Linear Algebra II

Martin Otto

Summer Term 2011

Contents

Notational conventions

In this second part, we adopt the following conventions and notational simplifications over and above conventions established in part I.

- all bases are understood to be labelled bases, with individual basis vectors listed as for instance in $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$.

- we extend the idea of summation notation to products, writing for instance

$$\prod_{i=1}^{n} \lambda_i \quad \text{for} \quad \lambda_1 \cdot \ldots \cdot \lambda_n.$$

and also to sums and direct sums of subspaces, as in

$$\sum_{i=1}^{k} U_i \quad \text{for} \quad U_1 + \cdots + U_k$$

and

$$\bigoplus_{i=1}^{k} U_i \quad \text{for} \quad U_1 \oplus \cdots \oplus U_k.$$

- in rings (like $\operatorname{Hom}(V, V)$) we use exponentiation notation for iterated products, always identifying the power to exponent zero with the unit element (neutral element w.r.t. multiplication). For instance, if φ is an endomorphism of V, we write

$$\varphi^k \quad \text{for} \quad \underbrace{\varphi \circ \cdots \circ \varphi}_{k \text{ times}} \quad \text{where } \varphi^0 = \operatorname{id}_V.$$

Chapter 1

Eigenvalues and Diagonalisation

One of the core topics of linear algebra concerns the choice of suitable bases for the representation of linear maps. For an endomorphism $\varphi\colon V \to V$ of the n-dimensional $\mathbb{F}$-vector space V, we want to find a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of V that is specially adapted to make the matrix representation $A_\varphi = [\![\varphi]\!]_B^B \in \mathbb{F}^{(n,n)}$ as simple as possible for the analysis of the map φ itself. This chapter starts from the study of *eigenvectors* of φ. Such an eigenvector is a vector $\mathbf{v} \neq \mathbf{0}$ that is mapped to a scalar multiple of itself under φ. So $\varphi(\mathbf{v}) = \lambda\mathbf{v}$ for some $\lambda \in \mathbb{F}$, the corresponding *eigenvalue*. If $\mathbf{v}$ is an eigenvector (with eigenvalue λ), it spans a one-dimensional subspace $U = \{\mu\mathbf{v}\colon \mu \in \mathbb{F}\} \subseteq V$, which is *invariant under φ* (or preserved by φ) in the sense that $\varphi(\mathbf{u}) = \lambda\mathbf{u} \in U$ for all $\mathbf{u} \in U$. In other words, in the direction of $\mathbf{v}$, φ acts as a rescaling with factor λ.

But for instance over the standard two-dimensional $\mathbb{R}$-vector space $\mathbb{R}^2$, an endomorphism need not have any eigenvectors. Consider the example of a rotation through an angle $0 < \alpha < \pi$; this linear map preserves no 1-dimensional subspaces at all. In contrast, we shall see later that any endomorphism of $\mathbb{R}^3$ must have at least one eigenvector. In the case of a non-trivial rotation, for instance, the axis of rotation gives rise to an eigenvector with eigenvalue 1.

In fact eigenvalues and eigenvectors often have further significance, either geometrically or in terms of the phenomena modelled by a linear map. To give an example, consider the homogeneous linear differential equation of the

harmonic oscillator

$$\frac{\mathrm{d}^2}{\mathrm{d}t^2} f(t) + cf(t) = 0$$

with a positive constant $c \in \mathbb{R}$ and for C^∞ functions $f \colon \mathbb{R} \to \mathbb{R}$ (modelling, for instance, the position of a mass attached to a spring as a function of time t). One may regard this as the problem of finding the eigenvectors for eigenvalue $-c$ of the linear operator $\frac{\mathrm{d}^2}{\mathrm{d}t^2}$ that maps a C^∞ function to its second derivative. In this case, there are two linearly independent solutions $f(t) = \sin(\sqrt{c}\,t)$ and $f(t) = \cos(\sqrt{c}\,t)$ which span the solution space of this differential equation. The eigenvalue $-c$ of $\frac{\mathrm{d}^2}{\mathrm{d}t^2}$ that we look at here is related to the frequency of the oscillation (which is $\sqrt{c}/2\pi$).

In quantum mechanics, states of a physical system are modelled as vectors of a $\mathbb{C}$-vector space (e.g., of wave functions). Associated physical quantities (observables) are described by linear (differential) operators on such states, which are endomorphisms of the state space. The eigenvectors of these operators are the possible values for measurements of those observables. Here one seeks bases of the state space made up from eigenvectors (eigenstates) associated with particular values for the quantity under consideration via their eigenvalues. W.r.t. such a basis an arbitrary state can be represented as a linear combination (superposition) of eigenstates, accounting for components which each have their definite value for that observable, but mixed in a composite state with different possible outcomes for its measurement.

If V has a basis consisting of eigenvectors of an endomorphism $\varphi \colon V \to V$, then w.r.t. that basis, φ is represented by a diagonal matrix, with the eigenvalues as entries on the diagonal, and all other entries equal to 0. Matrix arithmetic is often much simpler for diagonal matrices, and it therefore makes sense to apply a corresponding basis transformation to achieve diagonal form where this is possible.

Example 1.0.1 Consider for instance a square matrix $A \in \mathbb{R}^{(n,n)}$ (or $\mathbb{C}^{(n,n)}$). Suppose that there is a regular matrix C (for a basis transformation) such that

$$\hat{A} = CAC^{-1} = \begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & & 0 \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & \lambda_n \end{pmatrix}$$

is a diagonal matrix. Suppose now that we want to evaluate powers of A: $A^0 = E_n$, $A^1 = A$, $A^2 = AA$, ... We note that the corresponding powers of $\hat{A}$ are very easily computed. One checks that

$$\hat{A}^\ell = \begin{pmatrix} \lambda_1^\ell & 0 & \cdots & 0 \\ 0 & \lambda_2^\ell & & 0 \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n^\ell \end{pmatrix}.$$

Then

$$A^\ell = (C^{-1}\hat{A}C)^\ell = \underbrace{(C^{-1}\hat{A}C) \cdot \ldots \cdot (C^{-1}\hat{A}C)}_{\ell \text{ times}} = C^{-1}\hat{A}^\ell C$$

is best computed via the detour through $\hat{A}$.

Example 1.0.2 Consider the linear differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{v}(t) = A\mathbf{v}(t)$$

for a vector-valued C^∞ function $\mathbf{v} \colon t \mapsto \mathbf{v}(t) \in \mathbb{R}^n$, where the matrix $A \in \mathbb{R}^{(n,n)}$ is fixed. In close analogy with the one-dimensional case, one can find a solution using the exponential function. Here the exponential function has to be considered as a matrix-valued function on matrices, $B \mapsto e^B$, defined by its series expansion. The function $t \mapsto \mathbf{v}(t) := e^{tA}\mathbf{v}_0$ solves the differential equation for initial value $\mathbf{v}(0) = \mathbf{v}_0$. Here e^{tA} stands for the series

$$\sum_{k=0}^\infty \frac{t^k A^k}{k!} \in \mathbb{R}^{(n,n)},$$

which can be shown to converge for all t. The value $\mathbf{v}(t) = e^{tA}\mathbf{v}_0$ is the result of applying the matrix e^{tA} to the vector $\mathbf{v}_0$.

