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \mathbf{y} = c^2 x_1 y_1 - x_2 y_2 - x_3 y_3 - x_4 y_4$$

where c is the speed of light. Vectors are *time-*, *space-*, or *light-like* depending on whether $B(\mathbf{x},\mathbf{x})$ is positive, negative or zero. For instance $\mathbf{x} = 3c^{-1}\mathbf{e}_1 + 2\mathbf{e}_2 + 2\mathbf{e}_3 + \mathbf{e}_4$ is light-like.

Proof. For simplicity, let $V = \mathbb{R}^n$ and write $B(\mathbf{x},\mathbf{y}) = \mathbf{x}^T A \mathbf{y}$ where A is symmetric.

1. Define rank $B :=$ rank A and observe this is independent of basis (exercises).

2. Let β, γ be diagonalizing bases, ordered according to whether B is positive, negative or zero.

   $$\beta = \{\mathbf{v}_1,\dots,\mathbf{v}_p,\mathbf{v}_{p+1},\dots,\mathbf{v}_r,\mathbf{v}_{r+1},\dots,\mathbf{v}_n\}$$
   $$\gamma = \{\mathbf{w}_1,\dots,\mathbf{w}_q,\mathbf{w}_{q+1},\dots,\mathbf{w}_r,\mathbf{w}_{r+1},\dots,\mathbf{w}_n\}$$

   Here $r =$ rank B in accordance with part 1: our goal is to prove that $p = q$.

3. Assume $p < q$, define the matrix C and check what follows:

   $$C = \begin{pmatrix} \mathbf{v}_1^T A \\ \vdots \\ \mathbf{v}_p^T A \\ \mathbf{w}_{q+1}^T A \\ \vdots \\ \mathbf{w}_r^T A \end{pmatrix} \in M_{(r-q+p)\times n}(\mathbb{R})$$

   (a) rank $C \le r - q + p < r \implies$ null $C > n - r$, thus

   $$\exists \mathbf{x} \in \mathbb{R}^n \text{ such that } C\mathbf{x} = \mathbf{0} \text{ and } \mathbf{x} \notin \mathrm{Span}\{\mathbf{v}_{r+1},\dots,\mathbf{v}_n\}$$

   (b) The first p entries of $C\mathbf{x} = \mathbf{0}$ mean that $\mathbf{x} \in \mathrm{Span}\{\mathbf{v}_{p+1},\dots,\mathbf{v}_r,\mathbf{v}_{r+1},\dots,\mathbf{v}_n\}$ and so $B(\mathbf{x},\mathbf{x}) < 0$. Note how we use part (a) to get a strict inequality here.

   (c) Now write $\mathbf{x}$ with respect to γ: this time we see that $\mathbf{x} \in \mathrm{Span}\{\mathbf{w}_1,\dots,\mathbf{w}_q,\mathbf{w}_{r+1},\dots,\mathbf{w}_n\}$, whence $B(\mathbf{x},\mathbf{x}) \ge 0$: a contradiction. ∎

Quadratic Forms & Diagonalizing Conics

> **Definition 2.79.** To every symmetric bilinear form $B : V \times V \to \mathbb{F}$ is associated a *quadratic form*
>
> $$K : V \to \mathbb{F} : \mathbf{x} \mapsto B(\mathbf{x}, \mathbf{x})$$
>
> A function $K : V \to \mathbb{F}$ is termed a quadratic form when such a symmetric bilinear form exists.

Examples 2.80. 1. If B is a real inner product, then $K(\mathbf{v}) = \langle \mathbf{v}, \mathbf{v} \rangle = ||\mathbf{v}||^2$ is the square of the norm.

2. Let $\dim V = n$ and A be the matrix of B with respect to a basis β. By the symmetry of A,

$$K(\mathbf{x}) = \mathbf{x}^T A \mathbf{x} = \sum_{i,j=1}^{n} x_i A_{ij} x_j = \sum_{1 \le i \le j \le n} \tilde{a}_{ij} x_i x_j \quad \text{where } \tilde{a}_{ij} = \begin{cases} A_{ij} & \text{if } i = j \\ 2A_{ij} & \text{if } i \ne j \end{cases}$$

E.g., $K(\mathbf{x}) = 3x_1^2 + 4x_2^2 - 2x_1 x_2$ corresponds to the bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \left(\begin{smallmatrix} 3 & -1 \\ -1 & 4 \end{smallmatrix} \right) \mathbf{y}$

As a fun application, we consider the diagonalization of conics in $\mathbb{R}^2$. The general non-zero conic has equation

$$ax^2 + 2bxy + cy^2 + dx + ey + f = 0$$

where the first three terms comprise a quadratic form

$$K(\mathbf{x}) = ax^2 + 2bxy + cy^2 \rightsquigarrow B(\mathbf{v}, \mathbf{w}) = \mathbf{v}^T \begin{pmatrix} a & b \\ b & c \end{pmatrix} \mathbf{w}$$

If $\{\mathbf{v}_1, \mathbf{v}_2\}$ is a diagonalizing basis, then there exist scalars λ_1, λ_2 for which

$$K(t_1 \mathbf{v}_1 + t_2 \mathbf{v}_2) = \lambda_1 t_1^2 + \lambda_2 t_2^2$$

whence the general conic becomes

$$\lambda_1 t_1^2 + \lambda_2 t_2^2 + \mu_1 t_1 + \mu_2 t_2 = \eta, \quad \lambda_1, \lambda_2, \mu_1, \mu_2, \eta \in \mathbb{R}$$

If $\lambda_i \ne 0$, we may complete the square via the linear transformation $s_j = t_j + \frac{\mu_j}{2\lambda_j}$. The canonical forms are then recovered:

Parabola: λ_1 or $\lambda_2 = 0$ (but not both).

Ellipse: $\lambda_1 s_1^2 + \lambda_2 s_2^2 = k \ne 0$ where λ_1, λ_2, k have the same sign.

Hyperbola: $\lambda_1 s_1^2 + \lambda_2 s_2^2 = k \ne 0$ where λ_1, λ_2 have opposite signs.

Since B is symmetric, we may take $\{\mathbf{v}_1, \mathbf{v}_2\}$ to be an *orthonormal* basis of $\mathbb{R}^2$, whence any conic may be put in canonical form by applying only a rotation/reflection and translation (completing the square). Alternatively, we could diagonalize K using our earlier algorithm; this additionally permits shear transforms. By Sylvester's Law, the diagonal entries will have the same number of $(+, -, 0)$ terms regardless of the method, so the canonical form will be unchanged.

Examples 2.81. 1. We describe and plot the conic with equation $7x^2 + 24xy = 144$.

The matrix of the associated bilinear form is $\left(\begin{smallmatrix} 7 & 12 \\ 12 & 0 \end{smallmatrix}\right)$ which has orthonormal eigenbasis

$$\beta = \{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \frac{1}{5}\begin{pmatrix} 4 \\ 3 \end{pmatrix}, \frac{1}{5}\begin{pmatrix} -3 \\ 4 \end{pmatrix} \right\}$$

with eigenvalues $(\lambda_1, \lambda_2) = (16, -9)$. In the rotated basis, this is the canonical hyperbola

$$16t_1^2 - 9t_2^2 = 144 \iff \frac{t_1^2}{3^2} - \frac{t_2^2}{4^2} = 1$$

which is easily plotted. In case this is too fast, use the change of co-ordinate matrix to compute directly:

$$Q_\beta^\epsilon = \frac{1}{5}\begin{pmatrix} 4 & -3 \\ 3 & 4 \end{pmatrix} \implies \begin{pmatrix} t_1 \\ t_2 \end{pmatrix} = [\mathbf{x}]_\beta = Q_\epsilon^\beta \begin{pmatrix} x \\ y \end{pmatrix} = \frac{1}{5}\begin{pmatrix} 4x + 3y \\ -3x + 4y \end{pmatrix}$$

which quickly recovers the original conic by substitution.
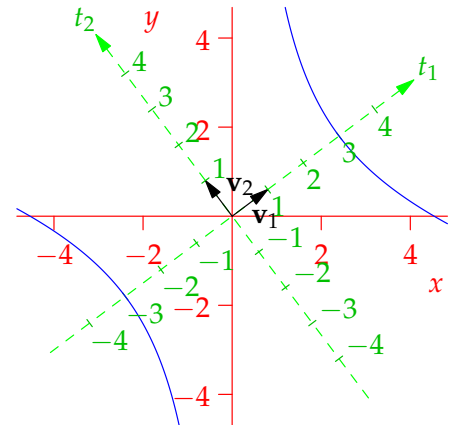
2. The conic defined by $K(\mathbf{x}) = x^2 + 12xy + 3y^2 = -33$ defines a hyperbola in accordance with Example 2.75. With respect to the basis $\gamma = \left\{ \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$, we see that

$$K(\mathbf{x}) = -11x^2 + 3(2x + y)^2 = -11t_1^2 + 3t_2^2 = -33$$
$$\iff \frac{t_1^2}{3} - \frac{t_2^2}{11} = 1$$

Even though γ is non-orthogonal, we can still plot the conic! If we instead used the orthonormal basis η, we'd obtain

$$K(\mathbf{x}) = (\sqrt{37} + 2)s_1^2 - (\sqrt{37} - 2)s_2^2 = -33$$

however the calculation to find η is time-consuming and the expressions for s_1, s_2 are extremely ugly.
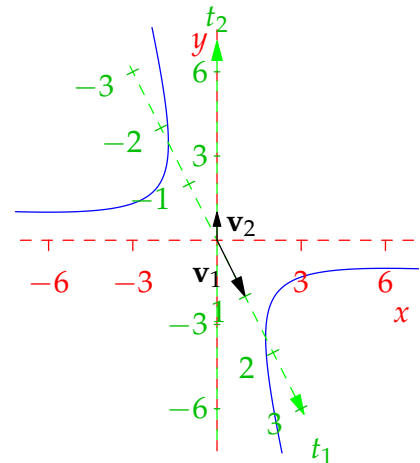
A similar approach can be applied to higher degree quadratic equations/manifolds: e.g. ellipsoids, paraboloids and hyperboloids in $\mathbb{R}^3$.

