

2 Inner Product Spaces

2.1 Real and Complex Inner Product Spaces

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

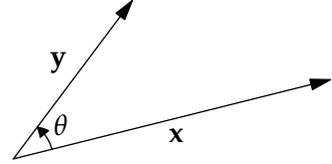
$$\mathbf{x} \cdot \mathbf{y} := x_1 y_1 + x_2 y_2$$

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 $\mathbf{x} \cdot \mathbf{y} = 0$.

As these formulae make clear, the dot product 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.



Definition 2.1. Suppose V is a vector space over \mathbb{F} ($= \mathbb{R}$ or \mathbb{C}). An *inner product* on V is a function $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{F}$ satisfying the following properties $\forall \mathbf{x}, \mathbf{y}, \mathbf{z} \in V, \lambda \in \mathbb{F}$,

(a) *Linear⁴ in the first slot:* $\langle \lambda \mathbf{x} + \mathbf{y}, \mathbf{z} \rangle = \lambda \langle \mathbf{x}, \mathbf{z} \rangle + \langle \mathbf{y}, \mathbf{z} \rangle$

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

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

We call $(V, \langle \cdot, \cdot \rangle)$ an *inner product space*.

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 plus 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 described as a *positive-definite, symmetric, bilinear form*.

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

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

Co-ordinates x_j, y_j are with respect to the standard basis $\epsilon = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$. We call $\|\mathbf{x}\|$ the *Euclidean norm*.

⁴**Warning!** In Physics, standard convention is to be *conjugate-linear* in the first slot: $\langle \lambda \mathbf{x} + \mathbf{y}, \mathbf{z} \rangle = \bar{\lambda} \langle \mathbf{x}, \mathbf{z} \rangle + \langle \mathbf{y}, \mathbf{z} \rangle$. If $\mathbb{F} = \mathbb{R}$ this makes no difference. The common Physics notation relates to ours via $\langle \mathbf{x} | \mathbf{y} \rangle = \langle \mathbf{y}, \mathbf{x} \rangle$.

Unless the inner product is stated explicitly, if we refer to \mathbb{R}^n as an inner product space, we mean Euclidean space. There are, however, many other ways to equip \mathbb{R}^n with an inner product...

Example 2.3. Define a new function on \mathbb{R}^2 via

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

It is easy to check that this satisfies the definition of an inner product:

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

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

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

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

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

So far, so good. What differs between $\langle \cdot, \cdot \rangle$ and the standard “dot” product is the notion of *orthogonality*. Consider, for instance, $\mathbf{x} = \frac{1}{2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\mathbf{y} = \frac{1}{2\sqrt{3}} \begin{pmatrix} 3 \\ -1 \end{pmatrix}$. With respect to $\langle \cdot, \cdot \rangle$, $\{\mathbf{x}, \mathbf{y}\}$ are *orthonormal*:

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

By contrast, there is nothing special about these vectors with respect to the standard dot product:

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

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

The example is what we call a *weighted inner product*: if we choose *weights* $a_1, \dots, a_n \in \mathbb{R}^+$, then

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

defines an inner product on \mathbb{R}^n . In this fashion, \mathbb{R}^n may be equipped with *infinitely many distinct inner products*. Even this idea generalizes:

Definition 2.4. A symmetric matrix $A \in M_n(\mathbb{R})$ is *positive-definite* if $\mathbf{x}^T A \mathbf{x} > 0$ for all non-zero $\mathbf{x} \in \mathbb{R}^n$.

If A is positive-definite, it is straightforward to check that $\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{y}^T A \mathbf{x}$ is an inner product on \mathbb{R}^n . In fact (Exercise 15) *all* inner products on \mathbb{R}^n arise in this fashion! The weighted inner products are when A is diagonal; Euclidean space has $A = I$.

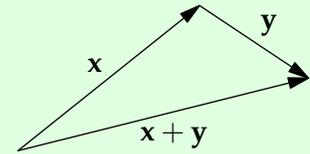
Example 2.5. The matrix $A = \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix}$ is positive-definite and thus defines an inner product

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

Lemma 2.6 (Basic properties). Let V be an inner product space, $\mathbf{x}, \mathbf{y}, \mathbf{z} \in V$ 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) $|\langle \mathbf{x}, \mathbf{y} \rangle| \leq \|\mathbf{x}\| \|\mathbf{y}\|$. Moreover, this is an **equality** if and only if \mathbf{x}, \mathbf{y} are parallel (linearly dependent).
6. (Triangle inequality) $\|\mathbf{x} + \mathbf{y}\| \leq \|\mathbf{x}\| + \|\mathbf{y}\|$. This is an **equality** if and only if \mathbf{x}, \mathbf{y} are parallel **and** point in the same direction.

As pictured, in \mathbb{R}^2 or \mathbb{R}^3 this simply says that any side of a triangle is no longer than the sum of the lengths of the other sides.



Be careful with notation: $|\lambda|$ is the *absolute value / modulus* of a scalar, whereas $\|\mathbf{x}\|$ is the *norm* of a vector.

In the real case, we may define the *angle* between non-zero vectors via $\cos \theta = \frac{\langle \mathbf{x}, \mathbf{y} \rangle}{\|\mathbf{x}\| \|\mathbf{y}\|}$ ($\in [-1, 1]$ by Cauchy–Schwarz!). Outwith Euclidean \mathbb{R}^2 and \mathbb{R}^3 , this notion is of limited application; elsewhere, *orthogonality* and *orthonormality* are typically all we need.

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 ($0\mathbf{x} + \mathbf{y} = \mathbf{0}$ is a linear dependence!). Otherwise,

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

Taking square-roots establishes the inequality. By part 2, equality forces $\|\mathbf{y}\|^2 \mathbf{x} = \langle \mathbf{x}, \mathbf{y} \rangle \mathbf{y}$: a linear dependence. Conversely, linear dependence (when $\mathbf{y} \neq \mathbf{0}$) means $\mathbf{x} = \lambda \mathbf{y}$ for some scalar λ : by part 3, the inequality becomes $|\lambda| \|\mathbf{y}\|^2 = |\lambda| \|\mathbf{y}\|^2$.

6. We establish the square of the required result.

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

Equality requires both equality in Cauchy–Schwarz (\mathbf{x}, \mathbf{y} parallel) and that $\langle \mathbf{x}, \mathbf{y} \rangle \geq 0$; 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 works perfectly when $\mathbb{C} = \mathbb{F}$. The subtle difference comes from how we expand linear combinations in the *second slot*. The proof is easy if you remember your complex conjugates; try it!

Lemma 2.7. A complex inner product is a positive-definite, conjugate-symmetric, **sesquilinear**⁵ form: it is conjugate-linear in the second slot,

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

When $\mathbb{F} = \mathbb{R}$ this is *bilinearity* (page 15): $\langle \mathbf{x}, \lambda \mathbf{y} + \mathbf{z} \rangle = \lambda \langle \mathbf{x}, \mathbf{y} \rangle + \langle \mathbf{x}, \mathbf{z} \rangle$. Since it costs nothing to write abstract inner product spaces as if they are complex, almost all results will cover the real and complex cases simultaneously, though occasionally a different proof will be required. If you don't feel confident with complex numbers, at first read let $\mathbb{F} = \mathbb{R}$ and delete all complex conjugates!

Definition 2.8. The *standard (Hermitian) inner product and norm* on \mathbb{C}^n are

$$\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{y}^* \mathbf{x} = \sum_{j=1}^n x_j \bar{y}_j = x_1 \bar{y}_1 + \cdots + x_n \bar{y}_n, \quad \|\mathbf{x}\| = \sqrt{\sum_{j=1}^n |x_j|^2}$$

where $\mathbf{y}^* = \bar{\mathbf{y}}^T$ is the *conjugate-transpose*⁶ of \mathbf{y} and $|x_j| = \sqrt{x_j \bar{x}_j}$ is the *modulus*.

Example 2.9. The vectors $\mathbf{x} = \begin{pmatrix} 1+i \\ i \end{pmatrix}$ and $\mathbf{y} = \begin{pmatrix} i \\ 1-i \end{pmatrix}$ are (Hermitian-)orthogonal:

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* \mathbf{x} = \begin{pmatrix} -i & 1+i \end{pmatrix} \begin{pmatrix} 1+i \\ i \end{pmatrix} = -i(1+i) + (1+i)i = 0$$

Weighted inner products may be defined as on page 16 (the weights a_j must still be *positive real numbers*):

$$\langle \mathbf{x}, \mathbf{y} \rangle := \sum_{j=1}^n a_j x_j \bar{y}_j = a_1 x_1 \bar{y}_1 + \cdots + a_n x_n \bar{y}_n$$

We may similarly define inner products in terms of positive-definite matrices $\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{y}^* A \mathbf{x}$.

Definition 2.10. A matrix $A \in M_n(\mathbb{C})$ is *self-adjoint (Hermitian)* if $A^* = A$. A self-adjoint matrix is *positive-definite* if $\mathbf{x}^* A \mathbf{x} > 0$ for all non-zero $\mathbf{x} \in \mathbb{C}^n$.

Self-adjoint means symmetric ($A^T = A$) if A is a real matrix.

Example 2.11. Exercise 6 verifies that $A = \begin{pmatrix} 3 & -i \\ i & 3 \end{pmatrix}$ is positive-definite. We therefore obtain a new inner product on \mathbb{C}^2 :

$$\langle \mathbf{x}, \mathbf{y} \rangle = \begin{pmatrix} \bar{y}_1 & \bar{y}_2 \end{pmatrix} \begin{pmatrix} 3 & -i \\ i & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = 3x_1 \bar{y}_1 + ix_1 \bar{y}_2 - ix_2 \bar{y}_1 + 3x_2 \bar{y}_2$$

⁵The prefix *sesqui-* means *one-and-a-half*; for instance a *sesquicentenary* is a 150 year anniversary.

⁶The physics convention (conjugate-linearity in the **first** slot) is arguably superior for Hermitian inner products, since $\langle \mathbf{x} | \mathbf{y} \rangle = \mathbf{x}^* \mathbf{y} = \bar{\mathbf{x}}^T \mathbf{y}$ keeps \mathbf{x}, \mathbf{y} in the same order, whereas our notation ($\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^* \mathbf{x}$) reverses them.

Further Examples of Inner Product Spaces

As always when discussing inner products, the field \mathbb{F} must be either \mathbb{R} or \mathbb{C} .

Definition 2.12 (Frobenius inner product). If $A, B \in M_{m \times n}(\mathbb{F})$, define

$$\langle A, B \rangle = \text{tr}(B^* A) = \sum_{j=1}^m \sum_{k=1}^n \overline{b_{kj}} a_{jk}$$

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 $M_{m \times n}(\mathbb{F}) \rightarrow \mathbb{F}^{m \times n}$ by stacking the columns of our matrices, then the Frobenius inner product is the standard (Hermitian/dot) inner product in disguise.

Example 2.13. In $M_{3 \times 2}(\mathbb{C})$,

$$\begin{aligned} \left\langle \begin{pmatrix} 1 & i \\ 2-i & 0 \\ 0 & -1 \end{pmatrix}, \begin{pmatrix} 0 & 7 \\ 1 & 2i \\ 3-2i & 4 \end{pmatrix} \right\rangle &= \text{tr} \begin{pmatrix} 0 & 1 & 3+2i \\ 7 & -2i & 4 \end{pmatrix} \begin{pmatrix} 1 & i \\ 2-i & 0 \\ 0 & -1 \end{pmatrix} = \text{tr} \begin{pmatrix} 2-i & -3-2i \\ 9-4i & -4+7i \end{pmatrix} \\ &= (2-i) + (-4+7i) = -2-8i \end{aligned}$$

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

Definition 2.14 (L^2 inner product). Given a *real* interval $[a, b]$, the function

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

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

Verifying Definition 2.1 is straightforward with a little analysis; for instance continuity allows us to conclude

$$\|f\|^2 = \int_a^b |f(x)|^2 dx = 0 \iff f(x) \equiv 0$$

This is our first example of an *infinite-dimensional* inner product space. With careful restriction, this works even for infinite intervals and a larger class of functions.⁷

Example 2.15. Let $f(x) = x$ and $g(x) = x^2$. These lie in the L^2 inner product space $C[-1, 1]$. Indeed,

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

With some simple scaling, we see that $\frac{1}{\|f\|}f = \sqrt{\frac{3}{2}}x$ and $\frac{1}{\|g\|}g = \sqrt{\frac{5}{2}}x^2$ are orthonormal.

⁷For 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 2.16 (ℓ^2 inner product). A sequence (x_n) of real or complex numbers is *square-summable* if the series $\sum |x_n|^2$ converges. It can be shown that the set of such sequences forms a vector space on which we can define an inner product⁸

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

In essence, this is the standard inner product on \mathbb{F}^n after letting $n \rightarrow \infty$! This example, and its L^2 cousin, are the prototypical *Hilbert spaces*, which have great application to differential equations, signal processing, etc. While we will certainly compute with these examples, a rigorous discussion requires significant input from analysis (convergence of series, completeness, integrability), which would take us too far afield.

Exercises 2.1. 1. Evaluate the inner product of each pair of vectors.

(a) $\mathbf{x} = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \mathbf{y} = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix}$ where $\langle \mathbf{x}, \mathbf{y} \rangle = 2x_1y_1 + 3x_2y_2 + x_3y_3$.

(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. Suppose \mathbf{z} is a constant vector in an inner product space V . The linearity condition says that the map $T_{\mathbf{z}} : V \rightarrow \mathbb{F}$ defined by

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

is linear. What, if anything, can you say about the function $U_{\mathbf{z}} : \mathbf{x} \mapsto \langle \mathbf{z}, \mathbf{x} \rangle$?

4. Define $\langle \mathbf{x}, \mathbf{y} \rangle := \sum_{j=1}^n x_j y_j$ 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 2.5 is positive-definite, so that

$$\langle \mathbf{x}, \mathbf{y} \rangle = 3x_1y_1 + x_1y_2 + x_2y_1 + x_2y_2$$

really does define 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!!!)