If it so happens that there is a basis relative to which A is similar to a diagonal matrix $\hat{A} = CAC^{-1}$, as in the previous example, then we may solve the differential equation $\frac{\mathrm{d}}{\mathrm{d}t}\hat{\mathbf{v}}(t) = \hat{A}\hat{\mathbf{v}}(t)$ instead, and merely have to express the initial value $\mathbf{v}_0$ in the new basis as $\hat{\mathbf{v}}_0 = C\mathbf{v}_0$. The evaluation of the exponential function on the diagonal matrix A is easy (as in the previous example):

$$\hat{A} = \begin{pmatrix} \lambda_1 & 0 & \cdots & 0 \\ 0 & \lambda_2 & & 0 \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_n \end{pmatrix} \Rightarrow e^{t\hat{A}} = \begin{pmatrix} e^{t\lambda_1} & 0 & \cdots & 0 \\ 0 & e^{t\lambda_2} & & 0 \\ \vdots & & \ddots & 0 \\ 0 & \cdots & 0 & e^{t\lambda_n} \end{pmatrix}.$$

We correspondingly find that the solution (in the adapted basis) is

$$\hat{\mathbf{v}}(t) = e^{t\hat{A}}\,\hat{\mathbf{v}}_0,$$

which then transforms back into the original basis according to

$$\mathbf{v}(t) = \big[C^{-1}e^{t\hat{A}}\,C\big]\mathbf{v}_0.$$

Convention: in this chapter, unless otherwise noted, we fix a finite-dimensional $\mathbb{F}$-vector space V of dimension greater than 0 throughout. We shall occasionally look at specific examples and specify concrete fields $\mathbb{F}$.

1.1 Eigenvectors and eigenvalues

Recall from chapter 3 in part I that $\mathrm{Hom}(V, V)$ stands for the space of all linear maps $\varphi\colon V \to V$ (vector space homomorphisms from V to V, i.e., endomorphisms). Recall that w.r.t. to the natural addition and scalar multiplication operations, $\mathrm{Hom}(V, V)$ forms an $\mathbb{F}$-vector space; while w.r.t. addition and composition it forms a ring. This structure is isomorphic to the corresponding structure on the space $\mathbb{F}^{(n,n)}$ of all $n \times n$ matrices over $\mathbb{F}$, which forms an $\mathbb{F}$-vector space w.r.t. component-wise addition and scalar multiplication; and a ring w.r.t. addition and matrix multiplication. Isomorphisms between $\mathrm{Hom}(V, V)$ and $\mathbb{F}^{(n,n)}$ are induced by choices of bases for V. For any fixed labelled basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ of V, the association between $\varphi \in \mathrm{Hom}(V, V)$ and $A = [\![\varphi]\!]_B^B$ provides a bijection that is compatible with all the above operations, and hence forms both a vector space isomorphism and a ring isomorphism.

Recall also from chapter 3 in part I how different choices of bases for V, say $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ and $\hat{B} = (\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_n)$, induce different matrix representations of endomorphisms. The relationship between the matrices $A = [\![\varphi]\!]_B^B$ and $\hat{A} = [\![\varphi]\!]_{\hat{B}}^{\hat{B}}$ is governed by the regular change of basis matrix C and its inverse C^{-1}

$$\hat{A} = CAC^{-1}.$$

Here $C = [\![\mathrm{id}_V]\!]_{\hat{B}}^B$ is the matrix representation of the identity w.r.t. bases B in the domain and $\hat{B}$ in the range; C^{-1} correspondingly is $[\![\mathrm{id}_V]\!]_B^{\hat{B}}$. Recall

from section 3.3.3 in part I that matrices A and $\hat{A}$ in such a relationship are called *similar*; and all matrices that occur as representations of a fixed endomorphism φ in this fashion precisely make up a similarity class of matrices.

Example 1.1.1 Consider the endomorphism φ of $\mathbb{R}^2$ that is a reflection in the axis through $(1,1) = \mathbf{e}_1 + \mathbf{e}_2$. Its matrix representations w.r.t. the standard basis and the basis formed by $\mathbf{e}_1 + \mathbf{e}_2$ and $\mathbf{e}_2 - \mathbf{e}_1$ are

$$A = [\![\varphi]\!]_B^B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \hat{A} = [\![\varphi]\!]_{\hat{B}}^{\hat{B}} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$
$$B = (\mathbf{e}_1, \mathbf{e}_2) \qquad\qquad\qquad \hat{B} = (\mathbf{e}_1 + \mathbf{e}_2, \mathbf{e}_1 - \mathbf{e}_2).$$

The 1-dimensional subspaces spanned by $\mathbf{e}_1 + \mathbf{e}_2$ and by $\mathbf{e}_2 - \mathbf{e}_1$, respectively, are both preserved by φ, as both vectors are eigenvectors.

In contrast, the rotation through 90 degrees $(\pi/2)$, whose representation w.r.t. the standard basis is

$$A' = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix},$$

has no eigenvectors (no 1-dimensional subspaces mapped into themselves). Correspondingly there is no basis w.r.t. which it would be represented by a diagonal matrix. In other words, while A is similar to the diagonal matrix $\hat{A}$, the matrix A' is not similar to any diagonal matrix.

Definition 1.1.2 Let $\varphi \in \operatorname{Hom}(V, V)$.

(i) A vector $\mathbf{v} \in V$ is an *eigenvector* [Eigenvektor] of φ iff $\mathbf{v} \neq \mathbf{0}$ and $\varphi(\mathbf{v}) = \lambda \mathbf{v}$ for some $\lambda \in \mathbb{F}$.

(ii) A scalar λ is an *eigenvalue* [Eigenwert] for φ iff there is an eigenvector $\mathbf{v}$ with $\varphi(\mathbf{v}) = \lambda \mathbf{v}$.

One also speaks of an eigenvector $\mathbf{v}$ with eigenvalue λ, or of an eigenvector/eigenvalue pair $\mathbf{v}/\lambda$ in connection with an eigenvector $\mathbf{v}$ with $\varphi(\mathbf{v}) = \lambda \mathbf{v}$.

Observation 1.1.3 *φ may have 0 as an eigenvalue; the eigenvectors for eigenvalue 0 are precisely the the non-trivial elements of the kernel.*

If $\mathbf{v}$ is an eigenvector of φ, then the corresponding eigenvalue λ is uniquely determined (note that $\mathbf{v} \neq \mathbf{0}$). In contrast, if $\mathbf{v}$ is an eigenvector for eigenvalue λ, then so is any $\mu \mathbf{v}$ for $\mu \neq 0$.

φ may also have several linearly independent eigenvectors with the same eigenvalue; e.g. $\varphi = \operatorname{id}_V$ has all vectors in $V \setminus \{\mathbf{0}\}$ as eigenvectors for eigenvalue 1.

The problems of finding (eigenvectors and) eigenvalues for given endomorphisms is referred to as the *eigenvalue problem*[Eigenwertproblem].

In fact, the eigenvectors (if any) for a given eigenvalue λ almost form a subspace – we just need to fill in $\mathbf{0}$.

Definition 1.1.4 For $\varphi \in \mathrm{Hom}(V, V)$ and $\lambda \in \mathbb{F}$, the *eigenspace* [Eigenraum] w.r.t. λ is the subspace

$$V_\lambda := \{\mathbf{v} \in V : \varphi(\mathbf{v}) = \lambda \mathbf{v}\} \subseteq V.$$

If $V_\lambda \neq \{\mathbf{0}\}$, then its dimension $\dim(V_\lambda)$ is called the (geometric) *multiplicity* of the eigenvalue λ.

Exercise 1.1.1 Check that
  (i) V_λ as defined above is indeed a subspace of V.
  (ii) $V_0 = \ker(\varphi)$.
 (iii) λ is an eigenvalue of φ iff $V_\lambda \neq \{\mathbf{0}\}$.
 (iv) if λ is an eigenvalue of φ, then the eigenvectors with eigenvalue λ are precisely the vectors in $V_\lambda \setminus \{\mathbf{0}\}$.
  (v) if λ is an eigenvalue of φ, then the restriction of φ to V_λ is $\lambda \, \mathrm{id}_{V_\lambda}$.