Exercises 2.8 1. Prove that the sum of any two bilinear forms is bilinear, and that any scalar multiple of a bilinear form is bilinear: thus the set of bilinear forms on V is a vector space.

(*You can't use matrices here, since V could be infinite-dimensional!*)

2. Compute the matrix of the bilinear form

$$B(\mathbf{x}, \mathbf{y}) = x_1 y_1 - 2x_1 y_2 + x_2 y_1 - x_3 y_3$$

on $\mathbb{R}^3$ with respect to the basis $\beta = \left\{ \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \right\}$.

3. Check that the function $B(f, g) = f'(0)g''(0)$ is a bilinear form on the vector space of twice-differentiable functions. Find the matrix of B with respect to $\beta = \{\cos t, \sin t, \cos 2t, \sin 2t\}$ when restricted to the subspace Span β.

4. For each matrix A, find a diagonal matrix D and an invertible matrix Q such that $Q^T A Q = D$.

(a) $\begin{pmatrix} 1 & 3 \\ 3 & 2 \end{pmatrix}$ (b) $\begin{pmatrix} 3 & 1 & 2 \\ 1 & 4 & 0 \\ 2 & 0 & -1 \end{pmatrix}$

5. If K is a quadratic form and $K(\mathbf{x}) = 2$, what is the value of $K(3\mathbf{x})$?

6. If $\mathbb{F}$ does not have characteristic 2, and $K(\mathbf{x}) = B(\mathbf{x}, \mathbf{x})$ is a quadratic form, prove that we can recover the bilinear form B via

$$B(\mathbf{x}, \mathbf{y}) = \frac{1}{2}\left(K(\mathbf{x} + \mathbf{y}) - K(\mathbf{x}) - K(\mathbf{y})\right)$$

7. If $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \mathbf{y}$ is a bilinear form on $\mathbb{F}^2$, compute the quadratic form $K(\mathbf{x})$.

8. Suppose B is a symmetric bilinear form on a real finite-dimensional space. With reference to the proof of Sylvester's Law, explain why rank B is independent of the choice of diagonalizing basis.

9. If char $\mathbb{F} \neq 2$, apply the diagonalizing algorithm to the symmetric bilinear form $B = \mathbf{x}^T \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \mathbf{y}$ on $\mathbb{F}^2$. What goes wrong if char $\mathbb{F} = 2$?

10. Describe and plot the following conics:

(a) $x^2 + y^2 + xy = 6$

(b) $35x^2 + 120xy = 4x + 3y$

11. Suppose that a non-empty, non-degenerate[15] conic C in $\mathbb{R}^2$ has the form $ax^2 + 2bxy + cy^2 + dx + ey + f = 0$, where at least one of $a, b, c \neq 0$, and define $\Delta = b^2 - ac$. Prove that:

- C is a parabola if and only if $\Delta = 0$;
- C is an ellipse if and only if $\Delta < 0$;
- C is a hyperbola if and only if $\Delta > 0$.

(Hint: λ_1, λ_2 are the eigenvalues of a symmetric matrix, so. . .)

---

[15] The conic contains at least two points and cannot be factorized as a product of two straight lines: for example, the following are disallowed;

- $x^2 + y^2 + 1 = 0$ is empty (unless one allows conics over $\mathbb{C}$ . . .)
- $x^2 + y^2 = 0$ contains only one point;
- $x^2 - xy - x + y = (x - 1)(x - y) = 0$ is the product of two lines.

3 Canonical Forms

3.1 Jordan Forms & Generalized Eigenvectors

Throughout this course we've concerned ourselves with variations of a general question: for a given map $T \in \mathcal{L}(V)$, find a basis β such that the matrix $[T]_\beta$ is as close to diagonal as possible. In this chapter we see what is possible when T is non-diagonalizable.

Example 3.1. The matrix $A = \begin{pmatrix} -8 & 4 \\ -25 & 12 \end{pmatrix} \in M_2(\mathbb{R})$ has characteristic equation

$$p(t) = (-8 - t)(12 - t) + 4 \cdot 25 = t^2 - 4t + 4 = (t - 2)^2$$

and thus a single eigenvalue $\lambda = 2$. It is non-diagonalizable since the eigenspace is one-dimensional

$$E_2 = \mathcal{N} \begin{pmatrix} -10 & 4 \\ -25 & 10 \end{pmatrix} = \mathrm{Span} \begin{pmatrix} 2 \\ 5 \end{pmatrix}$$

However, if we consider a basis $\beta = \{\mathbf{v}_1, \mathbf{v}_2\}$ where $\mathbf{v}_1 = \begin{pmatrix} 2 \\ 5 \end{pmatrix}$ is an eigenvector, then $[L_A]_\beta$ is upper-triangular, which is better than nothing! How simple can we make this matrix? Let $\mathbf{v}_2 = \begin{pmatrix} x \\ y \end{pmatrix}$, then

$$A\mathbf{v}_2 = \begin{pmatrix} -8x + 4y \\ -25x + 12y \end{pmatrix} = 2 \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} -10x + 4y \\ -25x + 10y \end{pmatrix} = 2\mathbf{v}_2 + (-5x + 2y)\mathbf{v}_1$$

$$\implies [L_A]_\beta = \begin{pmatrix} 2 & -5x + 2y \\ 0 & 2 \end{pmatrix}$$

Since $\mathbf{v}_2$ cannot be parallel to $\mathbf{v}_1$, the only thing we cannot have is a diagonal matrix. The next best thing is for the upper right corner be 1; for instance we could choose

$$\beta = \{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \begin{pmatrix} 2 \\ 5 \end{pmatrix}, \begin{pmatrix} 1 \\ 3 \end{pmatrix} \right\} \implies [L_A]_\beta = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}$$

Definition 3.2. A *Jordan block* is a square matrix of the form

$$J = \begin{pmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{pmatrix}$$

where all non-indicated entries are zero. Any 1×1 matrix is also a Jordan block.

A *Jordan canonical form* is a block-diagonal matrix $\mathrm{diag}(J_1, \ldots, J_m)$ where each J_k is a Jordan block.

A *Jordan canonical basis* for $T \in \mathcal{L}(V)$ is a basis β of V such that $[T]_\beta$ is a Jordan canonical form.

If a map is diagonalizable, then any eigenbasis is Jordan canonical and the corresponding Jordan canonical form is diagonal. What about more generally? Does every non-diagonalizable map have a Jordan canonical basis? If so, how can we find such?

Example 3.3. It can easily be checked that $\beta = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\} = \left\{ \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}$ is a Jordan canonical basis for

$$A = \begin{pmatrix} -1 & 2 & 3 \\ -4 & 5 & 4 \\ -2 & 1 & 4 \end{pmatrix}$$

(really $L_A \in \mathcal{L}(\mathbb{R}^3)$). Indeed

$$A\mathbf{v}_1 = 2\mathbf{v}_1, \quad A\mathbf{v}_2 = 3\mathbf{v}_2, \quad A\mathbf{v}_3 = \begin{pmatrix} 4 \\ 5 \\ 3 \end{pmatrix} = \begin{pmatrix} 1+3 \\ 2+3 \\ 0+3 \end{pmatrix} = \mathbf{v}_2 + 3\mathbf{v}_3 \implies [L_A]_\beta = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 3 & 1 \\ 0 & 0 & 3 \end{pmatrix}$$

Generalized Eigenvectors

Example 3.3 was easy to check, but how would we go about finding a suitable β if we were merely given A? We brute-forced this in Example 3.1, but such is not a reasonable approach in general. Eigenvectors get us some of the way:

- $\mathbf{v}_1$ is an eigenvector in Example 3.1, but $\mathbf{v}_2$ is not.

- $\mathbf{v}_1$ and $\mathbf{v}_2$ are eigenvectors in Example 3.3, but $\mathbf{v}_3$ is not.

The practical question is how to fill out a Jordan canonical basis once we have a maximal independent set of eigenvectors. We now define the necessary objects.

Definition 3.4. Suppose $T \in \mathcal{L}(V)$ has an eigenvalue λ. Its *generalized eigenspace* is

$$K_\lambda := \{\mathbf{x} \in V : (T - \lambda I)^k(\mathbf{x}) = \mathbf{0} \text{ for some } k \in \mathbb{N}\} = \bigcup_{k \in \mathbb{N}} \mathcal{N}(T - \lambda I)^k$$

A *generalized eigenvector* is any non-zero $\mathbf{v} \in K_\lambda$.

As with eigenspaces, the generalized eigenspaces of $A \in M_n(\mathbb{F})$ are those of the map $L_A \in \mathcal{L}(\mathbb{F}^n)$.

It is easy to check that our earlier Jordan canonical bases consist of generalized eigenvectors.

Example 3.1: We have one eigenvalue $\lambda = 2$. Since $(A - 2I)^2 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$ is the zero matrix, every non-zero vector is a generalized eigenvector; plainly $K_2 = \mathbb{R}^2$.

Example 3.3: We see that

$$(A - 2I)\mathbf{v}_1 = \mathbf{0}, \quad (A - 3I)\mathbf{v}_2 = \mathbf{0}, \quad (A - 3I)^2\mathbf{v}_3 = (A - 3I)\mathbf{v}_2 = \mathbf{0}$$

whence β is a basis of generalized eigenvectors. Indeed

$$K_3 = \text{Span}\{\mathbf{v}_1, \mathbf{v}_2\}, \quad K_2 = E_2 = \text{Span}\{\mathbf{v}_3\}$$

though verifying this with current technology is a little awkward...