⁸Neither of these facts are obvious; for instance, we'd need to see that the sum of two square-summable sequences is also square-summable, and that being square-summable forces the inner product series to converge...

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

8. Prove parts 1, 2 and 3 of Lemma 2.6.

9. (Pythagorean Theorem) Let V be an inner product space. If \mathbf{x}, \mathbf{y} are orthogonal, prove

$$\|\mathbf{x} + \mathbf{y}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2$$

10. Use elementary algebra to prove the Cauchy–Schwarz inequality for vectors $\mathbf{x} = \begin{pmatrix} a \\ b \end{pmatrix}$ and $\mathbf{y} = \begin{pmatrix} c \\ d \end{pmatrix}$ 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 2.6?

12. The *polarization identities* demonstrate that if you know the length of every vector, then you know the inner product! Verify them both:

(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, $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{4} \sum_{k=1}^4 i^k \|\mathbf{x} + i^k \mathbf{y}\|^2$

13. Use the Cauchy–Schwarz inequality on a suitable inner product space to prove that

$$\int_0^2 \frac{\sqrt{x}}{x+1} dx \leq \frac{2}{\sqrt{3}}$$

14. (Fourier Series basis) Let $m \in \mathbb{Z}$ and consider the complex-valued function $f_m(x) = \frac{1}{\sqrt{2\pi}} e^{imx}$. Prove that the functions f_m are pairwise orthonormal

$$\langle f_m, f_n \rangle = \begin{cases} 1 & \text{if } m = n \\ 0 & \text{if } m \neq n \end{cases}$$

with respect to the L^2 inner product on $C[-\pi, \pi]$.

(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 $\langle \cdot, \cdot \rangle$ be an inner product on \mathbb{F}^n . Define the matrix $A \in M_n(\mathbb{F})$ by $A_{jk} = \langle \mathbf{e}_k, \mathbf{e}_j \rangle$ where $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ 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 *self-adjoint* ($A^* = A$) and *positive-definite* ($\mathbf{x} \neq \mathbf{0} \implies \mathbf{x}^* A \mathbf{x} > 0$).

More generally, if $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ is a basis of any finite-dimensional inner product space, 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 Subspaces and the Gram–Schmidt Process

In finite dimensions, linear algebra can sometimes feel purely algorithmic: evaluate the co-ordinates of every vector with respect to a basis $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$,

$$\mathbf{x} = \sum_{k=1}^n a_k \mathbf{v}_k = a_1 \mathbf{v}_1 + \dots + a_n \mathbf{v}_n \rightsquigarrow [\mathbf{x}]_\beta = \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \quad (*)$$

and use matrix multiplication to represent linear maps. In practice, finding the co-ordinates can be tedious (many row operations). If, however, β is an *orthogonal* basis, then this becomes almost trivial.

Lemma 2.17. *Let V be an inner product space and $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ an orthogonal spanning set of non-zero vectors:*

$$V = \text{Span } \beta \quad \text{and} \quad \langle \mathbf{v}_j, \mathbf{v}_k \rangle = \begin{cases} 0 & \text{if } j \neq k \\ \|\mathbf{v}_j\|^2 \neq 0 & \text{if } j = k \end{cases}$$

Then β is a basis of V and each $\mathbf{x} \in V$ has unique representation

$$\mathbf{x} = \sum_{k=1}^n \frac{\langle \mathbf{x}, \mathbf{v}_k \rangle}{\|\mathbf{v}_k\|^2} \mathbf{v}_k \quad \left(a_k = \frac{\langle \mathbf{x}, \mathbf{v}_k \rangle}{\|\mathbf{v}_k\|^2} \text{ in } (*) \right)$$

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

Proof. Since β spans V , any given $\mathbf{x} \in V$ may be written as a linear combination (*). Simply take the inner product with each basis vector in turn, and apply the orthogonality of β :

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

Finally, let $\mathbf{x} = \mathbf{0}$ to see that β is linearly independent ($a_1 \mathbf{v}_1 + \dots + a_n \mathbf{v}_n = \mathbf{0} \implies a_k = 0$). ■

Examples 2.18. 1. Let $\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = x_1 \mathbf{e}_1 + x_2 \mathbf{e}_2$ with respect to the standard orthonormal basis $\beta = \{\mathbf{e}_1, \mathbf{e}_2\}$ of \mathbb{R}^2 . We easily check that

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

2. In Euclidean \mathbb{R}^3 , $\beta = \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\} = \left\{ \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 3 \\ 6 \\ -5 \end{pmatrix} \right\}$ is an orthogonal set and thus a basis.

We compute the co-ordinates of $\mathbf{x} = \begin{pmatrix} 7 \\ 4 \\ 2 \end{pmatrix}$ with respect to β :

$$\begin{aligned} \mathbf{x} &= \sum_{k=1}^3 \frac{\langle \mathbf{x}, \mathbf{v}_k \rangle}{\|\mathbf{v}_k\|^2} \mathbf{v}_k = \frac{7+8+6}{1+4+9} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + \frac{14-4+0}{4+1+0} \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix} + \frac{21+24-10}{9+36+25} \begin{pmatrix} 3 \\ 6 \\ -5 \end{pmatrix} \\ &= \frac{3}{2} \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix} + 2 \begin{pmatrix} 2 \\ -1 \\ 0 \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 3 \\ 6 \\ -5 \end{pmatrix} \implies [\mathbf{x}]_\beta = \begin{pmatrix} 3/2 \\ 2 \\ 1/2 \end{pmatrix} \end{aligned}$$

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

3. With respect to the standard inner product on \mathbb{C}^2 , $\beta = \{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \begin{pmatrix} 1 \\ i \end{pmatrix}, \begin{pmatrix} i \\ 1 \end{pmatrix} \right\}$ forms an orthogonal basis (each vector has norm $\sqrt{2}$). Given $\mathbf{x} = \begin{pmatrix} 1+i \\ 3i \end{pmatrix}$, observe that

$$\begin{aligned} \sum_{k=1}^2 \frac{\langle \mathbf{x}, \mathbf{v}_k \rangle}{\|\mathbf{v}_k\|^2} \mathbf{v}_k &= \frac{\langle \begin{pmatrix} 1+i \\ 3i \end{pmatrix}, \begin{pmatrix} 1 \\ i \end{pmatrix} \rangle}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} + \frac{\langle \begin{pmatrix} 1+i \\ 3i \end{pmatrix}, \begin{pmatrix} i \\ 1 \end{pmatrix} \rangle}{2} \begin{pmatrix} i \\ 1 \end{pmatrix} \\ &= \frac{(1+i)(1) + (3i)(-i)}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} + \frac{(1+i)(-i) + (3i)(1)}{2} \begin{pmatrix} i \\ 1 \end{pmatrix} \\ &= \frac{4+i}{2} \begin{pmatrix} 1 \\ i \end{pmatrix} + \frac{1+2i}{2} \begin{pmatrix} i \\ 1 \end{pmatrix} = \mathbf{x} \end{aligned}$$

Otherwise said, $[\mathbf{x}]_\beta = \frac{1}{2} \begin{pmatrix} 4+i \\ 1+2i \end{pmatrix}$.

Orthogonal Subspaces, Complements and Projections

Given the practical benefit of an orthogonal/orthonormal basis, the obvious question is “How do we find one?” For this purpose, and for other reasons, it is helpful to discuss how subspaces interact with the structure of an inner product.

Definition 2.19. Let V be an inner product space and U a subspace. The *orthogonal complement* U^\perp is

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

More generally, subspaces U, W are *orthogonal* if

$$\forall \mathbf{u} \in U, \mathbf{w} \in W, \langle \mathbf{u}, \mathbf{w} \rangle = 0$$

Example 2.20. Given the subspace $U = \text{Span} \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \right\}$ of Euclidean \mathbb{R}^3 , we see that

$$U^\perp = \left\{ \mathbf{x} \in \mathbb{R}^3 : \mathbf{x} \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = 0 = \mathbf{x} \cdot \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \right\} = \text{Span} \left(\begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix} \right)$$

The example certainly suggests that U^\perp is a subspace. This, and much else, holds in general.

Lemma 2.21. Suppose U is a subspace of an inner product space V . Then:

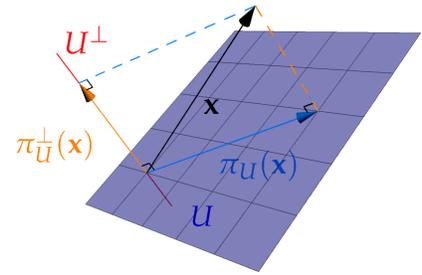
- (a) U^\perp is a subspace of V .
- (b) $U \cap U^\perp = \{\mathbf{0}\}$.
- (c) $U \leq (U^\perp)^\perp$. In particular, U and U^\perp are orthogonal subspaces (Definition 2.19).
- (d) If $V = U \oplus U^\perp$, then $U = (U^\perp)^\perp$.

We leave the proof to Exercise 6. Theorem 2.22 and Corollary 2.27 will show that $V = U \oplus U^\perp$ (and thus $(U^\perp)^\perp = U$) always holds when $\dim U$ is finite. It is perhaps surprising that (c) need not be equality in infinite dimensions (Exercise 8).

Orthogonal Projections As in the example and part (d) of the Lemma, it is common to have a direct sum decomposition $V = U \oplus U^\perp$. Recall what this means:

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

In such a situation, the vectors $\mathbf{u} = \pi_U(\mathbf{x})$, $\mathbf{w} = \pi_{U^\perp}(\mathbf{x})$ are called the *orthogonal projections* of \mathbf{x} onto the subspaces U, U^\perp . Indeed the orthogonal projections can be thought of as *functions*



$$\pi_U : V \rightarrow U, \quad \pi_{U^\perp} : V \rightarrow U^\perp$$

Orthogonal projections are crucial both theoretically and in applications. We'll discuss them repeatedly in future sections.

As our next result shows, orthogonal projections are simple when U has a *finite orthogonal basis*.

Theorem 2.22. *Let V be an inner product space. Suppose $\beta = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ is a set of orthogonal non-zero vectors and let $U = \text{Span } \beta$. Then $V = U \oplus U^\perp$.*

More specifically, every $\mathbf{x} \in V$ decomposes uniquely in the form $\mathbf{x} = \mathbf{u} + \mathbf{w}$ where

$$\mathbf{u} = \pi_U(\mathbf{x}) = \sum_{k=1}^n \frac{\langle \mathbf{x}, \mathbf{u}_k \rangle}{\|\mathbf{u}_k\|^2} \mathbf{u}_k \in U \quad \text{and} \quad \mathbf{w} = \pi_{U^\perp}(\mathbf{x}) = \mathbf{x} - \mathbf{u} \in U^\perp$$

Proof. Since U is an inner product space in its own right, Lemma 2.17 says that β is a basis of U .

Now consider \mathbf{u}, \mathbf{w} as defined. Plainly $\mathbf{u} \in \text{Span } \beta = U$. For each \mathbf{u}_j , the orthogonality of β tells us that

$$\begin{aligned} \langle \mathbf{w}, \mathbf{u}_j \rangle &= \langle \mathbf{x} - \mathbf{u}, \mathbf{u}_j \rangle = \langle \mathbf{x}, \mathbf{u}_j \rangle - \langle \mathbf{u}, \mathbf{u}_j \rangle = \langle \mathbf{x}, \mathbf{u}_j \rangle - \left\langle \sum_{k=1}^n \frac{\langle \mathbf{x}, \mathbf{u}_k \rangle}{\|\mathbf{u}_k\|^2} \mathbf{u}_k, \mathbf{u}_j \right\rangle \\ &= \langle \mathbf{x}, \mathbf{u}_j \rangle - \sum_{k=1}^n \frac{\langle \mathbf{x}, \mathbf{u}_k \rangle}{\|\mathbf{u}_k\|^2} \langle \mathbf{u}_k, \mathbf{u}_j \rangle \\ &= \langle \mathbf{x}, \mathbf{u}_j \rangle - \langle \mathbf{x}, \mathbf{u}_j \rangle = 0 \end{aligned}$$

Since \mathbf{w} is orthogonal to a basis of U , it follows that \mathbf{w} is orthogonal to any element of U ; we conclude that $\mathbf{w} \in U^\perp$.

Finally, suppose $\mathbf{x} = \mathbf{u}_1 + \mathbf{w}_1$ for some $\mathbf{u}_1 \in U, \mathbf{w}_1 \in U^\perp$. Then

$$\mathbf{x} = \mathbf{u}_1 + \mathbf{w}_1 = \mathbf{u} + \mathbf{w} \implies \mathbf{u}_1 - \mathbf{u} = \mathbf{w} - \mathbf{w}_1 \in U \cap U^\perp = \{\mathbf{0}\}$$

by Lemma 2.21. We conclude that the decomposition $\mathbf{x} = \mathbf{u} + \mathbf{w}$ is unique, and that $V = U \oplus U^\perp$ is indeed a direct sum. ■

Example 2.23. Revisiting Example 2.20, we project $\mathbf{x} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$ onto $U = \text{Span} \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} \right\}$:

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

The Gram–Schmidt Process

Theorem 2.22 tells us how to compute the orthogonal projections corresponding to $V = U \oplus U^\perp$, provided U has a finite, orthogonal basis. Given the usefulness of such a set-up, our next goal is to see how to construct such. Thankfully there exists a straightforward algorithm.

Theorem 2.24 (Gram–Schmidt). Suppose $S = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ is a **finite** linearly independent subset of an inner product space V . Inductively construct a sequence of vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$:

$$\mathbf{u}_k := \mathbf{s}_k - \sum_{j=1}^{k-1} \frac{\langle \mathbf{s}_k, \mathbf{u}_j \rangle}{\|\mathbf{u}_j\|^2} \mathbf{u}_j \quad (\mathbf{u}_k = \pi_{U_{k-1}}^\perp(\mathbf{s}_k) \text{ where } U_{k-1} = \text{Span}\{\mathbf{u}_1, \dots, \mathbf{u}_{k-1}\})$$

Then $\beta := \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ is an orthogonal basis of $\text{Span } S$.