These definitions and observations lead to the following reformulation of the eigenvalue (and eigenvector) problems.

Proposition 1.1.5 *The eigenspace V_λ is the kernel of the endomorphism $\varphi - \lambda \, \mathrm{id}_V$:*
$$V_\lambda = \ker\left(\varphi - \lambda \, \mathrm{id}_V\right).$$

Proof.   This is obvious from the definitions, noting that $\varphi - \lambda \, \mathrm{id}_V$ is the linear map

$$\begin{aligned} \varphi - \lambda \, \mathrm{id}_V : V &\longrightarrow V \\ \mathbf{v} &\longmapsto \varphi(\mathbf{v}) - \lambda \mathbf{v}. \end{aligned}$$

Clearly $\varphi(\mathbf{v}) = \lambda \mathbf{v}$ iff $\varphi(\mathbf{v}) - \lambda \mathbf{v} = \mathbf{0}$ iff $\mathbf{v} \in \ker(\varphi - \lambda \, \mathrm{id}_V)$.
$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

If $A = [\![\varphi]\!]_B^B \in \mathbb{F}^{(n,n)}$ is the matrix representation of φ w.r.t. any chosen basis B of V, then the map $\varphi - \lambda \, \mathrm{id}_V$ is represented — w.r.t. the same basis — by the matrix $A - \lambda E_n$, where E_n is the n-dimensional unit matrix. Whether

or not this map has a non-trivial kernel is determined by its rank, and hence by its determinant $\det(A - \lambda E_n) = |A - \lambda E_n| = \det(\varphi - \lambda \mathrm{id})$. Compare section 4.1 in part I for this. Thinking of λ as a parameter in this condition, we replace it by a variable x say, and regard the resulting determinant as a polynomial in this variable.

Definition 1.1.6 For $A \in \mathbb{F}^{(n,n)}$, the polynomial obtained as the determinant of the matrix $A - xE_n$,

$$
p_A = \det(A - xE_n) = \begin{vmatrix} (a_{11} - x) & a_{12} & a_{13} & \ldots & a_{1n} \\ a_{21} & (a_{22} - x) & a_{23} & \ldots & a_{2n} \\ \vdots & & \ddots & & \vdots \\ a_{n1} & & \ldots & & (a_{nn} - x) \end{vmatrix}
$$

is called the *characteristic polynomial* [charakteristisches Polynom] of A.

Some coefficients of the characteristic polynomial are easy to compute.

Exercise 1.1.2 Show that for $A \in \mathbb{F}^{(n,n)}$, if $p_A = a_n X^n + a_{n-1} X^{n-1} + \cdots + a_0$, then

$$
a_n = (-1)^n \quad , \quad a_{n-1} = (-1)^{n-1} \sum_i a_{ii} \quad , \quad a_0 = |A|.
$$

Observation 1.1.7 *If $\hat{A} = CAC^{-1}$ for a regular matrix C, then $p_A = p_{\hat{A}}$. This follows from the observation that*

$$
\hat{A} - xE_n = CAC^{-1} - xE_n = CAC^{-1} - xCE_nC^{-1} = C(A - xE_n)C^{-1},
$$

and therefore $|\hat{A} - xE_n| = |C||A - xE_n||C^{-1}| = |C||A - xE_n||C|^{-1} = |A - xE_n|$.

Thus, similar matrices have the same characteristic polynomial, and we may associate the characteristic polynomial directly with φ rather than with its matrix representations.

We shall put all considerations involving polynomials like p_φ on a more formal footing in the next section.

Definition 1.1.8 The *characteristic polynomial* of $\varphi \in \mathrm{Hom}(V, V)$ is $p_\varphi = |A - xE_n|$, where $A = [\![\varphi]\!]_B^B$ [for any choice of basis B for V, with result independent of that choice].

Theorem 1.1.9 *For any $\varphi \in \mathrm{Hom}(V, V)$ and $\lambda \in \mathbb{F}$: λ is an eigenvalue of φ iff $p_\varphi(\lambda) = 0$, i.e., if λ is a zero of the characteristic polynomial p_φ.*

Proof.   Let B be any basis of V, and $A = [\![\varphi]\!]_B^B$ be the corresponding matrix representation of φ. By the above, λ is an eigenvalue iff $V_\lambda \neq \{\mathbf{0}\}$ iff $\ker(\varphi - \lambda\,\mathrm{id}_V) \neq \{\mathbf{0}\}$ iff $A - \lambda E_n$ is not regular iff $|A - \lambda E_n| = 0$ iff $p_A(\lambda) = 0$ iff $p_\varphi(\lambda) = 0$.

$\square$

Definition 1.1.10 Let $\varphi \in \mathrm{Hom}(V, V)$. A subspace $U \subseteq V$ is called an *invariant subspace for φ*, or *invariant under φ*, if $\varphi(\mathbf{u}) \in U$ for all $\mathbf{u} \in U$.

Observation 1.1.11 *Any eigenspace V_λ of φ is an invariant subspace of φ.*

Exercise 1.1.3 Consider a rotation in $\mathbb{R}^3$ through angle α about the axis spanned by $\mathbf{a} \in \mathbb{R}^3 \setminus \{\mathbf{0}\}$ [w.l.o.g. fix $\mathbf{a} = \mathbf{e}_3$]. Determine *all* all its invariant subspaces

  (i) for $0 < \alpha < \pi$.

  (ii) for $\alpha = \pi$.

Lemma 1.1.12 *Suppose $V = U_1 \oplus U_2$ is the direct sum of invariant subspaces U_1 and U_2 for $\varphi \in \mathrm{Hom}(V, V)$. Then there is a basis for V such that φ is represented by a matrix A of the following block shape*

$$A = \begin{pmatrix} B_1^{(n_1, n_1)} & \mathbf{0}^{(n_1, n_2)} \\ \mathbf{0}^{(n_2, n_1)} & B_2^{(n_2, n_2)} \end{pmatrix}.$$

An analogous result obtains if V splits into the direct sum of three or more invariant subspaces.

Proof.   Let $\dim(U_i) = n_i$ for $i = 1, 2$. Choosing bases $(\mathbf{b}_1^{(i)}, \ldots, \mathbf{b}_{n_i}^{(i)})$ for for U_i, we join them to obtain a basis $B = (\mathbf{b}_1^{(1)}, \ldots, \mathbf{b}_{n_1}^{(1)}, \mathbf{b}_1^{(2)}, \ldots, \mathbf{b}_{n_2}^{(2)})$ for $V = U_1 \oplus U_2$.

Let $\varphi_i \in \mathrm{Hom}(U_i, U_i)$ be the restriction of φ to U_i; $i = 1, 2$. If B_i is the matrix representation of φ_i, then the representation of φ w.r.t. B is as claimed.

$\square$

We look at the relationship between eigenspaces w.r.t. different eigenvalues.

Proposition 1.1.13 *Let $\varphi \in \mathrm{Hom}(V, V)$ and suppose $\lambda_1, \ldots, \lambda_m$ are distinct eigenvalues of φ. Then the sum of the corresponding eigenspaces V_{λ_i} for $i = 1, \ldots, m$ is direct, i.e., $V_{\lambda_1} + V_{\lambda_2} + \cdots + V_{\lambda_m} = V_{\lambda_1} \oplus V_{\lambda_2} \oplus \cdots \oplus V_{\lambda_m}$.*

Proof.   By induction on $m \geqslant 2$. Consider the case of $m = 2$. We need to show that $V_{\lambda_1} \cap V_{\lambda_2} = \{\mathbf{0}\}$. Let $\mathbf{v} \in V_{\lambda_1} \cap V_{\lambda_2}$. Then, as $\mathbf{v} \in V_{\lambda_1}$, $\varphi(\mathbf{v}) = \lambda_1 \mathbf{v}$; and, as also $\mathbf{v} \in V_{\lambda_2}$, $\varphi(\mathbf{v}) = \lambda_2 \mathbf{v}$. Hence $\lambda_1 \mathbf{v} = \lambda_2 \mathbf{v}$ and $(\lambda_1 - \lambda_2)\mathbf{v} = \mathbf{0}$. As $\lambda_1 \neq \lambda_2$, this implies $\mathbf{v} = \mathbf{0}$.