In order to easily compute generalized eigenspaces, it is useful to invoke the main result of this section. We postpone the proof for a while due to its meatiness.

> **Theorem 3.5.** *Suppose that the characteristic polynomial of* $T \in \mathcal{L}(V)$ *splits over* $\mathbb{F}$:
>
> $$p(t) = (\lambda_1 - t)^{m_1} \cdots (\lambda_k - t)^{m_k}$$
>
> *where the* λ_j *are the* distinct *eigenvalues of* T *with algebraic multiplicities* m_j. *Then:*
>
> 1. *For each eigenvalue;* (a) $K_\lambda = \mathcal{N}(T - \lambda I)^m$ *and* (b) $\dim K_\lambda = m$.
> 2. $V = K_{\lambda_1} \oplus \cdots \oplus K_{\lambda_k}$: *there exists a basis of generalized eigenvectors.*

Compare this with the statement on diagonalizability from the start of the course.

With regard to part 2; we shall eventually be able to *choose* this to be a Jordan canonical basis. In conclusion: a map has a Jordan canonical basis if and only if its characteristic polynomial splits.

Examples 3.6. 1. Observe how Example 3.3 works in this language:

$$A = \begin{pmatrix} -1 & 2 & 3 \\ -4 & 5 & 4 \\ -2 & 1 & 4 \end{pmatrix} \implies p(t) = (2-t)^1(3-t)^2$$

$$K_2 = \mathcal{N}(A - 2I)^1 = \text{Span} \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} \implies \dim K_2 = 1$$

$$K_3 = \mathcal{N}(A - 3I)^2 = \mathcal{N} \begin{pmatrix} 2 & -1 & -1 \\ 0 & 0 & 0 \\ 2 & -1 & -1 \end{pmatrix} = \text{Span} \left\{ \begin{pmatrix} 1 \\ 2 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\} \implies \dim K_3 = 2$$

$$\mathbb{R}^3 = K_2 \oplus K_3$$

2. We find the generalized eigenspaces of the matrix $A = \begin{pmatrix} 5 & 2 & -1 \\ 0 & 0 & 0 \\ 9 & 6 & -1 \end{pmatrix}$

The characteristic polynomial is

$$p(t) = \det(A - \lambda I) = -t \begin{vmatrix} 5-t & -1 \\ 9 & -1-t \end{vmatrix} = -t(t^2 - 5t + t - 5 + 9) = -(0-t)^1(2-t)^2$$

- $\lambda = 0$ has multiplicity 1; indeed $K_0 = \mathcal{N}(A - 0I)^1 = \mathcal{N}(A) = \text{Span} \begin{pmatrix} 1 \\ -1 \\ 3 \end{pmatrix}$ is just the eigenspace E_0.

- $\lambda = 2$ has multiplicity 2,

$$K_2 = \mathcal{N}(A - 2I)^2 = \mathcal{N} \begin{pmatrix} 3 & 2 & -1 \\ 0 & -2 & 0 \\ 9 & 6 & -3 \end{pmatrix}^2 = \mathcal{N} \begin{pmatrix} 0 & -4 & 0 \\ 0 & 4 & 0 \\ 0 & -12 & 0 \end{pmatrix} = \text{Span} \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\}$$

In this case the corresponding eigenspace is *one*-dimensional, $E_2 = \text{Span} \begin{pmatrix} 1 \\ 0 \\ 3 \end{pmatrix} \lneq K_2$, and the matrix is non-diagonalizable.

Observe also that $\mathbb{R}^3 = K_0 \oplus K_2$ in accordance with the Theorem.

Properties of Generalized Eigenspaces and the Proof of Theorem 3.5

A lot of work is required to justify our main result. Feel free to skip the proofs at first reading.

> **Lemma 3.7.** *Let λ be an eigenvalue of $\mathrm{T} \in \mathcal{L}(V)$. Then:*
>
> 1. *E_λ is a subspace of K_λ, which is itself a subspace of V.*
>
> 2. *K_λ is T-invariant.*
>
> 3. *Suppose K_λ is finite-dimensional and $\mu \neq \lambda$. Then:*
>
>   (a) *K_λ is $(\mathrm{T} - \mu\mathrm{I})$-invariant and the restriction of $\mathrm{T} - \mu\mathrm{I}$ to K_λ is an isomorphism.*
>   (b) *If μ is another eigenvalue, then $K_\lambda \cap K_\mu = \{\mathbf{0}\}$. In particular K_λ contains no eigenvectors other than those in E_λ.*

Proof.   1. These are an easy exercise.

2. Let $\mathbf{x} \in K_\lambda$, then $\exists k$ such that $(\mathrm{T} - \lambda\mathrm{I})^k(\mathbf{x}) = \mathbf{0}$. But then

$$(\mathrm{T} - \lambda\mathrm{I})^k\big(\mathrm{T}(\mathbf{x})\big) = (\mathrm{T} - \lambda\mathrm{I})^k\big(\mathrm{T}(\mathbf{x}) - \lambda\mathbf{x} + \lambda\mathbf{x}\big)$$
$$= (\mathrm{T} - \lambda\mathrm{I})^{k+1}(\mathbf{x}) + \lambda(\mathrm{T} - \lambda\mathrm{I})^k(\mathbf{x}) = \mathbf{0}$$

Otherwise said, $\mathrm{T}(\mathbf{x}) \in K_\lambda$.

3.   (a) Let $\mathbf{x} \in K_\lambda$. Part 2 tells us that

$$(\mathrm{T} - \mu\mathrm{I})(\mathbf{x}) = \mathrm{T}(\mathbf{x}) - \mu\mathbf{x} \in K_\lambda$$

whence K_λ is $(\mathrm{T} - \mu\mathrm{I})$-invariant.

Suppose, for a contradiction, that $\mathrm{T} - \mu\mathrm{I}$ is *not injective* on K_λ. Then

$$\exists \mathbf{y} \in K_\lambda \setminus \{\mathbf{0}\} \text{ such that } (\mathrm{T} - \mu\mathrm{I})(\mathbf{y}) = \mathbf{0}$$

Let $k \in \mathbb{N}$ be minimal such that $(\mathrm{T} - \lambda\mathrm{I})^k(\mathbf{y}) = \mathbf{0}$ and let $\mathbf{z} = (\mathrm{T} - \lambda\mathrm{I})^{k-1}(\mathbf{y})$. Plainly $\mathbf{z} \neq \mathbf{0}$, for otherwise k is not minimal. Moreover,

$$(\mathrm{T} - \lambda\mathrm{I})(\mathbf{z}) = (\mathrm{T} - \lambda\mathrm{I})^k(\mathbf{y}) = \mathbf{0} \implies \mathbf{z} \in E_\lambda$$

Since $\mathrm{T} - \mu\mathrm{I}$ and $\mathrm{T} - \lambda\mathrm{I}$ commute, we can also compute the effect of $\mathrm{T} - \mu\mathrm{I}$;

$$(\mathrm{T} - \mu\mathrm{I})(\mathbf{z}) = (\mathrm{T} - \mu\mathrm{I})(\mathrm{T} - \lambda\mathrm{I})^{k-1}(\mathbf{y}) = (\mathrm{T} - \lambda\mathrm{I})^{k-1}(\mathrm{T} - \mu\mathrm{I})(\mathbf{x}) = \mathbf{0}$$

which says that $\mathbf{z}$ is an eigenvector in E_μ; if μ isn't an eigenvalue, then we already have our contradiction! Even if μ is an eigenvalue, $E_\mu \cap E_\lambda = \{\mathbf{0}\}$ provides the desired contradiction.

We conclude that $(\mathrm{T} - \mu\mathrm{I})_{K_\lambda} \in \mathcal{L}(K_\lambda)$ is injective. Since $\dim K_\lambda < \infty$, the restriction is automatically an isomorphism.

(b) This is another exercise.   ∎

Now to prove Theorem 3.5: remember that the characteristic polynomial of T is assumed to split.

Proof. (Part 1(a)) Fix an eigenvalue λ. By definition, we have $\mathcal{N}(T - \lambda I)^m \le K_\lambda$.

For the converse, parts 2 and 3 of the Lemma tell us (why?) that

$$p_\lambda(t) = (\lambda - t)^{\dim K_\lambda} \quad \text{from which} \quad \dim K_\lambda \le m \tag{$*$}$$

By the Cayley–Hamilton Theorem, T_{K_λ} satisfies its characteristic polynomial, whence

$$\forall \mathbf{x} \in K_\lambda, \ (\lambda I - T)^{\dim K_\lambda}(\mathbf{x}) = \mathbf{0} \implies K_\lambda \le \mathcal{N}(T - \lambda I)^m$$

(Parts 1(b) and 2) We prove simultaneously by induction on the number of distinct eigenvalues of T.

(*Base case*) If T has only one eigenvalue, then $p(t) = (\lambda - t)^m$. Another appeal to Cayley–Hamilton says $(T - \lambda I)^m(\mathbf{x}) = \mathbf{0}$ for all $\mathbf{x} \in V$. Thus $V = K_\lambda$ and $\dim K_\lambda = m$.

(*Induction step*) Fix k and suppose the results hold for maps with k distinct eigenvalues. Let T have distinct eigenvalues $\lambda_1, \dots, \lambda_k, \mu$, with multiplicities $m_1, \dots, m_k, m$ respectively. Define[16]

$$W = \mathcal{R}(T - \mu I)^m$$

The subspace W has the following properties, the first two of which we leave as exercises:

- W is T-invariant.
- $W \cap K_\mu = \{\mathbf{0}\}$ so that μ is not an eigenvalue of the restriction T_W.
- Each $K_{\lambda_j} \le W$: since $(T - \mu I)_{K_{\lambda_j}}$ is an isomorphism (Lemma part 3), we can invert,

$$\mathbf{x} \in K_{\lambda_j} \implies \mathbf{x} = (T - \mu I)^m \left((T - \mu I)_{K_{\lambda_j}}^{-1} \right)^m (\mathbf{x}) \in \mathcal{R}(T - \mu I)^m = W$$

We conclude that λ_j is an eigenvalue of the restriction T_W with generalized eigenspace K_{λ_j}.