In examples, we often choose non-zero multiples $\mathbf{u}_k = a_k \pi_{U_{k-1}}^\perp(\mathbf{s}_k)$ so as to avoid fractions. If you want an *orthonormal* basis, it is usually easier to scale everything after the algorithm is complete.

Example 2.25. $S = \{\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3\} = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \right\}$ is a linearly independent subset of \mathbb{R}^3 .

1. Choose $\mathbf{u}_1 = \mathbf{s}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}$. Observe that $\|\mathbf{u}_1\|^2 = 1$
2. Choose $\mathbf{u}_2 = \mathbf{s}_2 - \frac{\langle \mathbf{s}_2, \mathbf{u}_1 \rangle}{\|\mathbf{u}_1\|^2} \mathbf{u}_1 = \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix} - \frac{2}{1} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix}$. Observe that $\|\mathbf{u}_2\|^2 = 10$.
3. $\pi_{U_2}^\perp(\mathbf{s}_3) = \mathbf{s}_3 - \frac{\langle \mathbf{s}_3, \mathbf{u}_1 \rangle}{\|\mathbf{u}_1\|^2} \mathbf{u}_1 - \frac{\langle \mathbf{s}_3, \mathbf{u}_2 \rangle}{\|\mathbf{u}_2\|^2} \mathbf{u}_2 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} - \frac{1}{1} \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} - \frac{2}{10} \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix} = \frac{2}{5} \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix}$. We choose $\mathbf{u}_3 = \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix}$.

Plainly $\beta = \left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 3 \end{pmatrix}, \begin{pmatrix} 0 \\ 3 \\ 1 \end{pmatrix} \right\}$ is an orthogonal basis of \mathbb{R}^3 .

Proof. For each $k \leq n$, define $S_k = \{\mathbf{s}_1, \dots, \mathbf{s}_k\}$ and $\beta_k = \{\mathbf{u}_1, \dots, \mathbf{u}_k\}$. We prove by induction that each β_k is an orthogonal basis of $U_k := \text{Span } \beta_k = \text{Span } S_k$. The Theorem is the terminal case $k = n$.

(Base case $k = 1$) Certainly $\beta_1 = \{\mathbf{s}_1\} = \{\mathbf{u}_1\}$ is an orthogonal set and $U_1 = \text{Span } \beta_1 = \text{Span } S_1$.

(Induction step) Fix $k \geq 2$ and assume β_{k-1} is an orthogonal basis of $U_{k-1} = \text{Span } S_{k-1}$. By Theorem 2.22, $\mathbf{u}_k = \pi_{U_{k-1}}^\perp(\mathbf{s}_k) \in U_{k-1}^\perp$. We also see that $\mathbf{u}_k \neq \mathbf{0}$, for if not,

$$\mathbf{s}_k = \sum_{j=1}^{k-1} \frac{\langle \mathbf{s}_k, \mathbf{u}_j \rangle}{\|\mathbf{u}_j\|^2} \mathbf{u}_j \in U_{k-1} = \text{Span } S_{k-1}$$

and S would be linearly dependent. It follows that β_k is an orthogonal set of non-zero vectors, and thus (Theorem 2.22) a basis for $U_k = \text{Span } \beta_k$. Moreover,

$$\mathbf{s}_k = \mathbf{u}_k + \sum_{j=1}^{k-1} \frac{\langle \mathbf{s}_k, \mathbf{u}_j \rangle}{\|\mathbf{u}_j\|^2} \mathbf{u}_j \in U_k \implies \text{Span } S_k \leq U_k$$

Since these spaces have the same (finite) dimension k , we conclude that $U_k = \text{Span } S_k$. ■

Example 2.26. We work in the space of real polynomials $P(\mathbb{R})$ equipped with the L^2 inner product $\langle f, g \rangle = \int_0^1 f(x)g(x) dx$. Let $S = \{1, x, x^2\}$ and apply the algorithm (we write $\beta = \{f_1, f_2, f_3\}$):

1. Choose $f_1(x) = 1$. We have $\|f_1\|^2 = \int_0^1 1 dx = 1$.
2. We project the second polynomial x away from $\text{Span}f_1$:

$$\pi_{\text{Span}f_1}^\perp(x) = x - \frac{\langle x, f_1 \rangle}{\|f_1\|^2} f_1 = x - \int_0^1 x dx = x - \frac{1}{2}$$

For simplicity, choose $f_2(x) = 2x - 1$ with $\|f_2\|^2 = \int_0^1 (2x - 1)^2 dx = \frac{1}{3}$.

3. Repeat, projecting the third supplied polynomial x^2 away from $\text{Span}\{f_1, f_2\}$:

$$\begin{aligned} \pi_{\text{Span}\{f_1, f_2\}}^\perp(x^2) &= x^2 - \frac{\langle x^2, f_1 \rangle}{\|f_1\|^2} f_1 - \frac{\langle x^2, f_2 \rangle}{\|f_2\|^2} f_2 \\ &= x^2 - \int_0^1 x^2 dx - \frac{\int_0^1 x^2 (2x - 1) dx}{1/3} (2x - 1) \\ &= x^2 - x + \frac{1}{6} \end{aligned}$$

We choose $f_3(x) = 6x^2 - 6x + 1$ with $\|f_3\|^2 = \int_0^1 (6x^2 - 6x + 1)^2 dx = \frac{1}{5}$

It follows that $\text{Span} S$ has an orthonormal basis

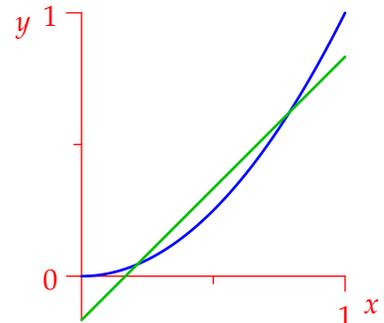
$$\beta = \left\{ 1, \sqrt{3}(2x - 1), \sqrt{5}(6x^2 - 6x + 1) \right\}$$

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

For a taste of how Gram–Schmidt and orthogonal projections may be applied, we use this example to find a **straight-line approximation** to the **parabola** x^2 on the interval $[0, 1]$.

Let $U = \text{Span}\{1, x\}$. As we've just seen, this has orthonormal basis $\{1, \sqrt{3}(2x - 1)\}$. Now compute:

$$\begin{aligned} \pi_U(x^2) &= \langle x^2, 1 \rangle + \langle x^2, \sqrt{3}(2x - 1) \rangle \sqrt{3}(2x - 1) \\ &= \int_0^1 x^2 dx + 3 \left(\int_0^1 2x^3 - x^2 dx \right) (2x - 1) \\ &= x - \frac{1}{6} \end{aligned}$$



As we'll see later, $u(x) = x - \frac{1}{6}$ is the unique linear polynomial which minimizes the function $\|x^2 - u(x)\|^2 = \int_0^1 (x^2 - u(x))^2 dx$. This is a different criteria to what you would be used to in calculus, though $u(x)$ plainly approximates the parabola better than any tangent line would.

We finish with a book-keeping corollary.

Corollary 2.27. 1. Every finite-dimensional inner product space has an orthonormal basis.

2. 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 2.22.

Proof. 1. Simply apply Gram–Schmidt to any basis.

2. By part 1, U has an orthonormal basis. Theorem 2.22 therefore applies. ■

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 8 and 9). 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 this sum converges.

Exercises 2.2. 1. Apply Gram–Schmidt to obtain an orthogonal basis β for $\text{Span } S$. Then obtain the co-ordinate representation of the given vector with respect to β .

(a) $S = \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\}$ with the standard inner/dot product on \mathbb{R}^3 . $\mathbf{x} = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}$.

(b) $S = \left\{ \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \right\}$ with the inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{y}^T \begin{pmatrix} 3 & 1 \\ 1 & 1 \end{pmatrix} \mathbf{x}$ on \mathbb{R}^2 . $\mathbf{x} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$.

(c) $S = \left\{ \begin{pmatrix} 3 & 5 \\ -1 & 1 \end{pmatrix}, \begin{pmatrix} -1 & 9 \\ 5 & -1 \end{pmatrix}, \begin{pmatrix} 7 & -17 \\ 2 & -6 \end{pmatrix} \right\}$ with the Frobenius product (Definition 2.12) on $M_2(\mathbb{R})$. $X = \begin{pmatrix} -1 & 27 \\ -4 & 8 \end{pmatrix}$. (Remember that this is just the dot product in \mathbb{R}^4 in disguise!)

(d) $S = \left\{ \begin{pmatrix} 1 \\ i \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ i \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right\}$ with the standard/Hermitian inner product on \mathbb{C}^3 . $\mathbf{x} = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$.

(e) $S = \{1, x, x^2\}$ with the L^2 inner product $\langle f, g \rangle = \int_{-1}^1 f(x)g(x) dx$ and $f(x) = x^2$.

Important! You will likely need *much* more practice than this to get comfortable with Gram–Schmidt. Make up your own problems: start by changing the numbers in part (a)...

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 $U = \text{Span } S \leq \mathbb{R}^3$. 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 U^\perp to $U = \text{Span}\{x^2\} \leq P_2(\mathbb{R})$ with respect to the L^2 inner product $\langle f, g \rangle = \int_0^1 f(t)g(t) dt$.

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

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

In other words, the function $\phi_\beta : (V, \langle \cdot, \cdot \rangle) \rightarrow (\mathbb{F}^n, \widetilde{\langle \cdot, \cdot \rangle}) : \mathbf{x} \mapsto [\mathbf{x}]_\beta$ is an isomorphism of inner product spaces where $\widetilde{\langle \cdot, \cdot \rangle}$ is the standard (dot/Hermitian) inner product on \mathbb{F}^n .

5. Suppose V, W are finite dimensional inner product spaces with orthonormal bases $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\gamma = \{\mathbf{w}_1, \dots, \mathbf{w}_m\}$ respectively. Suppose also that $T \in \mathcal{L}(V, W)$.

Prove that the matrix $A = [T]_{\beta}^{\gamma} \in M_{m \times n}(\mathbb{F})$ of T with respect to these bases has jk^{th} entry

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

6. Prove all four parts of Lemma 2.21.

7. (a) Verify that the set $\beta = \{1, x\}$ is orthogonal with respect to the L^2 inner product $\langle f, g \rangle = \int_{-1}^1 f(t)g(t) dt$.

(b) Compute the projection $\pi_U(x^2)$ where $U = \text{Span } \beta$. Compare your answer to Example 2.26.

8. Let ℓ^2 be the set of square-summable sequences of real numbers (Definition 2.16). Consider the sequences $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \dots$, where \mathbf{u}_j is the zero sequence except for a single 1 in the j^{th} entry. For instance,

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

(a) Let $U = \text{Span}\{\mathbf{u}_j : j \in \mathbb{N}\}$. 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 $\ell^2 \neq U \oplus U^\perp$.

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

(a) Verify 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 $\text{Span } \beta$.

(c) Sketch a few of the Fourier approximations (sum up to $m = 5$ or $7 \dots$) and observe, when extended to \mathbb{R} , how they approximate a discontinuous periodic function.

10. (Hard) Consider Lemma 2.17 and Theorem 2.22 when β is an *infinite* set. Do these results still hold? Explain as much as you can.

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

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

Example 2.28. As a sanity check in \mathbb{R}^2 , let $A = \begin{pmatrix} 1 & 2 \\ 0 & 3 \end{pmatrix}$, $\mathbf{w} = \begin{pmatrix} x \\ y \end{pmatrix}$ and $\mathbf{v} = \begin{pmatrix} p \\ q \end{pmatrix}$. Then ($A^* = A^T$ here!),

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

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

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 2.29. 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 $T^* : W \rightarrow V$ (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 2.30 (Basic Properties). 1. *If an adjoint exists,⁹ 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: $[T^*]_{\gamma}^{\beta} = ([T]_{\beta})^*$.*

⁹Existence of adjoints is trickier, so we postpone this a little: see Corollary 2.37 and Exercise 12.

1. *Proof. (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 2.6 (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 2.6:

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

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

3. By Exercise 2.2.5, 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^* :

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

Example 2.31. Let $T = L_A \in \mathcal{L}(\mathbb{C}^3, \mathbb{C}^2)$ where $A = \begin{pmatrix} -i & 1 & -3 \\ 2 & 1-i & 4+2i \end{pmatrix}$.

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 !

Part 3 of the Theorem says that every linear map $T \in \mathcal{L}(V, W)$ between finite-dimensional spaces has an adjoint and moreover tells us how to compute it:

Find orthonormal bases β, γ ; compute the matrix $[T]_{\beta}^{\gamma}$ and its conjugate-transpose $([T]_{\beta}^{\gamma})^*$; 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! In practice, it is often better to try a modified approach; see for instance part 2(b) of the next Example.

Examples 2.32. Let $T = \frac{d}{dx} \in \mathcal{L}(P_1(\mathbb{R}))$ be the derivative operator; $T(a + bx) = b$. For variety, we treat $P_1(\mathbb{R})$ as an inner product space in two ways.

1. First 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 2.26, 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 $\gamma = \{g_1, g_2\} = \left\{1, \sqrt{3}f_2\right\} = \left\{1, \sqrt{3}(2x - 1)\right\}$ is orthonormal. Observe that

$$\begin{aligned} T(g_1) &= 0, \quad T(g_2) = 2\sqrt{3} \implies [T]_{\gamma} = \begin{pmatrix} 0 & 2\sqrt{3} \\ 0 & 0 \end{pmatrix} \implies [T^*]_{\gamma} = \begin{pmatrix} 0 & 0 \\ 2\sqrt{3} & 0 \end{pmatrix} \\ \implies T^*(a + bx) &= T^*\left(a + \frac{b}{2} + \frac{b}{2\sqrt{3}}\sqrt{3}(2x - 1)\right) = \left(a + \frac{b}{2}\right) \cdot 2\sqrt{3}g_2 \\ &= 3(2a + b)(2x - 1) \end{aligned}$$

- (b) Use the orthogonal basis β and the projection formula (Theorem 2.22). With $p(x) = a + bx$,

$$\begin{aligned} T^*(p) &= \frac{\langle T^*(p), f_1 \rangle}{\|f_1\|^2} f_1 + \frac{\langle T^*(p), f_2 \rangle}{\|f_2\|^2} f_2 = \langle p, T(1) \rangle + \langle p, T(2x - 1) \rangle \cdot 3(2x - 1) \\ &= \langle p, 0 \rangle + 3 \langle p, 2 \rangle (2x - 1) = 3 \left(\int_0^1 2(a + bx) dx \right) (2x - 1) \\ &= 3(2a + b)(2x - 1) \end{aligned}$$

Note the advantage 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 observation however is that the two examples produce *different* maps T^* : the adjoint of T depends on the inner product!