The induction step is similar. Assume the claim is true for a sum of less than m distinct eigenspaces. We then show that also $V_{\lambda_1} \cap \sum_{i=2}^{m} V_{\lambda_i} = \{\mathbf{0}\}$. Notice, that by the inductive hypotheses, we already have that $\sum_{i=2}^{m} V_{\lambda_i} = \bigoplus_{i=2}^{m} V_{\lambda_i}$, which in particular implies that any $\mathbf{v} \in \sum_{i=2}^{m} V_{\lambda_i}$ has a unique decomposition as $\mathbf{v} = \sum_{i=2}^{m} \mathbf{u}_i$ with $\mathbf{u}_i \in V_{\lambda_i}$.

So let $\mathbf{v} \in V_{\lambda_1} \cap \sum_{i=2}^{m} V_{\lambda_i}$. Then $\mathbf{v} = \sum_{i=2}^{m} \mathbf{u}_i$ with $\mathbf{u}_i \in V_{\lambda_i}$ and by linearity, $\varphi(\mathbf{v}) = \sum_{i=2}^{m} \varphi(\mathbf{u}_i) = \sum_{i=2}^{m} \lambda_i \mathbf{u}_i$. On the other hand, as $\mathbf{v} \in V_{\lambda_1}$, $\varphi(\mathbf{v}) = \lambda_1 \mathbf{v}$. Hence $\sum_{i=2}^{m}(\lambda_1 - \lambda_i)\mathbf{u}_i = \mathbf{0}$. Uniqueness of decompositions in $\bigoplus_{i=2}^{m} V_{\lambda_i}$, applied to $\mathbf{0}$, implies that $(\lambda_1 - \lambda_i)\mathbf{u}_i = \mathbf{0}$ for $i = 2, \ldots, m$, and as $\lambda_i \neq \lambda_1$, it follows that the $\mathbf{u}_i$ are all equal to $\mathbf{0}$. Therefore $\mathbf{v} = \mathbf{0}$.

$\square$

Remark 1.1.14 If $\dim(V) = n$, it follows that $\varphi \in \mathrm{Hom}(V, V)$ can have at most n distinct eigenvalues. This is because each eigenvalue λ gives rise to a non-trivial eigenspace $V_\lambda \neq \{\mathbf{0}\}$ of dimension $\dim(V_\lambda) \geqslant 1$. As the sum of distinct eigenspaces is direct, the dimension of this sum is the sum of their dimensions. If φ has m distinct eigenvalues, therefore, the sum of the corresponding eigenspaces has dimension at least m. So $m \leqslant n$ follows.

We can already describe some cases in which $\varphi \in \mathrm{Hom}(V, V)$ can be *diagonalised*, i.e., is represented by a diagonal matrix for a suitable choice of basis.

Proposition 1.1.15 *Let $\dim(V) = n$. Then each one of the following conditions guarantees that there is a basis of V w.r.t. which φ is represented by a diagonal matrix (φ is diagonalisable):*

(i) *φ has n distinct eigenvalues.*

(ii) *p_φ has n distinct zeroes.*

(iii) *φ has m distinct eigenvalues $\lambda_1, \ldots, \lambda_m$ such that $\sum_{i=1}^{m} \dim(V_{\lambda_i}) = n$.*

Proof. (i) and (ii) are equivalent by Theorem 1.1.9. Both describe a special case of condition (iii), namely where $m = n$ and hence $\dim(V_{\lambda_i}) = 1$.

So we concentrate on (iii). Firstly, the V_{λ_i} are invariant subspaces, and their direct sum is equal to V by the assumption about dimensions. By Lemma 1.1.12, it suffices to argue that the restriction φ_i of φ to V_{λ_i} admits a diagonal representation. Any basis for V_{λ_i} will do, as by definition of V_{λ_i}, $\varphi_i = \lambda_i \,\mathrm{id}$ is represented by $\lambda_i E_{n_i}$ if $n_i = \dim(V_{\lambda_i})$.

$\square$

1.2 Polynomials

The role of the characteristic polynomial in connection with Theorem 1.1.9 suggests that an algebraic understanding of polynomials in one variable holds the key to a more detailed analysis of diagonalisability. This section provides the basis for this.

Definition 1.2.1 Let $\mathbb{F}$ be a field. A *polynomial* [Polynom] over $\mathbb{F}$ in one variable X is a formal sum $p = \sum_{i=0}^{m} a_i X^i$ with coefficients $a_i \in \mathbb{F}$. We write $\mathbb{F}[X]$ for the set of all polynomials in variable X over $\mathbb{F}$.

Convention: X^0 is identified with $1 \in \mathbb{F}$, X^1 with X, and any X^i for $i \geqslant 2$ will be regarded as an m-fold product of X, once we endow $\mathbb{F}[X]$ with a multiplication operation. We speak of the exponent i in X^i as the *power* or *order* of X in that term. The coefficient a_i in $p = \sum_{i=0}^{m} a_i X^i$ is called the coefficient of order or power i. The non-zero coefficient of the highest order in p is called the *leading coefficient* [Leitkoeffizient] of p.

We identify all polynomials whose coefficients are all 0 (i.e., we do not distinguish between $0, 0+0X, 0+0X+0X^2$, etc), and regard them as different representations of the *null polynomial* or *zero polynomial*, written just 0. The null polynomial is the only polynomial without a leading coefficient.

A polynomial is called *normalised* if it is the null polynomial or else if its leading coefficient is 1.

Definition 1.2.2 The *degree* [Grad] of the polynomial $p = \sum_{i=0}^{m} a_i X^i \in \mathbb{F}[X]$, denoted $d(p)$ is the power of the leading coefficient, if p is not the null polynomial. For the null polynomial we put $d(0) := -\infty$. [1]

---

[1] This convention will have the advantage that it gives a degree to the null polynomial that is smaller than that of any other polynomial (even up to addition of any $n \in \mathbb{N}$).

Polynomials of degree 0 and the null polynomial, i.e., polynomials $p = a_0 X^0 = a_0$ are called *constant polynomials* and identified with $a_0 \in \mathbb{F}$. In this way we regard $\mathbb{F}$ as a subset of $\mathbb{F}[X]$.

Polynomials of degree 1, $p = a_0 + a_1 X$ with $a_1 \neq 0$, are called *linear*; those of degree 2, $p = a_0 + a_1 X + a_2 X^2$ with $a_2 \neq 0$, *quadratic*; etc.

It is important to note that we do *not* identify a polynomial $p = \sum_i a_i X^i$ in $\mathbb{F}[X]$ with the *polynomial function*

$$\begin{aligned}\check{p} \colon \mathbb{F} &\longrightarrow \mathbb{F} \\ \lambda &\longmapsto p(\lambda) := \sum_i a_i \lambda^i\end{aligned} \qquad (*)$$

which is an element of $\mathrm{Pol}(\mathbb{F}) \subseteq \mathcal{F}(\mathbb{F}, \mathbb{F})$ (familiar as an $\mathbb{F}$-vector space from last term).