Since T_W has k distinct eigenvalues, the induction hypotheses apply:

$$W = K_{\lambda_1} \oplus \cdots \oplus K_{\lambda_k} \quad \text{and} \quad p_W(t) = (\lambda_1 - t)^{\dim K_{\lambda_1}} \cdots (\lambda_k - t)^{\dim K_{\lambda_k}}$$

Since $W \cap K_\mu = \{\mathbf{0}\}$ it is enough finally to use the rank–nullity theorem and count dimensions:

$$\dim V = \text{rank}(T - \mu I)^m + \text{null}(T - \mu I)^m = \dim W + \dim K_\mu = \sum_{j=1}^{k} \dim K_{\lambda_j} + \dim K_\mu$$

$$\overset{(*)}{\le} m_1 + \cdots + m_k + m = \deg(p(t)) = \dim V$$

The inequality is thus an *equality*; each $\dim K_{\lambda_j} = m_j$ and $\dim K_\mu = m$. We conclude that

$$V = K_{\lambda_1} \oplus \cdots \oplus K_{\lambda_k} \oplus K_\mu$$

which completes the induction step and thus the proof. Whew! ∎

---

[16]This is yet another argument where we consider a suitable subspace to which we can apply an induction hypothesis; recall the spectral theorem, Schur's lemma, bilinear form diagonalization, etc. Theorem 3.12 will provide one more!

Cycles of Generalized Eigenvectors

By Theorem 3.5, for every linear map whose characteristic polynomial splits there exists generalized eigenbasis. This isn't the same as a Jordan canonical basis, but we're very close!

Example 3.8. The matrix $A = \left(\begin{smallmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{smallmatrix}\right) \in M_3(\mathbb{R})$ is a single Jordan block, whence there is a single generalized eigenspace $K_5 = \mathbb{R}^3$ and the standard basis $\epsilon = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is Jordan canonical.

The crucial observation for what follows is that one of these vectors $\mathbf{e}_3$ *generates* the others via repeated applications of $A - 5I$:

$$\mathbf{e}_2 = (A - 5I)\mathbf{e}_3, \qquad \mathbf{e}_1 = (A - 5I)\mathbf{e}_2 = (A - 5I)^2\mathbf{e}_3$$

Definition 3.9. A *cycle of generalized eigenvectors* for a linear operator T is a set

$$\beta_{\mathbf{x}} := \left\{ (T - \lambda I)^{k-1}(\mathbf{x}), \dots, (T - \lambda I)(\mathbf{x}), \mathbf{x} \right\}$$

where the *generator* $\mathbf{x} \in K_\lambda$ is non-zero and k is minimal such that $(T - \lambda I)^k(\mathbf{x}) = \mathbf{0}$.

Note that the first element $(T - \lambda I)^{k-1}(\mathbf{x})$ is an *eigenvector*.
Our goal is to show that K_λ has a basis consisting of cycles of generalized eigenvectors; putting these together results in a Jordan canonical basis.

Lemma 3.10. *Let $\beta_{\mathbf{x}}$ be a cycle of generalized eigenvectors of T with length k. Then:*

1. *$\beta_{\mathbf{x}}$ is linearly independent and thus a basis of $\text{Span}\,\beta_{\mathbf{x}}$.*

2. *$\text{Span}\,\beta_{\mathbf{x}}$ is T-invariant. With respect to $\beta_{\mathbf{x}}$, the matrix of the restriction of T is the $k \times k$ Jordan block $[T_{\text{Span}\,\beta_{\mathbf{x}}}]_{\beta_{\mathbf{x}}} = \begin{pmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{pmatrix}$.*

In what follows, it will be useful to consider the linear map $U = T - \lambda I$. Note the following:

- The nullspace of U is the *eigenspace*: $\mathcal{N}(U) = E_\lambda \leq K_\lambda$.

- T commutes with U: that is $TU = UT$.

- $\beta_{\mathbf{x}} = \{U^{k-1}(\mathbf{x}), \dots, U(\mathbf{x}), \mathbf{x}\}$; that is, $\text{Span}\,\beta_{\mathbf{x}} = \langle \mathbf{x} \rangle$ is the U-cyclic subspace generated by $\mathbf{x}$.

Proof. 1. Feed the linear combination $\sum_{j=0}^{k-1} a_j U^j(\mathbf{x}) = \mathbf{0}$ to U^{k-1} to obtain

$$a_0 U^{k-1}(\mathbf{x}) = \mathbf{0} \implies a_0 = 0$$

Now feed the same combination to U^{k-2}, etc., to see that all coefficients $a_j = 0$.

2. Since T and U commute, we see that

$$T(U^j(\mathbf{x})) = U^j(T(\mathbf{x})) = U^j((U + \lambda I)(\mathbf{x})) = U^{j+1}(\mathbf{x}) + \lambda U^j(\mathbf{x}) \in \text{Span}\,\beta_{\mathbf{x}}$$

This justifies both T-invariance and the Jordan block claim! ∎

71

The basic approach to finding a Jordan canonical basis is to find the generalized eigenspaces and play with cycles until you find a basis for each K_λ. *Many* choices of canonical basis exist for a given map! We'll consider a more systematic method in the next section.

Examples 3.11. 1. The characteristic polynomial of $A = \left(\begin{smallmatrix} 1 & 0 & 2 \\ 0 & 1 & 6 \\ 6 & -2 & 1 \end{smallmatrix}\right) \in M_3(\mathbb{R})$ splits:

$$p(t) = (1-t)\begin{vmatrix} 1-t & 6 \\ -2 & 1-t \end{vmatrix} + 2\begin{vmatrix} 0 & 1-t \\ 6 & -2 \end{vmatrix} = (1-t)\left((1-t)^2 + 12 - 12\right) = (1-t)^3$$

With only one eigenvalue we see that $K_1 = \mathbb{R}^3$. Simply choose any vector in $\mathbb{R}^3$ and see what $U = A - I$ does to it! For instance, with $\mathbf{x} = \mathbf{e}_1$,

$$\beta_\mathbf{x} = \left\{U^2\left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix}\right), U\left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix}\right)\right\} = \left\{\left(\begin{smallmatrix} 12 \\ 36 \\ 0 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 0 \\ 6 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 \\ 0 \\ 0 \end{smallmatrix}\right)\right\}$$

provides a Jordan canonical basis of $\mathbb{R}^3$. We conclude

$$A = QJQ^{-1} = \left(\begin{smallmatrix} 12 & 0 & 1 \\ 36 & 0 & 0 \\ 0 & 6 & 0 \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{smallmatrix}\right)\left(\begin{smallmatrix} 12 & 0 & 1 \\ 36 & 0 & 0 \\ 0 & 6 & 0 \end{smallmatrix}\right)^{-1}$$

In practice, almost any choice of $\mathbf{x} \in \mathbb{R}^3$ will generate a cycle of length three!

2. The matrix $B = \left(\begin{smallmatrix} 7 & 1 & -4 \\ 0 & 3 & 0 \\ 8 & 1 & -5 \end{smallmatrix}\right) \in M_3(\mathbb{R})$ has characteristic equation

$$p(t) = (3-t)(t^2 - 2t - 3) = -(t+1)^1(t-3)^2$$

$\dim K_{-1} = 1 \implies K_{-1} = E_{-1} = \mathrm{Span}\left(\begin{smallmatrix} 1 \\ 0 \\ 2 \end{smallmatrix}\right)$, spanned by a cycle of length one.

Since $\dim K_3 = 2$, we have

$$K_3 = \mathcal{N}(B - 3I)^2 = \mathcal{N}\left(\begin{smallmatrix} 4 & 1 & -4 \\ 0 & 0 & 0 \\ 8 & 1 & -8 \end{smallmatrix}\right)^2 = \mathcal{N}\left(\begin{smallmatrix} -16 & 0 & 16 \\ 0 & 0 & 0 \\ -32 & 0 & 32 \end{smallmatrix}\right) = \mathrm{Span}\left\{\left(\begin{smallmatrix} 1 \\ 0 \\ 1 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 1 \\ 0 \end{smallmatrix}\right)\right\}$$

This is spanned by a cycle of length two: $\left(\begin{smallmatrix} 1 \\ 0 \\ 1 \end{smallmatrix}\right)$ is an eigenvector and

$$\left(\begin{smallmatrix} 1 \\ 0 \\ 1 \end{smallmatrix}\right) = (B - 3I)\left(\begin{smallmatrix} 0 \\ 1 \\ 0 \end{smallmatrix}\right)$$

We conclude that $\beta = \left\{\left(\begin{smallmatrix} 1 \\ 0 \\ 2 \end{smallmatrix}\right), \left(\begin{smallmatrix} 1 \\ 0 \\ 1 \end{smallmatrix}\right), \left(\begin{smallmatrix} 0 \\ 1 \\ 0 \end{smallmatrix}\right)\right\}$ is a Jordan canonical basis for B, and that

$$B = QJQ^{-1} = \left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 2 & 1 & 0 \end{smallmatrix}\right)\left(\begin{smallmatrix} -1 & 0 & 0 \\ 0 & 3 & 1 \\ 0 & 0 & 3 \end{smallmatrix}\right)\left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \\ 2 & 1 & 0 \end{smallmatrix}\right)^{-1}$$

3. Let $T = \frac{d}{dx}$ on $P_3(\mathbb{R})$. With respect to the standard basis $\epsilon = \{1, x, x^2, x^3\}$,

$$A = [T]_\epsilon = \left(\begin{smallmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{smallmatrix}\right)$$

With only one eigenvalue $\lambda = 0$, we have a single generalized eigenspace $K_0 = P_3(\mathbb{R})$. It is easy to check that $f(x) = x^3$ generates a cycle of length three and thus a Jordan canonical basis:

$$\beta = \{6, 6x, 3x^2, x^3\} \implies [T]_\beta = \left(\begin{smallmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{smallmatrix}\right) = \left(\begin{smallmatrix} 6 & 0 & 0 & 0 \\ 0 & 6 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 1 \end{smallmatrix}\right)^{-1}\left(\begin{smallmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 \end{smallmatrix}\right)\left(\begin{smallmatrix} 6 & 0 & 0 & 0 \\ 0 & 6 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 1 \end{smallmatrix}\right)$$

Our final results state that this process works generally.