Why should we care about adjoints?

Adjoint might seem merely to be an abstraction for its own sake. A convincing explanation of why adjoints are useful takes a lot of work, but 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 (Lemma 2.17) how these simplify computations.

The upcoming *spectral theorem* 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 2.33. 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 2.34. 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:

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

The Riesz Representation Theorem

This powerful result demonstrates a natural relationship between an inner product space and its dual $V^* = \mathcal{L}(V, \mathbb{F})$: every linear map $V \rightarrow \mathbb{F}$ arises by taking the inner product with a fixed vector.

Theorem 2.35. If V is finite-dimensional and $g : V \rightarrow \mathbb{F}$ is linear, then there exists a unique $\mathbf{y} \in V$ such that $g(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$.

Example 2.36. $g(p) := \int_0^1 p(x) dx$ is a linear map $g : P_2(\mathbb{R}) \rightarrow \mathbb{R}$. Equip $P_2(\mathbb{R})$ with the inner product for which the standard basis $\{1, x, x^2\}$ is orthonormal. Then

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

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

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 $\text{Span}\{\mathbf{y}\}^\perp \dots$

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

$$V = \mathcal{N}(g) \oplus \mathcal{N}(g)^\perp \quad \text{where} \quad \dim \mathcal{N}(g)^\perp = 1 \quad (\text{rank-nullity theorem and Corollary 2.27})$$

Let $\mathbf{u} \in \mathcal{N}(g)^\perp$ be a unit vector 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

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

The uniqueness of \mathbf{y} follows from the cancellation property (Lemma 2.6, part 4). ■

Due to this tight correspondence, the map is often decorated as g_y . Riesz's theorem indeed says that $y \mapsto g_y$ 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 2.37. *Every linear map on a finite-dimensional inner product space has an adjoint.*

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

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

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

To finish, we simply check that T^* really is the adjoint:

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

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 dot product, $T \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2x+y \\ x-3y \end{pmatrix}$ and $\mathbf{x} = \begin{pmatrix} 3 \\ 5 \end{pmatrix}$

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

(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 $A = \begin{pmatrix} 1 & 1 \\ 4 & 3 \end{pmatrix}$ 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

$$\langle \mathbf{x}, \mathbf{y} \rangle = 4x_1y_1 + x_2y_2$$

(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 2.32, 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 2.30.
6. Prove the Fundamental Subspaces Theorem 2.33.
7. Recall Riesz's Theorem (2.35). For each inner product space V and linear transformation $g : V \rightarrow \mathbb{F}$, find the vector $\mathbf{y} \in V$ such that $g(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$ for all $\mathbf{x} \in V$.
- $V = \mathbb{R}^3$ with the standard inner product, and $g\left(\begin{pmatrix} x \\ y \\ z \end{pmatrix}\right) = x - 2y + 4z$
 - $V = \mathbb{C}^2$ with the standard inner product, and $g\left(\begin{pmatrix} z \\ w \end{pmatrix}\right) = iz - 2w$
 - $V = P_2(\mathbb{R})$ 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 Riesz's Theorem 2.35, explain why \mathbf{y} depends only on g , not \mathbf{u} .
(Hint: what freedom do you have with regard to the unit vector \mathbf{u} ?)
- (b) In the proof of Corollary 2.37, check that $g(\mathbf{x}) := \langle T(\mathbf{x}), \mathbf{z} \rangle$ is a linear map.
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.
- 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)
 - Prove that T^* is diagonalizable if and only if T is.
12. (Hard) Here are two linear maps which do not have an adjoint!
- Since $\epsilon = \{1, x, x^2, \dots\}$ 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

$$T(4 + 3x + 2x^5) = 4 + 3 + 5 = 9$$
 Let $\langle \cdot, \cdot \rangle$ 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*, not a polynomial: T^* therefore does not exist.
 - For a related challenge, recall the space ℓ^2 of square-summable real sequences. For any sequence $(x_n) \in \ell^2$, define $T \in \mathcal{L}(\ell^2)$ via

$$T(x_n) = \left(\sum_{n=1}^{\infty} \frac{1}{n} x_n, 0, 0, 0, 0, \dots \right)$$
 Find the adjoint T^* . If $U \leq \ell^2$ is the subspace whose elements have only finitely many non-zero terms, show that the restriction T_U does not have an adjoint.

2.4 Normal & Self-Adjoint Operators and the Spectral Theorem

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

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

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

If V is a *real* space, then these matrices are identical: $T^* = T$. In the complex case, we instead have

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

Operators this special deserve to be named...

Definition 2.38. Suppose $T \in \mathcal{L}(V)$ is a linear operator for which an adjoint exists. We say that T is:

Self-adjoint if $T^* = T$

Normal if $TT^* = T^*T$

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 *symmetric*: $A^T = A$. A complex self-adjoint matrix is *conjugate-symmetric (Hermitian)*: $A^* = A$.

If T is self-adjoint then it is certainly normal, but the converse is *false*.

Example 2.39. 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:

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

In finite dimensions, 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 2.40 (Spectral Theorem, version 1). Let T be a linear operator on a finite-dimensional inner product space V over a field \mathbb{F} . There are two cases:

1. $\mathbb{F} = \mathbb{R}$. T has an orthonormal basis of eigenvectors if and only if it is **self-adjoint**.
2. $\mathbb{F} = \mathbb{C}$. T has an orthonormal basis of eigenvectors if and only if it is **normal**.

The theorem gets its name from the *spectrum* (set of eigenvalues) of T . Versions of the spectral theorem also apply in certain infinite dimensional settings: such a discussion is beyond these notes.

Examples 2.41. 1. We diagonalize the self-adjoint linear map $T = L_A \in \mathcal{L}(\mathbb{R}^2)$ where $A = \begin{pmatrix} 6 & 3 \\ 3 & -2 \end{pmatrix}$.

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

$$\text{Eigenvalues } \lambda_1 = 7, \quad \lambda_2 = -3$$

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

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

2. The map $T = L_A \in \mathcal{L}(\mathbb{R}^2)$ where $A = \begin{pmatrix} 1 & 3 \\ 0 & -2 \end{pmatrix}$ 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 $A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ 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.

$$\text{Characteristic polynomial } p(t) = t^2 + 1 = (t-i)(t+i)$$

$$\text{Eigenvalues } \lambda_1 = i, \quad \lambda_2 = -i$$

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

Certainly $\langle \mathbf{w}_1, \mathbf{w}_2 \rangle = 0$, $\beta = \{\mathbf{w}_1, \mathbf{w}_2\}$ is orthonormal, and $[T]_\beta = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}$ 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 $A = \begin{pmatrix} 3 & i \\ -i & 3 \end{pmatrix}$ and consider the self-adjoint operator $T = L_A \in \mathcal{L}(\mathbb{C}^2)$.

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

$$\text{Eigenvalues } \lambda_1 = 2, \quad \lambda_2 = 4$$

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

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

Proving the Spectral Theorem for Self-Adjoint Operators

Lemma 2.42 (Basic properties of self-adjoint operators). Let $T \in \mathcal{L}(V)$ be self-adjoint. Then:

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.

1. *Proof.* 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 (λ, \mathbf{x}) is an eigenpair. Then

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

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 $V = \text{Span}\{\mathbf{x}\}$ 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 (λ, \mathbf{x}) where we may assume \mathbf{x} has unit length. Let $W = \text{Span}\{\mathbf{x}\}^\perp$. If $\mathbf{w} \in W$, then

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

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 2.42: 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 2.42; 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 2.42, we summarize some of the basic properties of normal operators. These also apply to self-adjoint operators as a special case.

Lemma 2.43 (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. $\langle \mathbf{x}, (T - tI)(\mathbf{y}) \rangle = \langle \mathbf{x}, T(\mathbf{y}) \rangle - t \langle \mathbf{x}, \mathbf{y} \rangle = \langle T^*(\mathbf{x}), \mathbf{y} \rangle - \langle \bar{t} \mathbf{x}, \mathbf{y} \rangle = \langle (T^* - \bar{t}I)(\mathbf{x}), \mathbf{y} \rangle$ shows that $T - tI$ has adjoint $T^* - \bar{t}I$. It is trivial to check that these commute.

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

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,

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

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 2.44 (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 2.41.2 provides a second, since the matrix A is already upper triangular with respect to the standard orthonormal basis. Here is another example.

Example 2.45. 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 $T \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 2a-b \\ -2a+5b \end{pmatrix}$

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

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

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

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

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

2. Let $T(f(x)) = f'(x) + 4xf(0)$ where $T \in \mathcal{L}(P_1(\mathbb{R}))$ 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 $ST = TS$.

(Hint: recall Theorem 2.30)

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

(Hint: Use Lemma 2.43 and the Fundamental Subspaces Theorem 2.33)

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

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

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

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

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

(b) If W is both T - and T^* -invariant, then $(T_W)^* = (T^*)_W$.

(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 2.42 part 3 and use part (a))

9. (Hard) 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 $W = \text{Span}\{\mathbf{x}\}^\perp$ 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 2.46. A *linear*¹⁰ *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 2.47. Let $T = L_A \in \mathcal{L}(\mathbb{R}^2)$, where $A = \frac{1}{5} \begin{pmatrix} 4 & -3 \\ 3 & 4 \end{pmatrix}$. Then

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

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 2.48. 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, these are often called *orthogonal operators/matrices*; this isn't necessary since *unitary* encompasses both real and complex spaces. An orthogonal matrix satisfies $A^T A = I$.

Example 2.49. The matrix $A = \frac{1}{3} \begin{pmatrix} i & 2+2i \\ 2-2i & i \end{pmatrix}$ 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 2.30, if we choose any orthonormal basis β of V , then

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

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 . As Exercise 13 shows, this isn't an empty requirement...

¹⁰There also exist *non-linear* isometries: for instance *translations* ($T(\mathbf{x}) = \mathbf{x} + \mathbf{a}$ for any constant \mathbf{a}) and *complex conjugation* ($T(\mathbf{x}) = \bar{\mathbf{x}}$) on \mathbb{C}^n . Together with linear isometries, these essentially comprise all isometries in finite dimensions.

We now tackle the promised correspondence between unitary operators and isometries.

Theorem 2.50. *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

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

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

2. For the converse we show that $I - T^*T = 0$ by evaluating it on a basis. Observe that $(I - T^*T)^* = I^* - (T^*T)^* = I - T^*T$ 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 ,

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

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 2.51. *In finite dimensions, being unitary is equivalent to each of the following:*

- (a) *Preservation of the inner product¹¹ (+).*
- (b) *The **existence** of an orthonormal basis $\beta = \{\mathbf{w}_1, \dots, \mathbf{w}_n\}$ such that $T(\beta) = \{T(\mathbf{w}_1), \dots, T(\mathbf{w}_n)\}$ is also orthonormal.*
- (c) *That **every** orthonormal basis β of V is mapped to an orthonormal basis $T(\beta)$.*

Exercise 9 establishes claims (b) and (c). 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 2.52. *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.

¹¹In 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 2.53. 1. The matrix $A_\theta = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \in M_2(\mathbb{R})$ is orthogonal for any θ . Example 2.47 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 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. $A = \frac{1}{\sqrt{6}} \begin{pmatrix} \sqrt{2} & \sqrt{3} & 1 \\ \sqrt{2} & 0 & -2 \\ -\sqrt{2} & \sqrt{3} & -1 \end{pmatrix} \in M_3(\mathbb{R})$ is orthogonal: check the columns!

3. $A = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix} \in M_2(\mathbb{C})$ is unitary: 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 compute the characteristic polynomial and the eigenvalues:

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

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

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

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:

$$\begin{aligned} \langle T(f(x)), T(g(x)) \rangle &= \langle e^{ix}f(x), e^{ix}g(x) \rangle = \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{ix}f(x)\overline{e^{ix}g(x)} dx \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} e^{ix}e^{-ix}f(x)\overline{g(x)} dx = \langle f(x), g(x) \rangle \end{aligned}$$

(b), (c) $C[-\pi, \pi]$ doesn't have an orthonormal basis: indeed there is no orthonormal set $\beta = \{f_k\}$ of which every continuous function is a (*finite*) linear combination.¹² We cannot therefore claim that T maps orthonormal bases to orthonormal bases!

¹²An *infinite* orthonormal set $\beta = \{f_k : k \in \mathbb{Z}\}$ can be found so that every continuous 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 a basis. 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!

Unitary and Orthogonal Equivalence

Suppose $A \in M_n(\mathbb{R})$ is symmetric (self-adjoint) $A^T = A$. By the spectral theorem, it has an orthonormal eigenbasis $\beta = \{\mathbf{w}_1, \dots, \mathbf{w}_n\}$: $A\mathbf{w}_j = \lambda_j\mathbf{w}_j$. 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 2.54. The matrix $A = \begin{pmatrix} 1+i & 1+i \\ -1-i & 1+i \end{pmatrix}$ 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

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

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 a general phenomenon...

Definition 2.55. Square matrices A, B are *unitarily equivalent* if there exists a unitary matrix U such that $B = U^*AU$. *Orthogonal equivalence* is similar: $B = U^T AU$.