In fact we need to keep the two notions separate, especially when dealing with finite fields. We saw, for instance, that over $\mathbb{F}_2$ the polynomial functions $p(x) = x^2$ and $p'(x) = x$ are the same; we do not, however, identify the polynomials X and X^2 in $\mathbb{F}_2[X]$. We shall see below that $\mathbb{F}[X]$ as well as $\mathrm{Pol}(\mathbb{F})$ carry structure as $\mathbb{F}$-vector spaces and as rings. With respect to both structures the association between formal polynomials and the polynomial functions they induce, is structure preserving in the sense of a (vector space or ring) homomorphism, but is not injective.

Consider, for instance, the $\mathbb{F}_2$-vector spaces $\mathrm{Pol}(\mathbb{F}_2)$ and $\mathbb{F}_2[X]$. $\mathrm{Pol}(\mathbb{F}_2) = \mathcal{F}(\mathbb{F}_2, \mathbb{F}_2)$ has dimension 2 (four elements), but $\mathbb{F}_2[X]$ will be infinite-dimensional. In particular, $p_0 = 1$, $p_1 = X$, $p_2 = X^2$, …, are all distinct and in fact linearly independent in $\mathbb{F}_2[X]$.

Exercise 1.2.1 Check that $\check{\ } \colon \mathbb{F}[X] \to \mathrm{Pol}(\mathbb{F})$ is bijective for $\mathbb{F} = \mathbb{R}$, but not for any $\mathbb{F}_p$ (p a prime).

We shall later look at evaluations of polynomials $p \in \mathbb{F}[X]$ not just over $\mathbb{F}$, but also over the ring $\mathrm{Hom}(V, V)$, or over the ring $\mathbb{F}^{(n,n)}$. Writing $p(\varphi)$ or $p(A)$ (which give values to be defined below) we may regard these as values of corresponding 'polynomial functions' analogous to $\check{p}$ but over domains other than $\mathbb{F}$. Generally we shall suppress the $\check{p}$ notation in all these cases and also return to writing just $p(\lambda)$ for the value of $\check{p}$ on argument λ.

1.2.1   Algebra of polynomials

Addition of polynomials.   Addition of formal polynomials in $\mathbb{F}[X]$ is defined in component-wise fashion. If $p = \sum_{i=0}^{m} a_i X^i$ and $q = \sum_{i=0}^{n} b_i X^i$, we

may firstly assume that $n = m$ by extending the polynomial of lower degree with coefficients 0 as necessary. We then put

$$p + q = \Big(\sum_{i=0}^{m} a_i X^i\Big) + \Big(\sum_{i=0}^{m} b_i X^i\Big) := \sum_{i=0}^{m}(a_i + b_i)X^i.$$

Note that $d(p + q) \leqslant \max(d(p), d(q))$. The degree may indeed drop through cancellation of coefficients: in particular, if $b_i = -a_i$ for all i, then they add up to the null polynomial.

In fact $\mathbb{F}[X]$ carries the structure of an $\mathbb{F}$-vector space, with a correspondingly defined component-wise scalar multiplication. Since $\mathbb{F} \subseteq \mathbb{F}[X]$ via constant polynomials, this multiplication may, however, also be considered as a special case of the more general multiplication between polynomials to be considered below.

Exercise 1.2.2 Define component-wise scalar multiplication similarly, and verify that this, together with addition, turns $\mathbb{F}[X]$ into an $\mathbb{F}$-vector space with null vector 0 (the null polynomial).

Multiplication of polynomials.   We define a multiplication operation on $\mathbb{F}[X]$ as follows. If $p = \sum_{i=0}^{m} a_i X^i$ and $q = \sum_{i=0}^{n} b_i X^i$, we put

$$p \cdot q = \Big(\sum_{i=0}^{m} a_i X^i\Big)\Big(\sum_{i=0}^{n} b_i X^i\Big) := \sum_{k=0}^{m+n}\Big(\sum_{i+j=k} a_i \cdot b_j\Big)X^k.$$

Note that this multiplication gives what one would expect when applying distributivity and commutativity rules together with the obvious $X^i X^j := X^{i+j}$ and re-grouping terms w.r.t. to orders X^k.

Observation 1.2.3 *Multiplication of polynomials is additive w.r.t. degrees:* $d(pq) = d(p) + d(q)$.

In case at least one of p or q is the null polynomial, we extend the convention regarding $d(0) = -\infty$ by the "natural" crutches that $-\infty + n = n + (-\infty) = -\infty + (-\infty) = -\infty$.

Proposition 1.2.4 $(\mathbb{F}[X], +, \cdot, 0, 1)$ *forms a commutative ring with neutral element* 0 *(the null polynomial) for addition, and neutral element* 1 *(the constant polynomial* 1*) for multiplication.*

Proof.   Check the axioms, compare section 1.3.3 in part I.

□

A *scalar multiplication* $\mathbb{F} \times \mathbb{F}[X] \to \mathbb{F}[X]$ is obtained as the restriction of the above multiplication of polynomials to the case where the first polynomial is a constant polynomial $p = a_0 \in \mathbb{F}$.

Proposition 1.2.5 *W.r.t. addition and the induced scalar multiplication of polynomials, $\mathbb{F}[X]$ forms an $\mathbb{F}$-vector space with null vector 0 (the null polynomial).*

Proof.   Check the axioms.

□

Proposition 1.2.6 *The map $\check{\ } : \mathbb{F}[X] \to \mathrm{Pol}(\mathbb{F})$ is compatible with the operations of addition and multiplication and thus constitutes*

   (a) *an $\mathbb{F}$-vector space homomorphism w.r.t. to the vector space structure of $\mathbb{F}[X]$ and $\mathrm{Pol}(\mathbb{F})$.*

   (b) *a ring homomorphism w.r.t. to the ring structure of $\mathbb{F}[X]$ and $\mathrm{Pol}(\mathbb{F})$.*

Exercise 1.2.3 Show that the $\mathbb{F}$-vector space $\mathbb{F}[X]$ is infinite-dimensional, by showing that the $p_i = X^i$ for $i \in \mathbb{N}$ are linearly independent.
   Compare this with $\mathrm{Pol}(\mathbb{F})$, for instance for $\mathbb{F} = \mathbb{F}_2$.

The following says that the product of two non-zero polynomials cannot be zero in $\mathbb{F}[X]$. Rings with this property are called *integral domains* [Integritätsbereiche]. That not all rings share this property, can be seen for instance in $\mathbb{Z}_q$ for q not a prime. Fields on the other hand always satisfy this condition (why?).

Proposition 1.2.7 *For any $p, q \in \mathbb{F}[X]$: $pq = 0 \;\Rightarrow\; p = 0$ or $q = 0$.*

Proof.   This is a consequence of Observation 1.2.3 on degrees under multiplication. In more detail: suppose $p, q \neq 0$; then $d(p), d(q) \geqslant 0$ and the corresponding coefficients $a_{d(p)}$ in p and $b_{d(q)}$ in q are both different from 0. Hence the coefficient of order $d = d(p) + d(q)$, which is $a_{d(p)} \cdot b_{d(q)}$, is different from zero, whence $d(pq) \geqslant 0$ and the product cannot be the null polynomial.
□

Exercise 1.2.4 Show that the only elements of $\mathbb{F}[X]$ that have an inverse w.r.t. multiplication are the non-zero constant polynomials (i.e., those in $\mathbb{F} \setminus \{0\} \subseteq \mathbb{F}[X]$).

1.2.2   Division of polynomials

An interesting topic in rings generally is *divisibility*. Over the familiar ring $(\mathbb{Z}, +, \cdot, 0, 1)$ of the integers, the investigation of divisibility gives rise to the notions of division with remainder, greatest common divisors, and prime numbers, among others. We shall encounter similar notions in the ring $\mathbb{F}[X]$.