> **Theorem 3.12.** Let $T \in \mathcal{L}(V)$ have an eigenvalue λ. If $\dim K_\lambda < \infty$, then there exists a basis $\beta_\lambda = \beta_{\mathbf{x}_1} \cup \cdots \cup \beta_{\mathbf{x}_n}$ of K_λ consisting of finitely many linearly independent cycles.

Intuition suggests that we create cycles $\beta_{\mathbf{x}_j}$ by starting with a basis of the eigenspace E_λ and extending backwards: for each $\mathbf{x}$, if $\mathbf{x} = (T - \lambda I)(\mathbf{y})$, then $\mathbf{x} \in \beta_{\mathbf{y}}$; now repeat until you have a maximum length cycle. This is essentially what we do, though a sneaky induction is required to make sure we keep track of everything and guarantee that the result really is a basis of K_λ.

Proof. We prove by induction on $m = \dim K_\lambda$.

(*Base case*)   If $m = 1$, then $K_\lambda = E_\lambda = \operatorname{Span} \mathbf{x}$ for some eigenvector $\mathbf{x}$. Plainly $\{\mathbf{x}\} = \beta_{\mathbf{x}}$.

(*Induction step*)   Fix $m \geq 2$. Write $n = \dim E_\lambda \leq m$ and $U = (T - \lambda I)_{K_\lambda}$.

  (i) For the induction hypothesis, suppose *every* generalized eigenspace with dimension $< m$ (for *any* linear map!) has a basis consisting of independent cycles of generalized eigenvectors.

 (ii) Define $W = \mathcal{R}(U) \cap E_\lambda$: that is

$$\mathbf{w} \in W \iff \begin{cases} U(\mathbf{w}) = \mathbf{0} \quad \text{and} \\ \mathbf{w} = U(\mathbf{v}) \text{ for some } \mathbf{v} \in K_\lambda \end{cases}$$

Let $k = \dim W$, choose a complementary subspace X such that $E_\lambda = W \oplus X$ and select a basis $\{\mathbf{x}_{k+1}, \ldots, \mathbf{x}_n\}$ of X. If $k = 0$, the induction step is finished (why?). Otherwise we continue...

(iii) The calculation in the proof of Lemma 3.10 (take $j = 1$) shows that $\mathcal{R}(U)$ is T-invariant; it is therefore the single generalized eigenspace $\tilde{K}_\lambda$ of $T_{\mathcal{R}(U)}$.

(iv) By the rank–nullity theorem,

$$\dim \mathcal{R}(U) = \operatorname{rank} U = \dim K_\lambda - \operatorname{null} U = m - \dim E_\lambda < m$$

By the induction hypothesis, $\mathcal{R}(U)$ has a basis of independent cycles. Since the last non-zero element in each cycle is an eigenvector, this basis consists of k distinct cycles $\beta_{\hat{\mathbf{x}}_1} \cup \cdots \cup \beta_{\hat{\mathbf{x}}_k}$ whose terminal vectors form a basis of W.

 (v) Since each $\hat{\mathbf{x}}_j \in \mathcal{R}(U)$, there exist vectors $\mathbf{x}_1, \ldots, \mathbf{x}_k$ such that $\hat{\mathbf{x}}_j = U(\mathbf{x}_j)$. Including the length-one cycles generated by the basis of X, the cycles $\beta_{\mathbf{x}_1}, \ldots, \beta_{\mathbf{x}_n}$ now contain

$$\dim \mathcal{R}(U) + k + (n - k) = \operatorname{rank} U + \operatorname{null} U = m$$

vectors. We leave as an exercise the verification that these vectors are linearly independent. ∎

> **Corollary 3.13.** Suppose that the characteristic polynomial of $T \in \mathcal{L}(V)$ splits (necessarily $\dim V < \infty$). Then there exists a Jordan canonical basis, namely the union of bases β_λ from Theorem 3.12.

Proof. By Theorem 3.5, V is the direct sum of generalized eigenspaces. By the previous result, each K_λ has a basis β_λ consisting of finitely many cycles. By Lemma 3.10, the matrix of T_{K_λ} has Jordan canonical form with respect to β_λ. It follows that $\beta = \bigcup \beta_\lambda$ is a Jordan canonical basis for T. ∎

Exercises 3.1 1. For each matrix, find the generalized eigenspaces K_λ, find bases consisting of unions of disjoint cycles of generalized eigenvectors, and thus find a Jordan canonical form J and invertible Q so that the matrix may be expressed as QJQ^{-1}.

(a) $A = \begin{pmatrix} 1 & 1 \\ -1 & 3 \end{pmatrix}$ (b) $B = \begin{pmatrix} 1 & 2 \\ 3 & 2 \end{pmatrix}$ (c) $C = \begin{pmatrix} 11 & -4 & -5 \\ 21 & -8 & -11 \\ 3 & -1 & 0 \end{pmatrix}$

(d) $D = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 0 & 2 & 1 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 1 & -1 & 3 \end{pmatrix}$

2. If $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ is a Jordan canonical basis, what can you say about $\mathbf{v}_1$? Briefly explain why the linear map $L_A \in \mathcal{L}(\mathbb{R}^2)$ where $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ has no Jordan canonical form.

3. Find a Jordan canonical basis for each linear map T:

(a) $T \in \mathcal{L}(P_2(\mathbb{R}))$ defined by $T(f(x)) = 2f(x) - f'(x)$
(b) $T(f) = f'$ defined on $\mathrm{Span}\{1, t, t^2, e^t, te^t\}$
(c) $T(A) = 2A + A^T$ defined on $M_2(\mathbb{R})$

4. In Example 3.11.1, suppose $\mathbf{x} = \begin{pmatrix} a \\ b \\ c \end{pmatrix}$. Show that almost any choice of a, b, c produces a Jordan canonical basis $\beta_\mathbf{x}$.

5. We complete the proof of Lemma 3.7.

(a) Prove part 1: that $E_\lambda \leq K_\lambda \leq V$.
(b) Verify that $T - \mu I$ and $T - \lambda I$ commute.
(c) Prove part 3(b): generalized eigenspaces for distinct eigenvalues have trivial intersection.

6. Consider the induction step in the proof of Theorem 3.5.

(a) Prove that W is T-invariant.
(b) Explain why $W \cap K_\mu = \{\mathbf{0}\}$.
(c) The assumption $p_W(t) = (\lambda_1 - t)^{\dim K_{\lambda_1}} \cdots (\lambda_k - t)^{\dim K_{\lambda_k}}$ near the end of the proof is the induction hypothesis for part 1(b). Why can't we also assume that $\dim K_{\lambda_j} = m_j$ and thus tidy the inequality argument near the end of the proof?

7. We finish some of the details of Theorem 3.12.

(a) In step (ii), suppose $\dim W = k = 0$. Explain why $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is in fact a basis of K_λ, so that the rest of the proof is unnecessary.
(b) In step (v), prove that the m vectors in the cycles $\beta_{\mathbf{x}_1}, \dots, \beta_{\mathbf{x}_n}$ are linearly independent.

   (*Hint: model your argument on part 1 of Lemma 3.10*)

3.2 Cycle Patterns and the Dot Diagram

In this section we obtain a useful result that helps us compute Jordan forms more efficiently and systematically. To give us some clues how to proceed, here is a lengthy example.

Example 3.14. Precisely three Jordan canonical forms $A, B, C \in M_3(\mathbb{R})$ correspond to the characteristic polynomial $p(t) = (5 - t)^3$:

$$A = \begin{pmatrix} 5 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 5 \end{pmatrix} \qquad B = \begin{pmatrix} 5 & 1 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & 5 \end{pmatrix} \qquad C = \begin{pmatrix} 5 & 1 & 0 \\ 0 & 5 & 1 \\ 0 & 0 & 5 \end{pmatrix}$$

In all three cases the standard basis $\beta = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ is Jordan canonical, so how do we distinguish things? By considering the *number* and *lengths* of the cycles of generalized eigenvectors.