As suggested by the name, this is indeed an equivalence relation (Exercise 4). The above discussion proves half the following:

Theorem 2.56. $A \in M_n(\mathbb{C})$ 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 $A = U^*DU$. Then

$$\begin{aligned} A^*A &= (U^*DU)^*U^*DU = U^*D^*UU^*DU = U^*\bar{D}DU = U^*D\bar{D}U = U^*DUU^*\bar{D}U \\ &= U^*DU(U^*DU)^* = AA^* \end{aligned}$$

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

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

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

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 $\begin{pmatrix} a & -e^{i\theta}\bar{b} \\ b & e^{i\theta}a \end{pmatrix}$ 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. Hence conclude that unitary equivalence is an equivalence relation on the set $M_n(\mathbb{C})$.

(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 $A^T A = I$.

*(Such **complex** orthogonal matrices don't have the same nice relationship with inner products)*

6. Supply the details of Exercise 2.53.1.

(Hints: $\beta = \{\mathbf{i}, \mathbf{j}\}$ is orthonormal, whence $\{A\mathbf{i}, A\mathbf{j}\}$ must be orthonormal. 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 2.53.4 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 2.51 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 $T(x_1, x_2, \dots) = (0, x_1, x_2, \dots)$. 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 $\beta_W = \{\mathbf{w}_1, \dots, \mathbf{w}_n\}$. In such a situation, we can define the *orthogonal projections* onto W and W^\perp via

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

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 :

$$V = X \oplus Y \iff \forall \mathbf{v} \in V, \exists \text{ unique } \mathbf{x} \in X, \mathbf{y} \in Y \text{ such that } \mathbf{v} = \mathbf{x} + \mathbf{y}$$

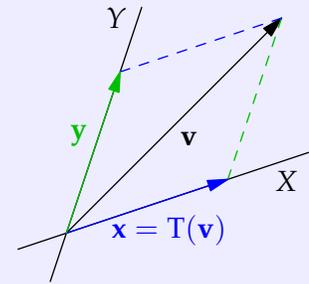
Definition 2.57. A linear map $T \in \mathcal{L}(V)$ is a *projection* if:

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

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 $\mathbf{x} + \mathbf{y} \mapsto \mathbf{x}$.

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 2.58. $A = \frac{1}{5} \begin{pmatrix} 6 & -2 \\ 3 & -1 \end{pmatrix}$ is a projection matrix with $\mathcal{R}(A) = \text{Span} \left(\begin{pmatrix} 2 \\ 1 \end{pmatrix} \right)$ and $\mathcal{N}(A) = \text{Span} \left(\begin{pmatrix} 1 \\ 3 \end{pmatrix} \right)$. 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 $\mathcal{R}(A) = \text{Span} \left(\begin{pmatrix} a \\ b \end{pmatrix} \right)$ and $\mathcal{N}(A) = \text{Span} \left(\begin{pmatrix} c \\ d \end{pmatrix} \right)$, then

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

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 2.59. $T \in \mathcal{L}(V)$ 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 .¹³ 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 2.60. An *orthogonal projection* is a projection $T \in \mathcal{L}(V)$ on an inner product space for which

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

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

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

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 (2.58 continued). The identity and zero matrices are both 2×2 orthogonal projection matrices, while those of type 3 are orthogonal if $\begin{pmatrix} a \\ b \end{pmatrix} \cdot \begin{pmatrix} c \\ d \end{pmatrix} = 0$: 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 $\sum_{j=1}^k \mathbf{w}_j \mathbf{w}_j^*$.

¹³In finite dimensions, the rank-nullity theorem and dimension counting finishes the proof here:

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

Theorem 2.61. A projection $T \in \mathcal{L}(V)$ is orthogonal if and only if it is self-adjoint $T = T^*$.

Proof. We prove the two directions separately.

(\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 (2.33),

$$\mathcal{N}(T) = \mathcal{N}(T^*) = \mathcal{R}(T)^\perp$$

Since T is a projection, $V = \mathcal{R}(T) \oplus \mathcal{N}(T) = \mathcal{R}(T) \oplus \mathcal{R}(T)^\perp$, from which (Lemma 2.21)

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

The language of projections allows us to rephrase the Spectral Theorem.

Theorem 2.62 (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 $\{\lambda_1, \dots, \lambda_k\}$ and corresponding eigenspaces E_1, \dots, E_k . Let $\pi_j \in \mathcal{L}(V)$ be the orthogonal projection onto E_j . Then:

1. $V = E_1 \oplus \dots \oplus E_k$ where each 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) $I_V = \pi_1 + \dots + \pi_k$
4. (Spectral decomposition) $T = \lambda_1 \pi_1 + \dots + \lambda_k \pi_k$

Proof. 1. V certainly equals the direct sum since T is diagonalizable. 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. $T(\mathbf{x}) = \sum_{j=1}^k T(\mathbf{x}_j) = \sum_{j=1}^k \lambda_j \mathbf{x}_j = \sum_{j=1}^k \lambda_j \pi_j(\mathbf{x})$.

Alternatively, for 2, 3 and 4 one can simply evaluate both sides on an orthonormal eigenbasis. ■

Examples 2.63. 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 $A = \begin{pmatrix} 10 & 2 \\ 2 & 7 \end{pmatrix}$ has spectrum $\{6, 11\}$ and orthonormal eigenvectors

$$\mathbf{w}_6 = \frac{1}{\sqrt{5}} \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \quad \mathbf{w}_{11} = \frac{1}{\sqrt{5}} \begin{pmatrix} 2 \\ 1 \end{pmatrix}$$

The corresponding projections therefore have matrices

$$\pi_6 = \mathbf{w}_6 \mathbf{w}_6^T = \frac{1}{5} \begin{pmatrix} 1 \\ -2 \end{pmatrix} (1 \ -2) = \frac{1}{5} \begin{pmatrix} 1 & -2 \\ -2 & 4 \end{pmatrix}, \quad \pi_{11} = \mathbf{w}_{11} \mathbf{w}_{11}^T = \frac{1}{5} \begin{pmatrix} 4 & 2 \\ 2 & 1 \end{pmatrix}$$

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 $B = \begin{pmatrix} 1+i & 1+i \\ -1-i & 1+i \end{pmatrix} \in M_2(\mathbb{C})$ has spectrum $\{2, 2i\}$ and corresponding orthonormal eigenvectors

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

The orthogonal projection matrices are therefore

$$\pi_2 = \mathbf{w}_2 \mathbf{w}_2^* = \frac{1}{2} \begin{pmatrix} i \\ 1 \end{pmatrix} (-i \ 1) = \frac{1}{2} \begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix} \quad \pi_{2i} = \mathbf{w}_{2i} \mathbf{w}_{2i}^* = \frac{1}{2} \begin{pmatrix} 1 & -i \\ i & 1 \end{pmatrix}$$

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 $C = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$ 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

$$\begin{aligned} \pi_2 &= \mathbf{u}\mathbf{u}^T = \frac{1}{3} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} (1 \ 1 \ 1) = \frac{1}{3} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix} \\ \pi_{-1} &= \mathbf{v}\mathbf{v}^T + \mathbf{w}\mathbf{w}^T = \frac{1}{2} \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix} (1 \ -1 \ 0) + \frac{1}{6} \begin{pmatrix} 1 \\ 1 \\ -2 \end{pmatrix} (1 \ 1 \ -2) \\ &= \frac{1}{2} \begin{pmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} + \frac{1}{6} \begin{pmatrix} 1 & 1 & -2 \\ 1 & 1 & -2 \\ -2 & -2 & 4 \end{pmatrix} = \frac{1}{3} \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix} \end{aligned}$$

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

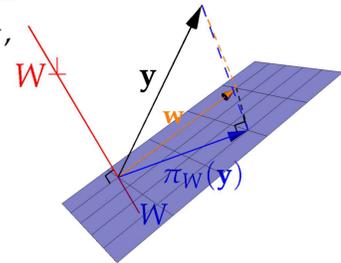
We finish this section with an important observation that drives much of the application of inner product spaces to other parts of mathematics and beyond. Throughout this discussion, X and Y denote inner product spaces.

Theorem 2.64. 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} :

$$\forall \mathbf{w} \in W, \|\mathbf{y} - \pi_W(\mathbf{y})\| \leq \|\mathbf{y} - \mathbf{w}\|$$

Proof. Apply the Pythagorean Theorem: since $\pi_W^\perp(\mathbf{y}) = \mathbf{y} - \pi_W(\mathbf{y}) \in W^\perp$,

$$\begin{aligned} \|\mathbf{y} - \mathbf{w}\|^2 &= \|\mathbf{y} - \pi_W(\mathbf{y}) + \pi_W(\mathbf{y}) - \mathbf{w}\|^2 \\ &= \|\mathbf{y} - \pi_W(\mathbf{y})\|^2 + \|\pi_W(\mathbf{y}) - \mathbf{w}\|^2 \\ &\geq \|\mathbf{y} - \pi_W(\mathbf{y})\|^2 \end{aligned}$$



with equality if and only if $\mathbf{w} = \pi_W(\mathbf{y})$. ■

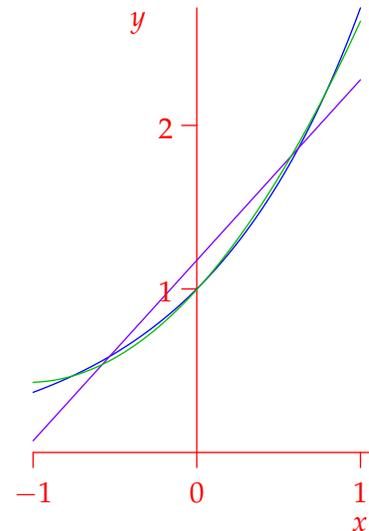
This set-up can be used to compute accurate approximations in many contexts.

Examples 2.65. 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 $W = \text{Span}\{1, x, x^2\}$, 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

$$\begin{aligned} p(x) &= \frac{1}{2} \langle 1, e^x \rangle + \frac{3}{2} \langle x, e^x \rangle x + \frac{5}{8} (3x^2 - 1) \langle 3x^2 - 1, e^x \rangle \\ &= \frac{1}{2} \int_{-1}^1 e^x dx + \frac{3}{2} x \int_{-1}^1 x e^x dx \\ &\quad + \frac{5}{8} (3x^2 - 1) \int_{-1}^1 (3x^2 - 1) e^x dx \\ &= \frac{1}{2} (e - e^{-1}) + 3e^{-1}x + \frac{5}{4} (e - 7e^{-1})(3x^2 - 1) \\ &\approx 1.18 + 1.10x + 0.179(3x^2 - 1) \\ &\approx 1 + 1.1x + 0.537x^2 \end{aligned}$$



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 space

$$\begin{aligned} W_n &= \text{Span} \{1, e^{ix}, e^{-ix}, \dots, e^{inx}, e^{-inx}\} \\ &= \text{Span} \{1, \cos x, \sin x, \cos 2x, \sin 2x, \dots, \cos nx, \sin nx\} \end{aligned}$$

with respect to the L^2 inner product on $[-\pi, \pi]$:

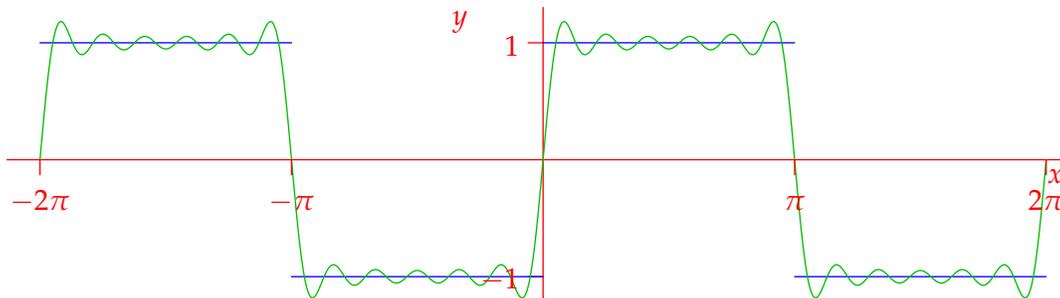
$$\begin{aligned} \mathcal{F}_n(x) &= \frac{1}{2\pi} \sum_{k=-n}^n \langle f(x), e^{ikx} \rangle e^{ikx} \\ &= \frac{1}{2\pi} \langle f(x), 1 \rangle + \frac{1}{\pi} \sum_{k=1}^n \langle f(x), \cos kx \rangle \cos kx + \langle f(x), \sin kx \rangle \sin kx \end{aligned}$$

By Theorem 2.64, 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, given $f(x) = \begin{cases} 1 & \text{if } 0 < x \leq \pi, \\ -1 & \text{if } -\pi < x \leq 0, \end{cases}$

$$\mathcal{F}_{2n-1}(x) = \frac{4}{\pi} \sum_{j=1}^n \frac{\sin(2j-1)x}{2j-1} = \frac{4}{\pi} \left(\sin x + \frac{\sin 3x}{3} + \frac{\sin 5x}{5} + \dots + \frac{\sin(2n-1)x}{2n-1} \right)$$



$y = f(x)$ and its eleventh Fourier approximation $y = \mathcal{F}_{11}(x)$ extended periodically

Exercises 2.6. 1. Compute the matrices of the orthogonal projection onto W viewed as a subspace of the standard inner product spaces \mathbb{R}^n or \mathbb{C}^n .

(a) $W = \text{Span} \left(\begin{pmatrix} 4 \\ -1 \end{pmatrix} \right)$ (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 matrix, 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 already have orthonormal eigenbases from Exercise 2.5.1)

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.
 (Hint: evaluate $\|T(\mathbf{r} - \lambda\mathbf{n})\|^2$ where $\mathbf{r} \in \mathcal{R}(T)$ and $\mathbf{n} \in \mathcal{N}(T)$ are unit vectors...)
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 $T = \lambda_1\pi_1 + \dots + \lambda_k\pi_k$ to prove the following.
 - (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 2.65.2
 - (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 a linear map $T \in \mathcal{L}(V, W)$ 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 see what can be done when $V \neq W$ and allow $\beta \neq \gamma$. The result is an approach that applies to *any* linear map between finite-dimensional inner product spaces.