The ring $\mathbb{F}[X]$ fails to be a field (check that $X \in \mathbb{F}[X]$ does not have a multiplicative inverse). Therefore, the question of whether for given p, q there is a polynomial r such that $q \cdot r = p$ is non-trivial.

We first define division with remainder for polynomials.

Definition 1.2.8 Let $p, q \in \mathbb{F}[X]$, $q \neq 0$. Then $s \in \mathbb{F}[X]$ is the result of *dividing p by q with remainder $r \in \mathbb{F}[X]$* if

$$p = sq + r \quad \text{where} \quad d(r) < d(q).$$

Lemma 1.2.9 *For $p, q \in \mathbb{F}[X]$ as above, p can be divided by q with remainder, with unique results $s, r \in \mathbb{F}[X]$.*

Proof.   We first consider the case that $d(p) < d(q)$. Then $s = 0$ and $r = p$ are admissible results. To see that they are uniquely determined, note that if $s \neq 0$ then necessarily $d(sq) \geqslant d(q)$. As $d(r) < d(q)$ is a requirement, we find that necessarily $d(sq + r) = d(sq) \geqslant d(q)$ (look at the leading coefficients). But then $d(sq + r) > d(p)$ shows that $p \neq sq + r$.

We now prove the general case by induction on $d(p) \geqslant 0$. The base case, for $d(p) = 0$, is easy. So we assume $d(p) \geqslant 1$ and that the claim has been established for all p' of degree less than $n = d(p)$. If $d(q) > n$ we are done by the above. So let $d(q) := m \leqslant n$. Let $p = \sum_{i=0}^{n} a_i X^i$ and $q = \sum_{j=0}^{m} b_j X^j$. If $p = sq + r$ with $d(r) < d(q)$ then $d(s)$ must be $k := n - m$. So we know that any result must be of the form $s = c_k X^k + s'$ where $d(s') < k$, $c_k \neq 0$.

Now $sq + r = (c_k X^k + s')q + r$ has leading coefficient $c_k b_m$ of order $m + k = n$. Necessarily therefore $c_k = a_n / b_m$. The task of finding s and r such that $p = sq + r$ with $d(r) < d(q)$ now reduces to finding s' and r' such that $p' = s'q + r'$ with $d(r') < d(q)$ for $p' = p - (a_n / b_m) X^k q$. Since p' is a polynomial of degree $d(p') < n$, there is a unique solution according to the inductive hypothesis.

$\square$

Definition 1.2.10 For $p, q \in \mathbb{F}[X]$, $q \neq 0$: q *divides* p, denoted $q|p$, iff $p = sq$ for some $s \in \mathbb{F}[X]$ (with remainder 0). We explicitly also include the case where $p = 0$: any non-zero polynomial divides the null polynomial 0. A polynomial that divides p is called a *divisor*.
    We do not regard the null polynomial as a divisor of any polynomial.

    Note that whenever q divides p then so does cq for any $c \in \mathbb{F} \backslash \{0\}$: divisors are best considered up to multiplication with non-zero constant polynomials because divisibility in the field $\mathbb{F}$ is trivial.

Exercise 1.2.5 Show for $p, q \in \mathbb{F}[X]$: if $p|q$ and $q|p$ then p and q are constant multiples of each other, i.e., $p = cq$ and $q = c^{-1}p$ for some $c \in \mathbb{F} \setminus \{0\}$. Hint: look at the degrees.

    We now link divisibility by simple degree 1 polynomials (linear factors, roots) to the existence of zeroes. While zeroes are defined in terms of the associated polynomial function, we thus link them to divisibility in $\mathbb{F}[X]$.

Definition 1.2.11 Let $p \in \mathbb{F}[X]$ and $\lambda \in \mathbb{F}$.
   (i) $\lambda \in \mathbb{F}$ is a *zero* [Nullstelle] of p iff $p(\lambda) = 0$.
  (ii) p has the *linear factor* [Linearfaktor] $(X - \lambda)$, or p has the *root* [Wurzel] λ, iff $(X - \lambda)$ divides p, i.e., iff $p = (X - \lambda)s$ for some $s \in \mathbb{F}[X]$.
 (iii) If λ is a root of p, then the *algebraic multiplicity* of λ is the maximal m such that $(X - \lambda)^m$ divides p.

Proposition 1.2.12 *For $p \in \mathbb{F}[X]$ of degree at least 1, and for any $\lambda \in \mathbb{F}$: λ is a zero of p iff p has $(X - \lambda)$ as a linear factor.*

Proof.   Clearly, if $p = (X - \lambda)s$, then $p(x) = (x - \lambda)s(x)$ and thus $p(\lambda) = 0$.
    Conversely, assume that $p(\lambda) = 0$. As $d(p) \geqslant 1$, we may divide p by $q = (X - \lambda)$ with remainder: $p = (X - \lambda)s + r$ where $d(r) < 1$ means that $r \in \mathbb{F}$. Hence $p(x) = (x - \lambda)s(x) + r$ and $p(\lambda) = r$. So $r = 0$ as λ is a zero of p, and therefore $(X - \lambda)$ divides p.
                                                                                   □

Exercise 1.2.6 Determine for which $\lambda \in \mathbb{R}$ the linear factor $X - \lambda$ divides the polynomial $p = X^3 + 2X^2 + 4X + 8 \in \mathbb{R}[X]$. One way of doing this is to compare coefficients in p and $(X - \lambda)(X^2 + \alpha X + \beta)$ and making suitable case distinctions regarding possible values for α, β, λ.

Primes are numbers that are not non-trivially divisible, or numbers that are irreducible by integer division. Irreducible polynomials are similarly defined. Trivial divisibility here concerns products involving a constant polynomial.

Definition 1.2.13 A polynomial p of degree $p \geqslant 1$ is called *irreducible* [irreduzibel] iff it is not divisible by any polynomial q of degree $0 < d(q) < p$.

Example 1.2.14 Any degree 1 polynomial, and in particular any linear factor $(X - \lambda)$, is irreducible.

Example 1.2.15 Over the real field $\mathbb{R}$, there are also irreducible polynomials of degree 2, like $p = X^2 + 1$. The same polynomial *is* reducible when viewed as a polynomial over the field $\mathbb{C}$: $X^2 + 1 = (X - i)(X + i)$ in $\mathbb{C}[X]$. Any polynomial of odd degree over $\mathbb{R}$ has a zero (its graph crosses the x-axis) and hence, by the last Proposition, is divisible by a linear factor; hence any polynomial of odd degree greater than 1 over $\mathbb{R}$ is reducible. More on irreducibility in $\mathbb{C}[X]$ and $\mathbb{R}[X]$ is presented in section 1.2.3 below.

Exercise 1.2.7 Consider $\mathbb{F}_2[X]$. Show that any non-linear polynomial
  (i) without the constant term 1, or
  (ii) with an odd number of powers X^i for $i \geqslant 1$ (with or without the constant term 1)
is reducible in $\mathbb{F}_2[X]$. Find all irreducible polynomials of degree up to 4. [Note that all irreducible polynomials must describe the constant polynomial function 1, which is represented not just by the constant polynomial 1.]

The following notion is helpful in analysing divisibility questions in rings.

Definition 1.2.16 Let $(A, +, \cdot, 0, 1)$ be a commutative ring. A non-empty subset $I \subseteq A$ is an *ideal* [Ideal] if it is closed under addition and under multiplication with arbitrary ring elements:
  (i) $a, b \in I \implies a + b \in I$.
  (ii) $a \in I, r \in A \implies ra \in I$.
An ideal $I \subseteq A$ is a *principal ideal* [Hauptideal] if $I = I_a := \{ra : r \in A\}$ consists of all the multiples of a fixed element a (the generator of I).