- A has eigenspace $E_5 = K_5 = \mathbb{R}^3$. Since $(A - 5I)\mathbf{v} = \mathbf{0}$ for all $\mathbf{v} \in \mathbb{R}^3$, we have maximum cycle-length one. We therefore need three distinct cycles to construct a Jordan basis, e.g.

$$\beta_{\mathbf{e}_1} = \{\mathbf{e}_1\}, \quad \beta_{\mathbf{e}_2} = \{\mathbf{e}_2\}, \quad \beta_{\mathbf{e}_3} = \{\mathbf{e}_3\} \implies \beta = \beta_{\mathbf{e}_1} \cup \beta_{\mathbf{e}_2} \cup \beta_{\mathbf{e}_3} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$$

- B has eigenspace $E_5 = \mathrm{Span}\{\mathbf{e}_1, \mathbf{e}_3\}$. By computing

$$\mathbf{v} = \begin{pmatrix} a \\ b \\ c \end{pmatrix} \implies (B - 5I)\mathbf{v} = \begin{pmatrix} b \\ 0 \\ 0 \end{pmatrix} \implies (B - 5I)^2\mathbf{v} = \mathbf{0}$$

we see that $\beta_{\mathbf{v}}$ is a cycle with maximum length two, provided $b \neq 0$ ($\mathbf{v} \notin E_5$). We therefore need two distinct cycles, of lengths two and one, to construct a Jordan basis, e.g.

$$\beta_{\mathbf{e}_2} = \{(B - 5I)\mathbf{e}_2, \mathbf{e}_2\} = \{\mathbf{e}_1, \mathbf{e}_2\}, \quad \beta_{\mathbf{e}_3} = \{\mathbf{e}_3\} \implies \beta = \beta_{\mathbf{e}_2} \cup \beta_{\mathbf{e}_3} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$$

- C has eigenspace $E_5 = \mathrm{Span}\,\mathbf{e}_1$. This time

$$\mathbf{v} = \begin{pmatrix} a \\ b \\ c \end{pmatrix} \implies (C - 5I)\mathbf{v} = \begin{pmatrix} b \\ c \\ 0 \end{pmatrix}, \quad (C - 5I)^2\mathbf{v} = \begin{pmatrix} c \\ 0 \\ 0 \end{pmatrix}, \quad (C - 5I)^3\mathbf{v} = \mathbf{0}$$

generates a cycle with maximum length two provided $c \neq 0$. Indeed this cycle is a Jordan basis, so one cycle is all we need:

$$\beta = \beta_{\mathbf{e}_3} = \{(C - 5I)^2\mathbf{e}_3, (C - 5I)\mathbf{e}_3, \mathbf{e}_3\} = \{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$$

Why is the example relevant? Suppose that $\dim_{\mathbb{R}} V = 3$ and that $T \in \mathcal{L}(V)$ has characteristic polynomial $p(t) = (5 - t)^3$. Theorem 3.12 tells us that T has a Jordan canonical form, and that is is moreover one of the above matrices A, B, C. Our goal is to develop a method whereby the pattern of cycle-lengths can be determined, thus allowing us to be able to discern which Jordan form is correct. As a side-effect, this will also demonstrate that the pattern of cycle lengths for a given T is independent of the Jordan basis so that, up to some reasonable restriction, the Jordan *form* of T is unique. To aid us in this endeavor, we require some terminology...

> **Definition 3.15.** Let V be finite dimensional and K_λ a generalized eigenspace of $T \in \mathcal{L}(V)$. Following the Theorem 3.12, assume that $\beta_\lambda = \beta_{\mathbf{x}_1} \cup \cdots \cup \beta_{\mathbf{x}_n}$ is a Jordan canonical basis of T_{K_λ}, where the cycles are arranged in *non-increasing* length. That is:
>
> 1. $\beta_{\mathbf{x}_j} = \{(T - \lambda I)^{k_j - 1}(\mathbf{x}_j), \dots, \mathbf{x}_j\}$ has length k_j, and
>
> 2. $k_1 \geq k_2 \geq \cdots \geq k_n$
>
> The *dot diagram* of T_{K_λ} is a representation of the elements of β_λ, one dot for each vector: the j^{th} column represents the elements of $\beta_{\mathbf{x}_j}$ arranged vertically with $\mathbf{x}_j$ at the bottom.

Given a linear map, our eventual goal is to identify the dot diagram as an intermediate step in the computation of a Jordan basis. First, however, we observe how the conversion of dot diagrams to a Jordan form is essentially trivial.

Example 3.16. Suppose $\dim V = 14$ and that $T \in \mathcal{L}(V)$ has the following eigenvalues and dot diagrams:

$$\lambda_1 = -4 \qquad \lambda_2 = 7 \qquad \lambda_3 = 12$$

Then generalized eigenspaces of T satisfy:

- $K_{-4} = \mathcal{N}(T + 4I)^2$ and $\dim K_{-4} = 6$;

- $K_7 = \mathcal{N}(T - 7I)^3$ and $\dim K_7 = 5$;

- $K_{12} = \mathcal{N}(T - 12I) = E_{12}$ and $\dim K_{12} = 3$;

T has a Jordan canonical basis β with respect to which its Jordan canonical form is

$$[T]_\beta = \begin{pmatrix} -4 & 1 & & & & & & & & & & & & \\ 0 & -4 & & & & & & & & & & & & \\ & & -4 & 1 & & & & & & & & & & \\ & & 0 & -4 & & & & & & & & & & \\ & & & & -4 & & & & & & & & & \\ & & & & & -4 & & & & & & & & \\ & & & & & & 7 & 1 & 0 & & & & & \\ & & & & & & 0 & 7 & 1 & & & & & \\ & & & & & & 0 & 0 & 7 & & & & & \\ & & & & & & & & & 7 & 1 & & & \\ & & & & & & & & & 0 & 7 & & & \\ & & & & & & & & & & & 12 & & \\ & & & & & & & & & & & & 12 & \\ & & & & & & & & & & & & & 12 \end{pmatrix}$$

Note how the sizes of the Jordan blocks are *non-increasing* within each eigenvalue. For instance, for $\lambda_1 = -4$, the sequence of cycle lengths (k_j) is $2 \geq 2 \geq 1 \geq 1$.

> **Theorem 3.17.** *Suppose β_λ is a Jordan canonical basis of T_{K_λ} as described in Definition 3.15, and suppose the i^{th} row of the dot diagram has r_i entries. Then:*
>
> 1. *For each $r \in \mathbb{N}$, the vectors associated to the dots in the first r rows form a basis of $\mathcal{N}(T - \lambda I)^r$.*
>
> 2. *$r_1 = \text{null}(T - \lambda I) = \dim V - \text{rank}(T - \lambda I)$*
>
> 3. *When $i > 1$, $r_i = \text{null}(T - \lambda I)^i - \text{null}(T - \lambda I)^{i-1} = \text{rank}(T - \lambda I)^{i-1} - \text{rank}(T - \lambda I)^i$*

Example (3.14 cont). We describe the dot diagrams of the three matrices A, B, C, along with the corresponding vectors in the Jordan canonical basis β and the values r_i.

$A:$ •   •   • x_1 x_2 x_3 e_1 e_2 e_3

     Since $A - 5I$ is the zero matrix, $r_1 = 3 - \text{rank}(A - 5I) = 3$. The dot diagram has one row, corresponding to three independent cycles of length one: $\beta = \beta_{e_1} \cup \beta_{e_2} \cup \beta_{e_3}$.

$B:$ •   • $(B - 5I)x_1$ x_2 e_1 e_3
      • x_1 e_2

     Row 1: $B - 5I = \left(\begin{smallmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{smallmatrix}\right) \implies \text{rank}(B - 5I) = 1$ and $r_1 = 3 - 1 = 2$. The first row $\{e_1, e_3\}$ is a basis of $E_5 = \mathcal{N}(B - 5I)$.

     Row 2: $(B - 5I)^2$ is the zero matrix, whence $r_2 = \text{rank}(B - 5I) - \text{rank}(B - 5I)^2 = 1 - 0 = 1$.

     The dot diagram corresponds to $\beta = \beta_{e_2} \cup \beta_{e_3} = \{e_1, e_2\} \cup \{e_3\}$.

$C:$ • $(C - 5I)^2 x_1$ e_1
     • $(C - 5I)x_1$ e_2
     • x_1 e_3

     Row 1: $C - 5I = \left(\begin{smallmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{smallmatrix}\right) \implies r_1 = 3 - \text{rank}(C - 5I) = 1$. The first row $\{e_1\}$ is a basis of $E_5 = \mathcal{N}(C - 5I)$.

     Row 2: $(C - 5I)^2 = \left(\begin{smallmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{smallmatrix}\right) \implies r_2 = \text{rank}(C - 5I) - \text{rank}(C - 5I)^2 = 2 - 1 = 1$. The first two rows $\{e_1, e_2\}$ form a basis of $\mathcal{N}(C - 5I)^2$.

     Row 3: $(C - 5I)^3$ is the zero matrix, whence $r_3 = \text{rank}(C - 5I)^2 - \text{rank}(C - 5I)^3 = 1 - 0 = 1$.

Proof. As previously, let $U = T - \lambda I$.

1. Since each dot represents a basis vector $U^p(v_j)$, any $v \in K_\lambda$ may be written uniquely as a linear combination of the dots. Applying U simply moves all the dots up a row and all dots in the top row to 0. It follows that $v \in \mathcal{N}(U^r) \iff$ it lies in the span of the first r rows. Since the dots are linearly independent, they form a basis.

2. By part 1, $r_1 = \dim \mathcal{N}(U) = \text{null}(T - \lambda I) = \dim V - \text{rank}(T - \lambda I)$.

3. More generally,

$$r_i = (r_1 + \cdots + r_i) - (r_1 + \cdots + r_{i-1}) = \dim \mathcal{N}(U^i) - \dim \mathcal{N}(U^{i-1})$$
$$= \text{null}(U^i) - \text{null}(U^{i-1}) = \text{rank}(T - \lambda I)^{i-1} - \text{rank}(T - \lambda I)^i$$

∎

Since the ranks of maps $(T - \lambda I)^i$ are independent of basis, so also is the dot diagram…

> **Corollary 3.18.** *For any eigenvalue λ, the dot diagram is uniquely determined by T and λ. If we list Jordan blocks for each eigenspace in non-increasing order, then the Jordan form of a linear map is unique up to the order of the eigenvalues.*

We now have a slightly more systematic method for finding Jordan canonical bases.