Example 2.66. Let $A = \begin{pmatrix} 3 & 1 \\ 2 & -2 \\ 1 & 3 \end{pmatrix}$ and consider the 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 $A\mathbf{v}_2 = 2\sqrt{3}\mathbf{w}_2$, we see that $[L_A]_{\beta}^{\gamma} = \begin{pmatrix} 4 & 0 \\ 0 & 2\sqrt{3} \\ 0 & 0 \end{pmatrix}$ is almost diagonal.

Our main result says that such bases always exist!

Theorem 2.67 (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 $\beta = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ of V and $\gamma = \{\mathbf{w}_1, \dots, \mathbf{w}_m\}$ of W , and positive scalars $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r$ such that

$$T(\mathbf{v}_j) = \begin{cases} \sigma_j \mathbf{w}_j & \text{if } j \leq r \\ \mathbf{0} & \text{otherwise} \end{cases} \quad \text{equivalently} \quad [T]_{\beta}^{\gamma} = \begin{pmatrix} \text{diag}(\sigma_1, \dots, \sigma_r) & \mathbf{O} \\ \mathbf{O} & \mathbf{O} \end{pmatrix}$$

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 $A \in M_{m \times n}(\mathbb{F})$ has rank $A = r$, then $A = P\Sigma Q^*$ where

$$\Sigma = [L_A]_{\beta}^{\gamma} = \begin{pmatrix} \text{diag}(\sigma_1, \dots, \sigma_r) & \mathbf{O} \\ \mathbf{O} & \mathbf{O} \end{pmatrix}, \quad P = (\mathbf{w}_1, \dots, \mathbf{w}_m), \quad Q = (\mathbf{v}_1, \dots, \mathbf{v}_n)$$

Since the columns of P, Q are orthonormal, both matrices are unitary.

Definition 2.68. The numbers $\sigma_1, \dots, \sigma_r$ are the *singular values* of T . If T is not maximum rank, we have additional zero singular values $\sigma_{r+1} = \dots = \sigma_{\min(m,n)} = 0$.

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 (Exercise 5).

Rank-one Decomposition If we write $g_j : V \rightarrow \mathbb{F}$ for the linear map $g_j : \mathbf{v} \mapsto \langle \mathbf{v}, \mathbf{v}_j \rangle$ (Theorem 2.35), 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 ($\mathbf{w}_j g_j : V \rightarrow W$). For matrices, g_j is just multiplication by the row vector \mathbf{v}_j^* and we may write¹⁴

$$A = \sum_{j=1}^r \sigma_j \mathbf{w}_j \mathbf{v}_j^*$$

Freedom of Choice While the singular values are uniquely determined, there is often some freedom regarding β and γ , particularly if an eigenspace of T^*T has dimension ≥ 2 . The block-diagonal matrix $[T]_\beta^\gamma$ remains independent of the bases chosen.

Example (2.66 cont). We apply the method in the Theorem to the matrix A .

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 $\sigma_1 = 4$, $\sigma_2 = 2\sqrt{3}$. Moreover

$$\mathbf{w}_1 = \frac{1}{\sigma_1} A \mathbf{v}_1 = \frac{1}{4\sqrt{2}} A \begin{pmatrix} 1 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \mathbf{w}_2 = \frac{1}{\sigma_2} A \mathbf{v}_2 = \frac{1}{2\sqrt{6}} A \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \frac{1}{\sqrt{6}} \begin{pmatrix} -1 \\ -2 \\ 1 \end{pmatrix}$$

are orthonormal; 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 a sum of rank-one matrices:

$$A = \sigma_1 \mathbf{w}_1 \mathbf{v}_1^T + \sigma_2 \mathbf{w}_2 \mathbf{v}_2^T = 4 \begin{pmatrix} \frac{1}{\sqrt{2}} \\ 0 \\ \frac{1}{\sqrt{2}} \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} + 2\sqrt{3} \begin{pmatrix} \frac{-1}{\sqrt{6}} \\ \frac{-2}{\sqrt{6}} \\ \frac{1}{\sqrt{6}} \end{pmatrix} \begin{pmatrix} \frac{-1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{pmatrix} = \begin{pmatrix} 2 & 2 \\ 0 & 0 \\ 2 & 2 \end{pmatrix} + \begin{pmatrix} 1 & -1 \\ 2 & -2 \\ -1 & 1 \end{pmatrix}$$

¹⁴Since β is orthonormal, it is common to write \mathbf{v}_j^* for the map $g_j = \langle \cdot, \mathbf{v}_j \rangle$ in general contexts. To those familiar with the dual space $V^* = \mathcal{L}(V, \mathbb{F})$, the set $\{g_1, \dots, g_n\} = \{\mathbf{v}_1^*, \dots, \mathbf{v}_n^*\}$ 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. Everything depends on the Spectral Theorem (2.40)!

1. T^*T is self-adjoint and therefore 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

$$\langle T(\mathbf{v}_j), T(\mathbf{v}_k) \rangle = \langle T^*T(\mathbf{v}_j), \mathbf{v}_k \rangle = \lambda_j \langle \mathbf{v}_j, \mathbf{v}_k \rangle = \begin{cases} \lambda_j & \text{if } j = k \\ 0 & \text{if } j \neq k \end{cases} \quad (*)$$

In particular, every eigenvalue is a *non-negative real number*: $\lambda_j = \|T(\mathbf{v}_j)\|^2 \geq 0$.

Since $\text{rank } T^*T = \text{rank } T = r$ (Exercise 6), exactly r eigenvalues are non-zero. By reordering the eigenbasis if necessary, we may assume

$$\lambda_1 \geq \dots \geq \lambda_r > 0$$

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 $[T]_\beta^\gamma = \begin{pmatrix} \text{diag}(\sigma_1, \dots, \sigma_r) & O \\ O & O \end{pmatrix}$, then $[T^*]_\gamma^\beta$ is *essentially the same matrix* just that its dimensions have been reversed ($m \times n \rightarrow n \times m$). But then

$$[T^*]_\gamma^\beta = \begin{pmatrix} \text{diag}(\sigma_1, \dots, \sigma_r) & O \\ O & O \end{pmatrix} \implies T^*T = \begin{pmatrix} \text{diag}(\sigma_1^2, \dots, \sigma_r^2) & O \\ O & O \end{pmatrix}$$

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 are column vectors and so the (change of co-ordinate) matrices P, Q are unitary. ■

Examples 2.69. 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 singular value decomposition

$$A = P \Sigma Q^* = \begin{pmatrix} \frac{2}{\sqrt{5}} & \frac{-1}{\sqrt{5}} \\ \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \end{pmatrix} \begin{pmatrix} 4 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{\sqrt{5}} & \frac{2}{\sqrt{5}} \\ \frac{-2}{\sqrt{5}} & \frac{1}{\sqrt{5}} \end{pmatrix}$$

Multiply this out to express 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. A 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 $\langle f, g \rangle = \int_0^1 f(x)g(x) dx$ on $P_2(\mathbb{R})$ and $P_1(\mathbb{R})$ admits orthonormal bases

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

If $T = \frac{d}{dx}$ is the derivative operator, then 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 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 $T \in \mathcal{L}(V, W)$ gives rise to a natural function from W back to V which behaves somewhat like an inverse even when T is non-invertible!

Definition 2.70. 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 $\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_r\} \rightarrow \text{Span}\{\mathbf{w}_1, \dots, \mathbf{w}_r\}$, the pseudoinverse really does invert T :

$$\begin{array}{ccc} V = \underbrace{\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_r\}}_{\mathcal{N}(T)^{\perp} = \mathcal{R}(T^{\dagger})} \oplus \underbrace{\text{Span}\{\mathbf{v}_{r+1}, \dots, \mathbf{v}_n\}}_{\mathcal{N}(T) = \mathcal{R}(T^{\dagger})^{\perp}} & & \{\mathbf{0}_V\} \\ \uparrow T & \text{bijection} & \downarrow T^{\dagger} \\ W = \underbrace{\text{Span}\{\mathbf{w}_1, \dots, \mathbf{w}_r\}}_{\mathcal{R}(T) = \mathcal{N}(T^{\dagger})^{\perp}} \oplus \underbrace{\text{Span}\{\mathbf{w}_{r+1}, \dots, \mathbf{w}_m\}}_{\mathcal{R}(T)^{\perp} = \mathcal{N}(T^{\dagger})} & & \{\mathbf{0}_W\} \end{array}$$

Otherwise said, the combinations are *orthogonal projections*: $T^{\dagger}T = \pi_{\mathcal{N}(T)^{\perp}}^{\perp}$ and $TT^{\dagger} = \pi_{\mathcal{R}(T)}$

Given a singular value decomposition of a matrix $A = P\Sigma Q^*$, its pseudoinverse is

$$A^\dagger = \sum_{j=1}^r \frac{1}{\sigma_j} \mathbf{v}_j \mathbf{w}_j^* = Q \tilde{\Sigma} P^* \quad \text{where} \quad \tilde{\Sigma} = \begin{pmatrix} \text{diag}(\sigma_1^{-1}, \dots, \sigma_r^{-1}) & O \\ O & O \end{pmatrix}$$

Examples 2.71. 1. (Example 2.66) The matrix $A = \begin{pmatrix} 3 & 1 \\ 2 & -2 \\ 1 & 3 \end{pmatrix}$ has pseudoinverse

$$\begin{aligned} A^\dagger &= \frac{1}{\sigma_1} \mathbf{v}_1 \mathbf{w}_1^* + \frac{1}{\sigma_2} \mathbf{v}_2 \mathbf{w}_2^* \\ &= \frac{1}{4\sqrt{2}\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} (1 \ 0 \ 1) + \frac{1}{2\sqrt{3}\sqrt{2}\sqrt{6}} \begin{pmatrix} -1 \\ 1 \end{pmatrix} (-1 \ -2 \ 1) \\ &= \frac{1}{8} \begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix} + \frac{1}{12} \begin{pmatrix} 1 & 2 & -1 \\ -1 & -2 & 1 \end{pmatrix} = \frac{1}{24} \begin{pmatrix} 5 & 4 & 1 \\ 1 & -4 & 5 \end{pmatrix} \end{aligned}$$

which is exactly what we would have found by computing $A^\dagger = Q \tilde{\Sigma} P^T$. Observe also 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 projections onto $\mathcal{N}(A)^\perp = \text{Span}\{\mathbf{v}_1, \mathbf{v}_2\} = \mathbb{R}^2$ and $\mathcal{R}(A) = \text{Span}\{\mathbf{w}_1, \mathbf{w}_2\} \leq \mathbb{R}^3$ respectively. Both spaces have dimension two since $\text{rank } A = 2$.

2. (Example 2.69.2) The pseudoinverse of $T = \frac{d}{dx} : P_2(\mathbb{R}) \rightarrow P_1(\mathbb{R})$,

$$\begin{aligned} T^\dagger(\sqrt{3}(2x-1)) &= \frac{1}{2\sqrt{15}} \sqrt{5}(6x^2-6x+1) = \frac{1}{2\sqrt{3}}(6x^2-6x+1) \\ T^\dagger(1) &= \frac{1}{2\sqrt{3}} \sqrt{3}(2x-1) = x - \frac{1}{2} \\ T^\dagger(a+bx) &= T^\dagger\left(a + \frac{b}{2} + \frac{b}{2\sqrt{3}} \sqrt{3}(2x-1)\right) = \left(a + \frac{b}{2}\right) \left(x - \frac{1}{2}\right) + \frac{b}{12}(6x^2-6x+1) \\ &= \frac{b}{2}x^2 + ax - \frac{a}{2} - \frac{b}{6} \end{aligned}$$

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 a singular value decomposition of each matrix and compute its pseudoinverse.

(a) $\begin{pmatrix} 1 & 1 \\ -1 & -1 \end{pmatrix}$ (b) $\begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & -1 \end{pmatrix}$ (c) $\begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & 0 \\ 1 & 0 & -1 \end{pmatrix}$

2. For each map, describe the singular value decomposition and compute the pseudoinverse.

(a) $T \in \mathcal{L}(\mathbb{R}^2, \mathbb{R}^3)$ where $T \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x \\ x+y \\ x-y \end{pmatrix}$

(b) $T : P_2(\mathbb{R}) \rightarrow P_1(\mathbb{R})$ 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 $T(f) = f' + 2f$

3. Suppose $A = P\Sigma Q^*$ is a singular value decomposition.
- Describe a singular value decomposition of A^* .
 - Explain why $A^\dagger = Q\tilde{\Sigma}P^*$ isn't a singular value decomposition of A^\dagger ; what would be a correct decomposition?
4. Suppose $T : V \rightarrow W$ 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.
5. (a) Suppose $T = \mathcal{L}(V)$ is normal. Prove that each \mathbf{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?
6. 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!)
7. Let V, W be finite-dimensional inner product spaces and $T \in \mathcal{L}(V, W)$.
- By evaluating on the basis γ in the singular value theorem, prove that $T^*TT^\dagger = T^\dagger TT^* = T^*$.
 - If T is injective, then T^*T is invertible and $T^\dagger = (T^*T)^{-1}T^*$.
 - If T is surjective, then TT^* is invertible and $T^\dagger = T^*(TT^*)^{-1}$.
 - If $A \in M_{m \times n}(\mathbb{F})$ has linearly independent *columns*, prove that $\pi_{\mathcal{R}(A)} = A(A^*A)^{-1}A^*$.
8. Suppose T is a linear map between finite-dimensional inner product spaces. A *least-squares solution* to the equation $T(\mathbf{x}) = \mathbf{b}$ is a vector \mathbf{x} which minimizes $\|T(\mathbf{x}) - \mathbf{b}\|$.
- Use Theorem 2.64 to prove that $\mathbf{x}_0 = T^\dagger(\mathbf{b})$ is a least-squares solution, and that any other has the form $\mathbf{x}_0 + \mathbf{n}$ for some $\mathbf{n} \in \mathcal{N}(T)$.
 - Prove that $\mathbf{x}_0 = T^\dagger(\mathbf{b})$ has smaller norm than any other least-squares solution.
 - If T is injective, prove that $\mathbf{x}_0 = T^\dagger(\mathbf{b})$ is the unique least-squares solution.
9. Find the minimal norm solution to the first system, and the least-squares solution to the second:

$$\begin{cases} 3x + 2y + z = 9 \\ x - 2y + 3z = 3 \end{cases} \quad \begin{cases} 3x + y = 1 \\ 2x - 2y = 0 \\ x + 3y = 0 \end{cases}$$

(Hint: All the required details are in the various incarnations of Example 2.66)

10. Compute the pseudoinverse of the matrix $\begin{pmatrix} -1 & 7 \\ 3 & 4 \\ 7 & 1 \end{pmatrix}$ and use it to find the point (x, y) which comes closest to solving the system

$$\begin{cases} -x + 7y = 1 \\ 3x + 4y = 0 \\ 7x + y = 0 \end{cases}$$

Least-Squares Regression (non-examinable)

Given a data set $\{(t_j, y_j) : 1 \leq j \leq m\}$, the least-squares method can be employed to find a *best-fitting line* $y = c_0 + c_1 t$ to the data. This *regression line* is used to predict y given a value of t .