Proposition 1.2.17 *Any ideal $I \subseteq \mathbb{F}[X]$ is a principal ideal, i.e., every ideal I in $\mathbb{F}[X]$ possesses a generator p such that $I = I_p = \{pr \colon r \in \mathbb{F}[X]\}$. Such p is uniquely determined by I up to multiplication by constants $c \in \mathbb{F}$; in particular there is a unique normalised p such that $I = I_p$.*

Proof.   Let $I \subseteq \mathbb{F}[X]$ be an ideal. If $I = \{0\}$, then 0 generates I.

Otherwise let $p \in I$ be an element of minimal degree in $I \setminus \{0\}$. We claim that $I = I_p$ for any such $p \in I$. Clearly $I_p \subseteq I$, as I is closed under arbitrary products with polynomials.

Let us show that $I \subseteq I_p$. Let $q \in I$. We may assume that $q \neq 0$ as $0 \in I_q$ anyway. We may divide q by p with remainder, $q = ps + r$ with $d(r) < d(p)$. But $r = q - pr \in I$, and hence $d(r) < d(p)$ implies that $r = 0$ by the choice of p, and hence $q \in I_p$.

For uniqueness up to constants: if $I \neq \{0\}$ and $I = I_p = I_q$ then $p|q$ and $q|p$ imply that they are constant multiples of each other.

$\square$

Definition 1.2.18 A *greatest common divisor* of two polynomials $r, s \in \mathbb{F}[X]$ is any polynomial $p \in \mathbb{F}[X]$ such that $p|r$, $p|s$ and for any other $q \in \mathbb{F}[X]$ that divides both r and r we have $q|p$. If the constant polynomial 1 is a greatest common divisor of r and s, then r and s are called *relatively prime*[teilerfremd].

Exercise 1.2.8 Show that any two distinct linear factors $(X - \lambda_1)$ and $(X - \lambda_2)$ in $\mathbb{F}[X]$ are relatively prime.

Lemma 1.2.19 *Any two non-constant polynomials possess a greatest common divisor q. This greatest common divisor r is unique up to multiplication by constants $c \in \mathbb{F} \setminus \{0\}$. If p is a greatest common divisor of s, t, then $p = gs + ht$ for suitable $g, h \in \mathbb{F}[X]$.*

Proof.   For given s, t consider $I := \{gs + ht \colon g, h \in \mathbb{F}[X]\}$. One checks that $I \subseteq \mathbb{F}[X]$ is an ideal. By the previous lemma, $I = I_p$ for some $p \in \mathbb{F}[X]$. We show that this p is a greatest common divisor of s and t.

Clearly $p|s$ and $p|t$ as $s, t \in I = I_p$. If $q|s$ and $q|t$ then, by distributivity, q divides any element $gs + ht$ of I, hence in particular also p.

For uniqueness up to constant factors, observe that any two greatest common divisors must divide each other.

$\square$

A fundamental property of primes with respect to (integer) divisibility is that if a prime divides a product of two numbers then it must divide at least one of those factors. A similar phenomenon obtains here.

Proposition 1.2.20 *Let q be irreducible, $q, r, s \in \mathbb{F}[X]$. If q divides rs then q divides r or q divides s.*

Proof.   Let q be irreducible and assume that q divides rs, w.l.o.g., $rs \neq 0$. As q is irreducible it is not constant, therefore at least one of r and s must also be non-constant. W.l.o.g. assume that $d(s) \geqslant 1$. Let p be a greatest common divisor of s and q. From the last lemma we know that p is of the form $p = gs + hq$ for suitable $g, h \in \mathbb{F}[X]$.

As $p|q$ there is $t \in \mathbb{F}[X]$ such that $q = pt$. But as q is irreducible, either p or t must be constant.

If $t = c$ is constant, then $q|p$ and hence $q|s$.

If $p = c$ is constant, then $c = gs + hq$ for suitable $g, h \in \mathbb{F}[X]$. Therefore $1 = (c^{-1}g)s + (c^{-1}h)q$ and $r = 1r = (c^{-1}g)rs + (c^{-1}h)qr$ is divisible by q, since rs and qr are both divisible by q.

$\square$

Remark 1.2.21 From the above, one can conclude that similar to the prime decomposition in the ring of integers, the ring $\mathbb{F}[X]$ admits (an essentially unique) decomposition into irreducible polynomials. Any polynomial $p \in \mathbb{F}[X] \setminus \{0\}$ has a representation as a product of irreducible polynomials. This decomposition is unique up to permutations and up to constants.

1.2.3   Polynomials over the real and complex numbers

What can we say about roots (or zeroes) of polynomials in $\mathbb{F}[X]$? The answer largely depends on $\mathbb{F}$. We collect the crucial facts for the familiar classical fields of the real and complex numbers, $\mathbb{R}$ and $\mathbb{C}$. These results go beyond the scope of this course.

Theorem 1.2.22 (Fundamental theorem of algebra)
Any non-constant polynomial in $\mathbb{C}[X]$ has a zero. Consequently, any non-constant polynomial $p \in \mathbb{C}[X]$ of degree $d(p) = n \geqslant 1$ splits into linear factors $p = z_0(X - z_1) \cdots (X - z_n)$ for constants $z_i \in \mathbb{C}$.

The first statement implies the second, via Proposition 1.2.12 and induction on the degree n.

Over $\mathbb{R}$ the situation is different. There are polynomials without zeroes, of any even degree. For instance, $p = (X^2 + 1)^m$ has degree $n = 2m$ and no zeroes. Any non-constant polynomial of odd degree, on the other hand, does have a zero and hence is divisible by a corresponding linear factor. [This can be derived via the intermediate value theorem, but also follows along the lines of the algebraic argument given below.] It turns out that the irreducible polynomials in $\mathbb{R}[X]$ are precisely the linear factors (of degree 1) and those degree 2 polynomials without zeroes, viz. the constant multiples of quadratic polynomials $p = X^2 + bX + c = (X + b/2)^2 + (c - b^2/4)$ for which $c > b^2/4$.

Theorem 1.2.23 *The irreducible polynomials in $\mathbb{R}[X]$ are precisely the linear polynomials and those quadratic polynomials that have no zeroes.*

The argument that any other non-constant real polynomial is reducible essentially follows from Theorem 1.2.22. Using Proposition 1.2.12 it remains to show that no polynomial p of even degree $n = 2m > 2$ is irreducible. For this we consider $\mathbb{R}[X]$ as a subset $\mathbb{R}[X] \subseteq \mathbb{C}[X]$. Over $\mathbb{C}$, p splits into linear factors. Up to a constant factor, $p = (X - z_1) \cdots (X - z_n)$ for $z_i \in \mathbb{C}$. But because $p \in \mathbb{R}[X]$, $p = \bar{p}$ where $\bar{p}$ is obtained by complex conjugation of all complex numbers in p, sending $z_j = x_j + iy_j$ to $\bar{z}_j = x_j - iy_j$. It therefore follows that all roots $z_j \notin \mathbb{R}$ of p come in complex conjugate pairs, i.e., p is a product of real linear factors and pairs of complex linear factors of the form $(X - z)(X - \bar{z})$. Any such pair, however, is still also a real polynomial of degree 2. If $z = x + iy$ then $(X - z)(X - \bar{z}) = X^2 - (z + \bar{z})X + z\bar{z} = X^2 - 2xX + (x^2 + y^2) \in \mathbb{R}[X]$. I.e., p splits into polynomials of degree 2 even in $\mathbb{R}[X]$. Some of these may indeed be irreducible over $\mathbb{R}$.

1.3   Upper triangle form

The irreducible factors of the characteristic polynomial determine whether φ can be represented by an upper triangle matrix or not.