Example 3.19. The matrix $A = \begin{pmatrix} 6 & 2 & -4 & -6 \\ 0 & 3 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 2 & 1 & -2 & -1 \end{pmatrix}$ has characteristic equation

$$p(t) = (3-t)^2 \begin{vmatrix} 6-t & -6 \\ 2 & -1-t \end{vmatrix} = (2-t)(3-t)^3$$

We have two generalized eigenspaces:

- $K_2 = E_2 = \mathcal{N}(A - 2I) = \mathcal{N}\begin{pmatrix} 4 & 2 & -4 & -6 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 2 & 1 & -2 & -3 \end{pmatrix} = \text{Span}\begin{pmatrix} 3 \\ 0 \\ 0 \\ 2 \end{pmatrix}$. The trivial dot diagram $\bullet$ corresponds to this single eigenvector.

- $K_3 = \mathcal{N}(A - 3I)^3$. To find the dot diagram, compute powers of $A - 3I$:

  Row 1: $A - 3I = \begin{pmatrix} 3 & 2 & -4 & -6 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 2 & 1 & -2 & -4 \end{pmatrix}$ has rank 2 and the first row has $r_1 = 4 - 2 = 2$ entries.

  Row 2: $(A - 3I)^2 = \begin{pmatrix} -3 & 0 & 0 & 6 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -2 & 0 & 0 & 4 \end{pmatrix}$ has rank 1 and the second row has $r_2 = 2 - 1 = 1$ entry.

  Since we now have three dots (equalling dim K_3), the algorithm terminates and the dot diagram for K_3 is $\bullet$ $\bullet$
  $\bullet$

  For the single dot in the second row, we choose something in $\mathcal{N}(A - 3I)^2$ which isn't an eigenvector; perhaps the simplest choice is $\mathbf{x}_1 = \mathbf{e}_2$, which yields the two-cycle

  $$\beta_{\mathbf{x}_1} = \{(A - 3I)\mathbf{x}_1, \mathbf{x}_1\} = \left\{ \begin{pmatrix} 2 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} \right\}$$

  To complete the first row, choose any eigenvector to complete the span: for instance $\mathbf{x}_2 = \begin{pmatrix} 0 \\ 2 \\ 1 \\ 0 \end{pmatrix}$.

We now have suitable cycles and a Jordan canonical basis/form:

$$\beta = \left\{ \begin{pmatrix} 3 \\ 0 \\ 0 \\ 2 \end{pmatrix}, \begin{pmatrix} 2 \\ 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 2 \\ 1 \\ 0 \end{pmatrix} \right\}, \qquad A = QJQ^{-1} = \begin{pmatrix} 3 & 2 & 0 & 0 \\ 0 & 0 & 1 & 2 \\ 0 & 0 & 0 & 1 \\ 2 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} 3 & 2 & 0 & 0 \\ 0 & 0 & 1 & 2 \\ 0 & 0 & 0 & 1 \\ 2 & 1 & 0 & 0 \end{pmatrix}^{-1}$$

Other choices are available! For instance, if we'd chosen the two-cycle generated by $\mathbf{x}_1 = \mathbf{e}_3$, we'd obtain a different Jordan basis but the same canonical form J:

$$\tilde{\beta} = \left\{ \begin{pmatrix} 3 \\ 0 \\ 0 \\ 2 \end{pmatrix}, \begin{pmatrix} -4 \\ 0 \\ 0 \\ -2 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 2 \\ 1 \\ 0 \end{pmatrix} \right\}, \qquad A = \begin{pmatrix} 3 & -4 & 0 & 0 \\ 0 & 0 & 0 & 2 \\ 0 & 0 & 1 & 1 \\ 2 & -2 & 0 & 0 \end{pmatrix} \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} 3 & -4 & 0 & 0 \\ 0 & 0 & 0 & 2 \\ 0 & 0 & 1 & 1 \\ 2 & -2 & 0 & 0 \end{pmatrix}^{-1}$$

We do one final example for a non-matrix map.

Example 3.20. Let $\epsilon = \{1, x, y, x^2, y^2, xy\}$ and define $T(f(x,y)) = 2\frac{\partial f}{\partial x} - \frac{\partial f}{\partial y}$ as a linear operator on $V = \mathrm{Span}_{\mathbb{R}}\,\epsilon$. The matrix and characteristic polynomial of T is easy to compute:

$$[T]_\epsilon = \begin{pmatrix} 0 & 2 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 & -1 \\ 0 & 0 & 0 & 0 & -2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \implies p(t) = t^6, \quad [T^2]_\epsilon = \begin{pmatrix} 0 & 0 & 0 & 8 & 2 & -4 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad [T^3]_\epsilon = O$$

There is only one eigenvalue $\lambda = 0$ and therefore one generalized eigenspace $K_0 = V$. We could keep working with matrices, but it is easy to translate the nullspaces of the matrices back to subspaces of V, from which the necessary data can be read off:

$$\mathcal{N}(T) = \mathrm{Span}\{1, x + 2y, x^2 + 4y^2 + 4xy\} \qquad \text{null}\,T = 3,\ \text{rank}\,T = 3,\ r_1 = 3$$
$$\mathcal{N}(T^2) = \mathrm{Span}\{1, x, y, x^2 + 2xy, 2y^2 + xy\} \qquad \text{null}\,T^2 = 5,\ \text{rank}\,T^2 = 1,\ r_2 = 3 - 1 = 2$$

We now have five dots; since $\dim K_0 = 6$, the last row has one, and the dot diagram is
$$\begin{matrix} \bullet & \bullet & \bullet \\ \bullet & \bullet & \\ \bullet & & \end{matrix}$$
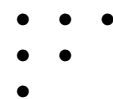
Since the first two rows span $\mathcal{N}(T^2)$, we may choose any $f_1 \notin \mathcal{N}(T^2)$ for the final dot: $f_1 = xy$ is suitable, from which the first column of the dot diagram becomes

$$\begin{matrix} T^2(xy) & \bullet & \bullet & & -4 & \bullet & \bullet \\ T(xy) & \bullet & & & 2y - x & \bullet & \\ xy & & & & xy & & \end{matrix}$$

Now choose the second dot on the second row to be anything in $\mathcal{N}(T^2)$ such that the first two rows span $\mathcal{N}(T^2)$: this time $f_2 = x^2 - 4y^2$ is suitable, and the diagram becomes:

$$\begin{matrix} T^2(xy) & T(x^2 - 4y^2) & \bullet & & -4 & 4x + 8y & \bullet \\ T(xy) & x^2 - 4y^2 & & & 2y - x & x^2 - 4y^2 & \\ xy & & & & xy & & \end{matrix}$$

The final dot is now chosen so that the first row spans $\mathcal{N}(T)$: this time $f_3 = x^2 + 4y^2 + 4xy$ works. The result is a Jordan canonical basis and form for T

$$\beta = \{-4,\ 2y - x,\ xy,\ 4x + 8y,\ x^2 - 4y^2,\ x^2 + 4y^2 + 4xy\}, \qquad J = [T]_\beta = \begin{pmatrix} \begin{matrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{matrix} & \\ & \begin{matrix} 0 & 1 \\ 0 & 0 \end{matrix} & \\ & & 0 \end{pmatrix}$$

As previously, many other choices of cycle-generators f_1, f_2, f_3 are available; while these result in different Jordan canonical bases, Corollary 3.18 assures us that we'll always obtain the same canonical form J.

Exercises 3.2 1. Let T be a linear operator whose characteristic polynomial splits. Suppose the eigenvalues and the dot diagrams for the generalized eigenspaces K_{λ_i} are as follows:

$$\lambda_1 = 2 \qquad \lambda_2 = 4 \qquad \lambda_3 = -3$$

$$
\begin{array}{ccc}
\bullet\ \bullet\ \bullet & \bullet\ \bullet & \bullet\ \bullet \\
\bullet\ \bullet & \bullet & \\
\bullet & \bullet &
\end{array}
$$

Find the Jordan form J of T.

2. Suppose T has Jordan canonical form

$$
J = \begin{pmatrix}
\begin{array}{ccc}
2 & 1 & 0 \\
0 & 2 & 1 \\
0 & 0 & 2
\end{array} & & & \\
& \begin{array}{cc} 2 & 1 \\ 0 & 2 \end{array} & & \\
& & 3 & \\
& & & 3
\end{pmatrix}
$$

(a) Find the characteristic polynomial of T.

(b) Find the dot diagram for each eigenvalue.

(c) For each eigenvalue find the smallest k_j such that $K_{\lambda_j} = \mathcal{N}(\mathrm{T} - \lambda_j \mathrm{I})^{k_j}$.

3. For each matrix A find a Jordan canonical form and an invertible Q such that $A = QJQ^{-1}$.

(a) $A = \begin{pmatrix} -3 & 3 & -2 \\ -7 & 6 & -3 \\ 1 & -1 & 2 \end{pmatrix}$
(b) $A = \begin{pmatrix} 0 & 1 & -1 \\ -4 & 4 & -2 \\ -2 & 1 & 1 \end{pmatrix}$
(c) $A = \begin{pmatrix} 0 & -3 & 1 & 2 \\ -2 & 1 & -1 & 2 \\ -2 & 1 & -1 & 2 \\ -2 & -3 & 1 & 4 \end{pmatrix}$

4. For each linear operator T, find a Jordan canonical form J and basis β:

(a) $\mathrm{T}(f) = f'$ on $\mathrm{Span}_{\mathbb{R}}\{e^t, te^t, t^2 e^t, e^{2t}\}$
(b) $\mathrm{T}(f(x)) = xf''(x)$ on $P_3(\mathbb{R})$
(c) $\mathrm{T}(f) = af_x + bf_y$ on $\mathrm{Span}_{\mathbb{R}}\{1, x, y, x^2, y^2, xy\}$. How does your answer depend on a, b?