The idea is to choose coefficients c_0, c_1 so as to minimize the sum of the squares of the **vertical deviations** of the line from the data set.

$$SS = \sum_{j=1}^m |y_j - c_0 - c_1 t_j|^2 = \|\mathbf{y} - A\mathbf{x}\|^2 \quad \text{where} \quad A = \begin{pmatrix} t_1 & 1 \\ \vdots & \vdots \\ t_m & 1 \end{pmatrix} \quad \mathbf{x} = \begin{pmatrix} c_1 \\ c_0 \end{pmatrix} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_m \end{pmatrix}$$

With the indicated matrix notation, we recognize this as a least-squares problem. Moreover, if there are at least two distinct t -values in the data set, then $\text{rank } A = 2$ is maximal and we have a unique best-fitting least-squares line: By Exercise 8 its coefficients are

$$\begin{pmatrix} c_1 \\ c_0 \end{pmatrix} = A^+ \mathbf{y} = (A^T A)^{-1} A^T \mathbf{y}$$

Example 2.72. Given the data set $\{(0, 1), (1, 1), (2, 0), (3, 2), (4, 2)\}$, we compute

$$A = \begin{pmatrix} 0 & 1 \\ 1 & 1 \\ 2 & 1 \\ 3 & 1 \\ 4 & 1 \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} 1 \\ 1 \\ 0 \\ 2 \\ 2 \end{pmatrix} \Rightarrow \mathbf{x}_0 = \begin{pmatrix} c_1 \\ c_0 \end{pmatrix} = (A^T A)^{-1} A^T \mathbf{y} = \begin{pmatrix} 30 & 10 \\ 10 & 5 \end{pmatrix}^{-1} \begin{pmatrix} 15 \\ 6 \end{pmatrix} = \frac{3}{10} \begin{pmatrix} 1 \\ 2 \end{pmatrix}$$

The regression line therefore has equation $y = \frac{3}{10}(t + 2)$.

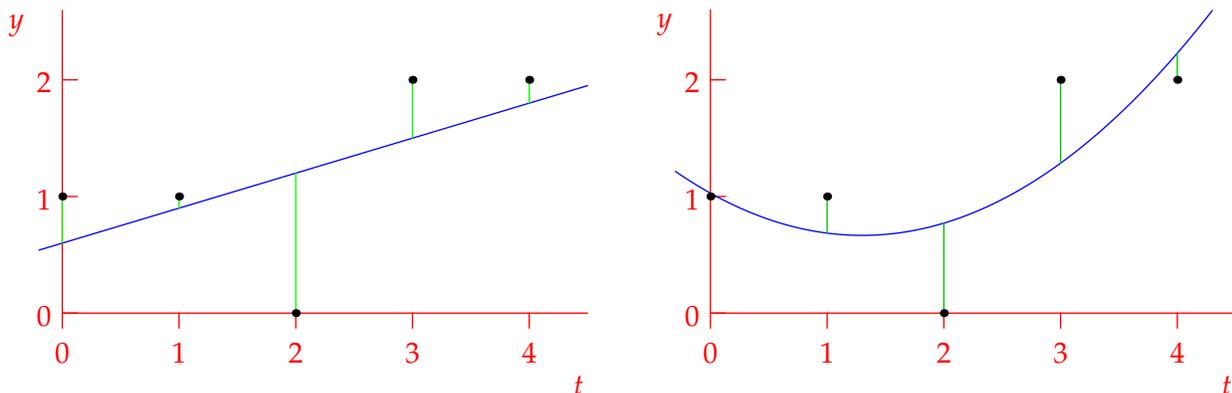
The process can be applied more generally to approximate using other functions. To find the best-fitting quadratic polynomial $y = c_0 + c_1 t + c_2 t^2$, we'd instead work with

$$A = \begin{pmatrix} t_1^2 & t_1 & 1 \\ \vdots & \vdots & \vdots \\ t_m^2 & t_m & 1 \end{pmatrix} \quad \mathbf{x} = \begin{pmatrix} c_2 \\ c_1 \\ c_0 \end{pmatrix} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_m \end{pmatrix} \Rightarrow \sum_{j=1}^m |y_j - c_0 - c_1 t_j - c_2 t_j^2|^2 = \|\mathbf{y} - A\mathbf{x}\|^2$$

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.



Extra Problems

11. Check the calculation for the best-fitting least-squares quadratic in Example 2.72.
12. Find the best-fitting least-squares linear and quadratic approximations to the data set

$$\{(1, 2), (3, 4), (5, 7), (7, 9), (9, 12)\}$$

13. 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:

$$c = \frac{\text{Cov}(t, y)}{\sigma_t^2} \quad \text{and} \quad d = \bar{y} - c\bar{t}$$

where

- $\bar{t} = \frac{1}{m} \sum_{j=1}^m t_j$ and $\bar{y} = \frac{1}{m} \sum_{j=1}^m y_j$ are the *means* (averages),
- $\sigma_t^2 = \frac{1}{m} \sum_{j=1}^m (t_j - \bar{t})^2$ is the *variance*,
- $\text{Cov}(t, y) = \frac{1}{m} \sum_{j=1}^m (t_j - \bar{t})(y_j - \bar{y})$ is the *covariance*.

2.8 Bilinear and Quadratic Forms

In this section we slightly generalize the idea of an inner product. Throughout, V is simply a vector space over a field \mathbb{F} : it need not be an inner product space.

Definition 2.73. A bilinear form $B : V \times V \rightarrow \mathbb{F}$ is linear in each entry: $\forall \mathbf{v}, \mathbf{x}, \mathbf{y} \in V, \lambda \in \mathbb{F}$,

$$B(\lambda \mathbf{x} + \mathbf{y}, \mathbf{v}) = \lambda B(\mathbf{x}, \mathbf{v}) + B(\mathbf{y}, \mathbf{v}), \quad B(\mathbf{v}, \lambda \mathbf{x} + \mathbf{y}) = \lambda B(\mathbf{v}, \mathbf{x}) + B(\mathbf{v}, \mathbf{y})$$

Additionally, B is *symmetric* if $\forall \mathbf{x}, \mathbf{y} \in V, B(\mathbf{x}, \mathbf{y}) = B(\mathbf{y}, \mathbf{x})$.

Examples 2.74. 1. If V is a *real* inner product space, then the inner product $\langle \cdot, \cdot \rangle$ 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 ,

$$B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 1 & 2 \\ 2 & 0 \end{pmatrix} \mathbf{y} = x_1 y_1 + 2x_1 y_2 + 2x_2 y_1$$

is a symmetric bilinear form; it isn't positive-definite so isn't an inner product (e.g. $B(\mathbf{j}, \mathbf{j}) = 0$).

As in the example, we often make use of a matrix.

Definition 2.75. Let B be a bilinear form on a finite-dimensional space with basis $\epsilon = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$. The *matrix of B with respect to ϵ* is $[B]_\epsilon = A \in M_n(\mathbb{F})$ whose ij^{th} entry is $A_{ij} = B(\mathbf{v}_i, \mathbf{v}_j)$.

Since B is bilinear, if $\mathbf{x}, \mathbf{y} \in V$ have 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,

$$B(\mathbf{y}, \mathbf{x}) = [\mathbf{y}]_\epsilon^T A [\mathbf{x}]_\epsilon = ([\mathbf{y}]_\epsilon^T A [\mathbf{x}]_\epsilon)^T = [\mathbf{x}]_\epsilon^T A^T [\mathbf{y}]_\epsilon$$

so that B is symmetric if and only if A is also. Finally, if β is another basis of V , then an appeal to the change of co-ordinate matrix Q_ϵ^β yields

$$\begin{aligned} B(\mathbf{x}, \mathbf{y}) &= [\mathbf{x}]_\epsilon^T A [\mathbf{y}]_\epsilon = (Q_\beta^\epsilon [\mathbf{x}]_\beta)^T A (Q_\beta^\epsilon [\mathbf{y}]_\epsilon) = [\mathbf{x}]_\beta^T (Q_\beta^\epsilon)^T A Q_\beta^\epsilon [\mathbf{y}]_\beta \\ &\implies [B]_\beta = (Q_\beta^\epsilon)^T [B]_\epsilon Q_\beta^\epsilon \end{aligned}$$

To summarize:

Lemma 2.76. 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. This can sometimes be done by brute force, but there is an algorithm...

Examples 2.77. Performing a sequence of *simultaneous* row and column operations will diagonalize any symmetric B . Moreover, only elementary operations/matrices $E_{ij}^{(\lambda)}$ of type III are required.¹⁵

1. We diagonalize Example 2.74.2 where $[B]_e = A = \begin{pmatrix} 1 & 2 \\ 2 & 0 \end{pmatrix}$. The trick is to **subtract twice column 1 from column 2**, and then do the same thing with rows: **subtract twice row 1 from row 2**.

$$E_{21}^{(-2)} A E_{12}^{(-2)} = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 2 & 0 \end{pmatrix} \begin{pmatrix} 1 & -2 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 2 & -4 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix} \quad (*)$$

We conclude that $[B]_\beta = (Q_\beta^\epsilon)^T [B]_e Q_\beta^\epsilon$ where $\beta = \{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -2 \\ 1 \end{pmatrix} \right\}$ is a diagonalizing basis. Otherwise said,

$$B(p_1 \mathbf{v}_1 + p_2 \mathbf{v}_2, q_1 \mathbf{v}_1 + q_2 \mathbf{v}_2) = p_1 q_1 - 4 p_2 q_2$$

If you *really* want, (*) can be rearranged to translate everything back to the standard basis ϵ :

$$\begin{aligned} B(\mathbf{x}, \mathbf{y}) &= \mathbf{x}^T \begin{pmatrix} 1 & 2 \\ 2 & 0 \end{pmatrix} \mathbf{y} = (x_1 \quad x_2) \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix} \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} \\ &= (x_1 + 2x_2 \quad x_2) \begin{pmatrix} 1 & 0 \\ 0 & -4 \end{pmatrix} \begin{pmatrix} y_1 + 2y_2 \\ y_2 \end{pmatrix} = (x_1 + 2x_2)(y_1 + 2y_2) - 4x_2 y_2 \end{aligned}$$

2. We diagonalize the bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y} = \mathbf{x}^T \begin{pmatrix} 1 & -2 & 3 \\ -2 & 0 & 4 \\ 3 & 4 & -1 \end{pmatrix} \mathbf{y}$.

- Add twice column 1 to column 2, rows similarly: $E_{21}^{(2)} A E_{12}^{(2)} = \begin{pmatrix} 1 & 0 & 3 \\ 0 & -4 & 10 \\ 3 & 10 & -1 \end{pmatrix}$
- Subtract thrice column 1 from column 3, rows similarly: $E_{31}^{(-3)} E_{21}^{(2)} A E_{12}^{(2)} E_{13}^{(-3)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -4 & 10 \\ 0 & 10 & -10 \end{pmatrix}$
- Add column 3 to column 2, rows similarly: $E_{23}^{(1)} E_{31}^{(-3)} E_{21}^{(2)} A E_{12}^{(2)} E_{13}^{(-3)} E_{32}^{(1)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 6 & 0 \\ 0 & 0 & -10 \end{pmatrix}$
- If β is the diagonalizing basis, then 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 $\beta = \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \end{pmatrix}, \begin{pmatrix} -3 \\ 0 \\ 1 \end{pmatrix} \right\}$.

If you're having trouble believing this, invert Q_β^ϵ as above 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 $\mathbb{F} = \mathbb{R}$, the Spectral Theorem says that every symmetric B may be diagonalized by an orthonormal basis (for the usual dot product). This is not what our algorithm produces! Our algorithm tends to be much faster (no solving polynomials) and applies over any field.

¹⁵Recall 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 2.78. The diagonalization algorithm is non-unique! Here are three ways to diagonalize the bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 1 & 6 \\ 6 & 3 \end{pmatrix} \mathbf{y} = x_1y_1 + 6(x_1y_2 + x_2y_1) + 3x_2y_2$ on \mathbb{R}^2 .

1. $\begin{pmatrix} 1 & 0 \\ -6 & 1 \end{pmatrix} \begin{pmatrix} 1 & 6 \\ 6 & 3 \end{pmatrix} \begin{pmatrix} 1 & -6 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & -33 \end{pmatrix} = [B]_\beta$ where $\beta = \left\{ \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -6 \\ 1 \end{pmatrix} \right\}$. This corresponds to

$$B(\mathbf{x}, \mathbf{y}) = (x_1 + 6x_2)(y_1 + 6y_2) - 33x_2y_2$$

2. $\begin{pmatrix} 1 & -2 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 6 \\ 6 & 3 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -2 & 1 \end{pmatrix} = \begin{pmatrix} -11 & 0 \\ 0 & 3 \end{pmatrix} = [B]_\gamma$ where $\gamma = \left\{ \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$. This corresponds to

$$B(\mathbf{x}, \mathbf{y}) = -11x_1y_1 + 3(2x_1 + x_2)(2y_1 + y_2)$$

3. Since $\mathbb{F} = \mathbb{R}$, we could instead apply the spectral theorem, though it is a lot more work for no benefit (we omit the disgusting expression for $B(\mathbf{x}, \mathbf{y})$ in these co-ordinates):

$$\eta = \left\{ \frac{1}{\sqrt{74 + 2\sqrt{37}}} \begin{pmatrix} 6 \\ 1 + \sqrt{37} \end{pmatrix}, \frac{1}{\sqrt{74 + 2\sqrt{37}}} \begin{pmatrix} -1 - \sqrt{37} \\ 6 \end{pmatrix} \right\}, \quad [B]_\eta = \begin{pmatrix} 2 + \sqrt{37} & 0 \\ 0 & 2 - \sqrt{37} \end{pmatrix}$$

Theorem 2.79. Suppose B is a bilinear form on 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 $n = \dim V$.

Base Case ($n = 1$) Plainly $B(x, y) = axy$ for some $a \in \mathbb{F}$; this is clearly symmetric.

Induction step Fix n , and assume that every non-zero symmetric bilinear form on a dimension n vector space over \mathbb{F} is diagonalizable. Suppose B is a non-zero symmetric bilinear form on a space V of dimension $n + 1$.

The aside (below) shows $\exists \mathbf{x} \in V$ such that $B(\mathbf{x}, \mathbf{x}) \neq 0$. Consider the linear map

$$T \in \mathcal{L}(V, \mathbb{F}), \quad T(\mathbf{v}) = B(\mathbf{x}, \mathbf{v})$$

Since $B(\mathbf{x}, \mathbf{x}) \neq 0$ we have $\text{rank } T = 1$, from which $\text{null } T = n$. Since B is symmetric when restricted to $\mathcal{N}(T)$, the induction hypothesis says there exists a basis β of $\mathcal{N}(T)$ diagonalizing $B_{\mathcal{N}(T)}$ is diagonal. But then B is diagonal with respect to the basis $\beta \cup \{\mathbf{x}\}$. ■

Fields of Characteristic Two The proof requires some $\mathbf{x} \in V$ 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 $\mathbf{x} = \mathbf{u} + \mathbf{v}$ looks like it does the job:

$$B(\mathbf{x}, \mathbf{x}) = B(\mathbf{u}, \mathbf{v}) + B(\mathbf{v}, \mathbf{u}) = 2B(\mathbf{u}, \mathbf{v}) \neq 0 \tag{†}$$

In a field of *characteristic two*, however, we have $2 = 1 + 1 = 0$, so (†) fails!

To see that this requirement isn't idle, consider the field $\mathbb{F} = \mathbb{Z}_2 = \{0, 1\}$ of remainders modulo 2, and the symmetric bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \mathbf{y}$ on the finite vector space $\mathbb{Z}_2^2 = \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix} \right\}$ over \mathbb{Z}_2 . Every element of this space satisfies $B(\mathbf{x}, \mathbf{x}) = 0$! It can be checked that the matrix of B is identical with respect to *any* basis of \mathbb{Z}_2^2 , whence B is symmetric but non-diagonalizable.

Sylvester's Law of Inertia

In Example 2.78 notice how the three diagonal matrix representations have something in common: each has exactly one positive and one negative entry. This is a general phenomenon:

Theorem 2.80 (Sylvester). Suppose B is a symmetric bilinear form on a real vector space V with diagonal matrix representation $\text{diag}(\lambda_1, \dots, \lambda_n)$. Then the number of entries λ_j which are positive/negative/zero is independent of the diagonal representation.

Definition 2.81. 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.

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 with a signature $(1, 3, 0)$ bilinear form,

$$B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} c^2 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \mathbf{y} = c^2 x_1 y_1 - x_2 y_2 - x_3 y_3 - x_4 y_4$$

where c is the speed of light. A vector \mathbf{x} is *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.

Sketch 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 $\text{rank } B := \text{rank } A$ and observe this is independent of basis (Exercise 8).
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 $r = \text{rank } B$ in accordance with part 1: our goal is to prove that $p = q$.

3. For contradiction, 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) $\text{rank } C \leq r - q + p < r \implies \text{null } C > n - r$, thus

$$\exists \mathbf{x} \in \mathbb{R}^n \text{ such that } C\mathbf{x} = \mathbf{0} \text{ and } \mathbf{x} \notin \text{Span}\{\mathbf{v}_{r+1}, \dots, \mathbf{v}_n\}$$

- (b) The first p entries of $C\mathbf{x} = \mathbf{0}$ force $\mathbf{x} \in \text{Span}\{\mathbf{v}_{p+1}, \dots, \mathbf{v}_r, \mathbf{v}_{r+1}, \dots, \mathbf{v}_n\}$ and so $B(\mathbf{x}, \mathbf{x}) < 0$. Note how part (a) guarantees a *strict* inequality here.

- (c) Finally, write \mathbf{x} with respect to γ . This time the last $r - q$ entries of $C\mathbf{x} = \mathbf{0}$ show that $\mathbf{x} \in \text{Span}\{\mathbf{w}_1, \dots, \mathbf{w}_q, \mathbf{w}_{r+1}, \dots, \mathbf{w}_n\}$, whence $B(\mathbf{x}, \mathbf{x}) \geq 0$, contradicting part (b). ■

Quadratic Forms & Diagonalizing Conics

Definition 2.82. To every symmetric bilinear form $B : V \times V \rightarrow \mathbb{F}$ is associated a *quadratic form*

$$K : V \rightarrow \mathbb{F} : \mathbf{x} \mapsto B(\mathbf{x}, \mathbf{x})$$

A function $K : V \rightarrow \mathbb{F}$ is termed a quadratic form when such a symmetric bilinear form exists.

Examples 2.83. 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 ,

$$K(\mathbf{x}) = \mathbf{x}^T A \mathbf{x} = \sum_{i,j=1}^n x_i A_{ij} x_j = \sum_{1 \leq i \leq j \leq n} \tilde{a}_{ij} x_i x_j \quad \text{where } \tilde{a}_{ij} = \begin{cases} A_{ij} & \text{if } i = j \\ 2A_{ij} & \text{if } i \neq j \end{cases}$$

E.g., $K(\mathbf{x}) = 3x_1^2 + 4x_2^2 - 2x_1x_2$ corresponds to the bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 3 & -1 \\ -1 & 4 \end{pmatrix} \mathbf{y}$

As a fun application, we diagonalize 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 non-zero quadratic form

$$K(\mathbf{x}) = ax^2 + 2bxy + cy^2 \Leftrightarrow B(\mathbf{v}, \mathbf{w}) = \mathbf{v}^T \begin{pmatrix} a & b \\ b & c \end{pmatrix} \mathbf{w}$$

If $\{\mathbf{v}_1, \mathbf{v}_2\}$ is a diagonalizing basis, then there exist scalars λ_1, λ_2 (not both zero since $K \neq 0$) 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 $\lambda_i \neq 0$, the linear transformation $s_j = t_j + \frac{\mu_j}{2\lambda_j}$ completes the square, recovering the familiar canonical forms:¹⁶

Parabola: $\lambda_1 s_1^2 = s_2 + k$ where $k, \lambda_1 \neq 0$ (or reverse $1 \leftrightarrow 2$).

Ellipse: $\lambda_1 s_1^2 + \lambda_2 s_2^2 = k \neq 0$ where λ_1, λ_2, k have the same sign.

Hyperbola: $\lambda_1 s_1^2 + \lambda_2 s_2^2 = k \neq 0$ where λ_1, λ_2 have opposite signs.

Since B is symmetric, there exists an *orthonormal* basis $\{\mathbf{v}_1, \mathbf{v}_2\}$ of \mathbb{R}^2 diagonalizing B : any conic can therefore be put in canonical form by applying only a rigid motion: a combination of rotation, reflection and translation (completing the square).

Alternatively, we could diagonalize K using our earlier algorithm; geometrically this also permits shear transforms. By Sylvester's Law, the diagonal entries have the same number of $(+, -, 0)$ terms regardless of the method, so the canonical form is unchanged.

¹⁶Any other possibility is *degenerate*: empty (e.g., $x^2 + y^2 = -1$), one or two points ($x^2 + y^2 = 0$ or $x^2 = 1$), or the product of two lines ($x^2 - y^2 = (x - y)(x + y) = 0$).

Examples 2.84. 1. We describe and plot the conic with equation $7x^2 + 24xy = 144$.
 The associated bilinear form has matrix $\begin{pmatrix} 7 & 12 \\ 12 & 0 \end{pmatrix}$ with orthonormal eigenbasis

$$\beta = \{\mathbf{v}_1, \mathbf{v}_2\} = \left\{ \frac{1}{5} \begin{pmatrix} 4 \\ 3 \end{pmatrix}, \frac{1}{5} \begin{pmatrix} -3 \\ 4 \end{pmatrix} \right\}$$

and eigenvalues $(\lambda_1, \lambda_2) = (16, -9)$. In the rotated basis, this is the canonical hyperbola

$$16t_1^2 - 9t_2^2 = 144 \iff \frac{t_1^2}{3^2} - \frac{t_2^2}{4^2} = 1$$

In case this is too fast, use the change of co-ordinate matrix to compute directly:

$$Q_\beta^\epsilon = \frac{1}{5} \begin{pmatrix} 4 & -3 \\ 3 & 4 \end{pmatrix} \implies \begin{pmatrix} t_1 \\ t_2 \end{pmatrix} = [\mathbf{x}]_\beta = Q_\beta^\epsilon \begin{pmatrix} x \\ y \end{pmatrix} = \frac{1}{5} \begin{pmatrix} 4x + 3y \\ -3x + 4y \end{pmatrix}$$

which recovers the original conic by substitution.

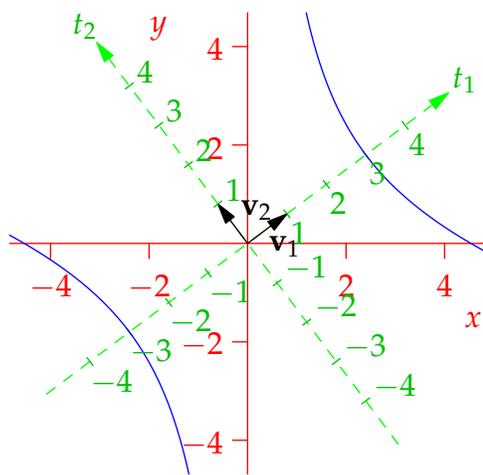
2. The conic defined by $K(\mathbf{x}) = x^2 + 12xy + 3y^2 = -33$ (Example 2.78) is also a hyperbola. With respect to the basis $\gamma = \left\{ \begin{pmatrix} 1 \\ -2 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\}$,

$$K(\mathbf{x}) = -11x^2 + 3(2x + y)^2 = -11t_1^2 + 3t_2^2 = -33 \iff \frac{t_1^2}{3} - \frac{t_2^2}{11} = 1$$

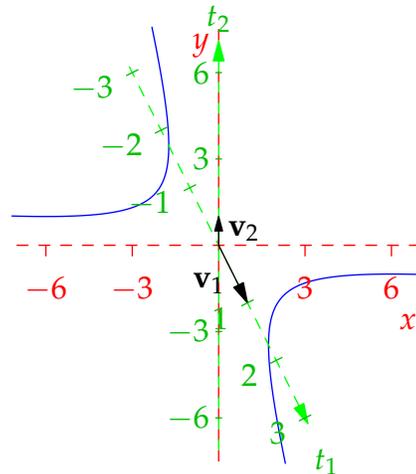
Even though γ is non-orthogonal, we can still plot the conic! If we instead used the orthonormal basis η , we'd obtain

$$K(\mathbf{x}) = (\sqrt{37} + 2)s_1^2 - (\sqrt{37} - 2)s_2^2 = -33$$

however finding η is time-consuming and the expressions for s_1, s_2 are extremely ugly.



Part 1



Part 2

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 (the set of bilinear forms on V is therefore a vector space).
(You can't use matrices here since V could be infinite-dimensional!)

2. Compute the matrix of the bilinear form

$$B(\mathbf{x}, \mathbf{y}) = x_1y_1 - 2x_1y_2 + x_2y_1 - x_3y_3$$

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 $B(f, g) = f'(0)g''(0)$ is a bilinear form on the vector space of twice-differentiable functions. Find the matrix of B with respect to $\beta = \{\cos t, \sin t, \cos 2t, \sin 2t\}$ on $\text{Span } \beta$.

4. For each matrix A , find a diagonal matrix D and an invertible matrix Q such that $Q^T A Q = D$. Hence describe a basis β which diagonalizes the bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T A \mathbf{y}$.

(a) $\begin{pmatrix} 1 & 3 \\ 3 & 2 \end{pmatrix}$ (b) $\begin{pmatrix} 5 & 1 \\ 1 & 3 \end{pmatrix}$ (c) $\begin{pmatrix} 3 & 1 & 2 \\ 1 & 4 & 0 \\ 2 & 0 & -1 \end{pmatrix}$

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

$$B(\mathbf{x}, \mathbf{y}) = \frac{1}{2} (K(\mathbf{x} + \mathbf{y}) - K(\mathbf{x}) - K(\mathbf{y}))$$

7. Consider the bilinear form $B(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \mathbf{y}$ on \mathbb{F}^2 .

(a) Compute the associated quadratic form $K(\mathbf{x})$.

(b) If $\text{char } \mathbb{F} \neq 2$, find a way to apply the diagonalizing algorithm to B . What goes wrong if $\text{char } \mathbb{F} = 2$?

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 $\text{rank } B$ is independent of the choice of diagonalizing basis.

9. Describe and plot the conics:

(a) $x^2 + y^2 + xy = 6$ (b) $35x^2 + 120xy = 4x + 3y$

10. Suppose $ax^2 + 2bxy + cy^2 + dx + ey + f = 0$ is a non-empty, non-degenerate conic C in \mathbb{R}^2 , and define its *discriminant* $\Delta = b^2 - ac$. Prove:

- 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 $\Delta > 0$.