5. (Generalized Eigenvector Method for ODEs) Let $A \in M_n(\mathbb{R})$ have an eigenvalue λ and suppose $\beta_{\mathbf{v}_0} = \{\mathbf{v}_{k-1}, \dots, \mathbf{v}_1, \mathbf{v}_0\}$ is a cycle of generalized eigenvectors for this eigenvalue. Show that

$$\mathbf{x}(t) := e^{\lambda t} \sum_{j=0}^{k-1} b_j(t)\mathbf{v}_j \quad \text{satisfies} \quad \mathbf{x}'(t) = A\mathbf{x} \iff \begin{cases} b_0'(t) = 0, \text{ and} \\ b_j'(t) = b_{j-1}(t) \text{ when } j \geq 1 \end{cases}$$

Use this method to solve the system of differential equations

$$\mathbf{x}' = \begin{pmatrix} 3 & 1 & 0 & 0 \\ 0 & 3 & 1 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix} \mathbf{x}$$

3.3 The Rational Canonical Form (non-examinable)

We finish the course with a very quick discussion of what can be done when the characteristic polynomial of a linear map does not split. In such a situation, we may assume that

$$p(t) = (-1)^n \left(\phi_1(t)\right)^{m_1} \cdots \left(\phi_k(t)\right)^{m_k} \tag{$*$}$$

where each $\phi_j(t)$ is an *irreducible monic polynomial* over the field.

Example 3.21. The following matrix has characteristic equation $p(t) = (t^2 + 1)^2(3 - t)$

$$A = \begin{pmatrix} 0 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 3 \end{pmatrix} \in M_5(\mathbb{R})$$

This doesn't split over $\mathbb{R}$ since $t^2 + 1 = 0$ has no real roots. It is, however, diagonalizable over $\mathbb{C}$.

A couple of basic facts from algebra:

- Every polynomial splits over $\mathbb{C}$: every $A \in M_n(\mathbb{C})$ therefore has a Jordan form.
- Every polynomial over $\mathbb{R}$ factorizes into linear or irreducible quadratic factors.

The question is how to deal with non-linear irreducible factors in the characteristic polynomial.

Definition 3.22. The monic polynomial $t^k + a_{k-1}t^{k-1} + \cdots + a_0$ has *companion matrix*

$$\begin{pmatrix} 0 & 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & 0 & & 0 & -a_1 \\ 0 & 1 & 0 & & 0 & -a_2 \\ \vdots & & & \ddots & & \vdots \\ 0 & 0 & 0 & & 0 & -a_{k-2} \\ 0 & 0 & 0 & \cdots & 1 & -a_{k-1} \end{pmatrix} \qquad \text{(when } k = 1, \text{ this is the } 1 \times 1 \text{ matrix } (-a_0))$$

If $T \in \mathcal{L}(V)$ has characteristic polynomial $(*)$, then a *rational canonical basis* is a basis for which

$$[T]_\beta = \begin{pmatrix} C_1 & O & \cdots & O \\ O & C_2 & & O \\ \vdots & & \ddots & \vdots \\ O & O & \cdots & C_r \end{pmatrix}$$

where each C_j is a companion matrix of some $(\phi_j(t))^{s_j}$ where $s_j \le m_j$. We call $[T]_\beta$ a *rational canonical form* of T.

We state the main result without proof:

Theorem 3.23. *A rational canonical basis exists for any linear operator* T *on a finite-dimensional vector space* V. *The canonical form is unique up to ordering of companion matrices.*

Example (3.21 cont). The matrix A is *already in rational canonical form*: the standard basis is rational canonical with three companion blocks,

$$C_1 = C_2 = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad C_3 = (3)$$

Example 3.24. Let $A = \begin{pmatrix} 4 & -3 \\ 2 & 2 \end{pmatrix} \in M_2(\mathbb{R})$. Its characteristic polynomial

$$p(t) = t^2 - 6t + 14 = (t-3)^2 + 5$$

doesn't split over $\mathbb{R}$ and so it has no eigenvalues. Instead simply pick a vector, $\mathbf{x} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ (say), define $\mathbf{y} = A\mathbf{x} = \begin{pmatrix} 4 \\ 2 \end{pmatrix}$, let $\beta = \{\mathbf{x}, \mathbf{y}\}$ and observe that

$$[\mathsf{L}_A]_\beta = \begin{pmatrix} 0 & -14 \\ 1 & 6 \end{pmatrix}$$

is a rational canonical form. Indeed this works for *any* $\mathbf{x} \neq \mathbf{0}$: if $\beta := \{\mathbf{x}, A\mathbf{x}\}$, then Cayley–Hamilton forces

$$A^2\mathbf{x} = (6A - 14I)\mathbf{x} = -14\mathbf{x} + 6A\mathbf{x} \implies [\mathsf{L}_A]_\beta = \begin{pmatrix} 0 & -14 \\ 1 & 6 \end{pmatrix}$$

whence β is a rational canonical basis and the form $[\mathsf{L}_A]_\beta$ is independent of $\mathbf{x}$!

A systematic approach to finding rational canonical forms is similar to that for Jordan forms: for each irreducible divisor of $p(t)$, the subspace $K_\phi = \mathcal{N}(\phi(\mathsf{T}))^m$ plays a role analogous to a generalized eigenspace; indeed $K_\lambda = K_\phi$ for the *linear* irreducible factor $\phi(t) = \lambda - t$!

We finish with two examples; hopefully the approach is intuitive, even without theoretical justification.

Examples 3.25. If the characteristic polynomial of $\mathsf{T} \in \mathcal{L}(\mathbb{R}^4)$ is

$$p(t) = (\phi(t))^2 = (t^2 - 2t + 3)^2 = t^4 - 4t^3 + 10t^2 - 12t + 9$$

then there are two possible rational canonical forms; here is an example of each.

1. If $A = \begin{pmatrix} 0 & -15 & 0 & -9 \\ 2 & 2 & -3 & 0 \\ 0 & -9 & 0 & -6 \\ -3 & 0 & 5 & 2 \end{pmatrix}$, then $\phi(A) = O$ is the zero matrix, whence $\mathcal{N}(\phi(A)) = \mathbb{R}^4$. Since $\phi(t)$ isn't the full characteristic polynomial, we expect there to be *two* independent cycles of length two in the canonical basis. Start with something simple as a guess:

$$\mathbf{x}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} \implies \mathbf{x}_2 = A\mathbf{x}_1 = \begin{pmatrix} 0 \\ 2 \\ 0 \\ -3 \end{pmatrix} \implies A\mathbf{x}_2 = \begin{pmatrix} -3 \\ 4 \\ 0 \\ -6 \end{pmatrix} = -3\mathbf{x}_1 + 2\mathbf{x}_2$$

Now make another choice that isn't in the span of $\{\mathbf{x}_1, \mathbf{x}_2\}$:

$$\mathbf{x}_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix} \implies \mathbf{x}_4 = A\mathbf{x}_3 = \begin{pmatrix} 0 \\ -3 \\ 0 \\ 5 \end{pmatrix} \implies A\mathbf{x}_4 = \begin{pmatrix} 0 \\ -6 \\ -3 \\ 10 \end{pmatrix} = -3\mathbf{x}_3 + 2\mathbf{x}_4$$

We therefore have a rational canonical basis $\beta = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4\}$ and

$$A = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & -3 \\ 0 & 0 & 1 & 0 \\ 0 & -3 & 0 & 5 \end{pmatrix} \begin{pmatrix} 0 & -3 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 0 & 0 & -3 \\ 0 & 0 & 1 & 2 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & -3 \\ 0 & 0 & 1 & 0 \\ 0 & -3 & 0 & 5 \end{pmatrix}^{-1}$$

Over $\mathbb{C}$, this example is diagonalizable. Indeed each of the 2×2 companion matrices is diagonalizable over $\mathbb{C}$.

2. Let $B = \begin{pmatrix} 0 & 0 & 2 & 1 \\ 1 & 1 & -1 & -1 \\ 0 & 1 & -2 & -16 \\ 0 & 0 & 1 & 5 \end{pmatrix}$. This time

$$\phi(B) = B^2 - 2B + 3I = \begin{pmatrix} 3 & 2 & -7 & -29 \\ -1 & 1 & 4 & 13 \\ 1 & -3 & -6 & -17 \\ 0 & 1 & 1 & 2 \end{pmatrix} \implies \mathcal{N}(\phi(B)) = \text{Span}\left\{ \begin{pmatrix} 3 \\ -1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 11 \\ -2 \\ 0 \\ 1 \end{pmatrix} \right\}$$

Anything not in this span will suffice as a generator for a single cycle of length four: e.g.,

$$\mathbf{x}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{x}_2 = B\mathbf{x}_1 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad \mathbf{x}_3 = B\mathbf{x}_2 = \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \end{pmatrix}, \quad \mathbf{x}_4 = B\mathbf{x}_3 = \begin{pmatrix} 2 \\ 0 \\ -1 \\ 1 \end{pmatrix}$$

$$B\mathbf{x}_4 = \begin{pmatrix} -1 \\ 2 \\ -14 \\ 4 \end{pmatrix} = -9 \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix} + 12 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix} - 10 \begin{pmatrix} 0 \\ 1 \\ 1 \\ 0 \end{pmatrix} + 4 \begin{pmatrix} 2 \\ 0 \\ -1 \\ 1 \end{pmatrix}$$

We therefore have a rational canonical basis $\beta = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4\}$ and

$$B = \begin{pmatrix} 1 & 0 & 0 & 2 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 & -9 \\ 1 & 0 & 0 & 12 \\ 0 & 1 & 0 & -10 \\ 0 & 0 & 1 & 4 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 2 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{pmatrix}^{-1}$$

In contrast to the first example, B isn't diagonalizable over $\mathbb{C}$. It has Jordan form $J = \begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 0 & 0 \\ 0 & 0 & \bar{\lambda} & 1 \\ 0 & 0 & 0 & \bar{\lambda} \end{pmatrix}$ where $\lambda = 1 + i\sqrt{2}$.