A matrix $A = (a_{ij}) \in \mathbb{F}^{(n,n)}$ is an *upper triangle matrix* if $a_{ij} = 0$ for $1 \leqslant j < i \leqslant n$. This means that all non-zero entries are on the diagonal or above. [An echelon matrix is a special case of an upper triangle matrix.]

Proposition 1.3.1 *The following are equivalent for $\varphi \in \mathrm{Hom}(V, V)$ with characteristic polynomial p_φ:*

(i) p_φ splits into linear factors.

(ii) there is a basis for V such that φ is represented by an upper triangle matrix A.

Note in particular, that by the fundamental theorem of algebra, any endomorphism of a $\mathbb{C}$-vector space admits a representation by an upper triangle matrix.

Proof.   (ii) $\Rightarrow$ (i) is straightforward from the definition of $p_\varphi(x) = p_A(x) = |A - xE_n|$. Recall that the determinant of an upper triangle matrix equals the product of the diagonal entries, in this case of the linear factors $(a_{ii} - X)$.

For (i) $\Rightarrow$ (ii) we proceed by induction on the dimension n of V. The case of $n = 1$ is trivial.

For the induction step, assume $n > 1$ and that the claim is true in dimension $n - 1$.

Let $(X - \lambda)$ be a linear factor of $p = p_\varphi$, $p = (X - \lambda)p'$. As λ is a zero of p_φ, there is an eigenvector $\mathbf{b}_1$ with eigenvalue λ: $\varphi(\mathbf{b}_1) = \lambda\mathbf{b}_1$. We choose $\mathbf{b}_1$ as our first basis vector.

Extending to a basis $B = (\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n)$ of V, we obtain a matrix representation $A = [\![\varphi]\!]_B^B$ in which $a_{11} = \lambda$ and all other entries in the first column equal to 0.

$$A = [\![\varphi]\!]_B^B = \begin{pmatrix} \lambda & a_{12} & \ldots & a_{1n} \\ 0 & & & \\ \vdots & & A' & \\ 0 & & & \end{pmatrix}$$

Let $V' = \operatorname{span}(\mathbf{b}_2, \ldots, \mathbf{b}_n)$ and $V_1 = \operatorname{span}(\mathbf{b}_1)$ so that $V = V_1 \oplus V'$. Note that $B' = (\mathbf{b}_2, \ldots, \mathbf{b}_n)$ is a basis for V'.

Define auxiliary linear maps $\varphi' \in \operatorname{Hom}(V', V')$ and $\psi \in \operatorname{Hom}(V', V_1)$ as follows. For $\mathbf{v} \in V'$, the vector $\varphi(\mathbf{v}) \in V = V_1 \oplus V'$ has a unique decomposition into a sum of vectors from V_1 and V', respectively. Define φ' and ψ such that for $\mathbf{v} \in V'$:

$$\varphi(\mathbf{v}) = \psi(\mathbf{v}) + \varphi'(\mathbf{v}) \text{ where } \psi(\mathbf{v}) \in V_1 \text{ and } \varphi'(\mathbf{v}) \in V'.$$

One checks that φ' and ψ are indeed linear. W.r.t. basis B' of V', φ' is represented by the matrix A' which is A with first row and first column removed.

If we expand $p_\varphi(x) = |A - xE_n|$ w.r.t. the first column, we find that $p_\varphi = (\lambda - X)p'$, with $p' = p_{A'} = p_{\varphi'}$ the characteristic polynomial of φ'.

Now p' splits into linear factors since $p = (X - \lambda)p'$ does. By the inductive hypothesis, V' has a basis $(\mathbf{b}'_2, \ldots, \mathbf{b}'_n)$ w.r.t. which φ' is represented by an upper triangle matrix. In terms of φ' and the $\mathbf{b}'_i$ this means that for $2 \leqslant i \leqslant n$:

$$\varphi'(\mathbf{b}'_i) \in \operatorname{span}(\mathbf{b}'_2, \ldots, \mathbf{b}'_i).$$

Finally the basis $(\mathbf{b}_1, \mathbf{b}'_2, \ldots, \mathbf{b}'_n)$ is a basis for V that is as desired for φ. Indeed, for $i = 1$, $\varphi(\mathbf{b}_1) = \lambda\mathbf{b}_1 \in \operatorname{span}(\mathbf{b}_1)$; and for $2 \leqslant i \leqslant n$,

$$\varphi(\mathbf{b}'_i) = \varphi'(\mathbf{b}'_i) + \psi(\mathbf{b}'_i) \in \operatorname{span}(\mathbf{b}'_2, \ldots, \mathbf{b}'_i) + \operatorname{span}(\mathbf{b}_1) = \operatorname{span}(\mathbf{b}_1, \mathbf{b}'_2, \ldots, \mathbf{b}'_i).$$

$\square$

Remark: In the situation of the induction step, the auxiliary map φ' can be found as a quotient map. Let U be the one-dimensional subspace $U = \operatorname{span}(\mathbf{v}_1)$. Then an alternative view of φ' is as the map

$$\begin{aligned} \varphi' : V/U &\longrightarrow V/U \\ \mathbf{v} + U &\longmapsto \varphi(\mathbf{v}) + U. \end{aligned}$$

One checks that this map is well defined, and corresponds to the above via an isomorphism of $V/U = (V' \oplus U)/U$ with V'. Compare section 2.6 in part I.

Corollary 1.3.2 *Let V be an n-dimensional $\mathbb{C}$-vector space, $\varphi \in \operatorname{Hom}(V, V)$. Then there is a basis B for V such that φ is represented by an upper triangle matrix A w.r.t. to B.*

The situation is obviously very different over $\mathbb{R}$. We again look at the simple example of a rotation through $\pi/2$ in $\mathbb{R}^2$. As this map does not have a single eigenvector, there is no first basis vector for achieving triangle form.

1.4   The Cayley–Hamilton Theorem

The Cayley–Hamilton Theorem is one of the key results in the analysis of eigenvalues and eigenspaces. It makes a surprising connection between the characteristic polynomial $p_\varphi \in \mathbb{F}[X]$, whose zeroes are the eigenvalues, with

the linear map that we obtain when we substitute φ itself for X in p_φ, or A for X in p_φ where A is a matrix representation of φ. .

So far we have thought of evaluating polynomials $p \in \mathbb{F}[X]$ over $\mathbb{F}$ — this is precisely what $\check{p}\colon \mathbb{F} \to \mathbb{F}$ stood for. But it also makes sense to evaluate a polynomial $p \in \mathbb{F}[X]$ on an endomorphism of an $\mathbb{F}$-vector space or on a square matrix over $\mathbb{F}$. The required operations of scalar multiplication (of matrices or endomorphisms), of multiplication (of matrices) or composition (of endomorphisms), and of addition (of matrices or endomorphisms) are well defined, and moreover obey the familiar laws of arithmetic in a ring.

Remark 1.4.1 The rings $\operatorname{Hom}(V, V)$ or $\mathbb{F}^{(n,n)}$ are not commutative — in general we cannot expect that $\varphi \circ \psi = \psi \circ \varphi$ or that $AB = BA$. We shall here not notice this lack of commutativity, because we shall always work w.r.t. a fixed endomorphism φ (or matrix A) and its powers, because we look at polynomials in a single variable. All the ring arithmetic we shall encounter, therefore, takes place within sub-rings of the form

$$\left\{ p(\varphi)\colon p \in \mathbb{F}[X] \right\} \subseteq \operatorname{Hom}(V, V)$$
$$\text{or} \quad \left\{ p(A)\colon p \in \mathbb{F}[X] \right\} \subseteq \mathbb{F}^{(n,n)}.$$

In restriction to these, composition and matrix multiplication, respectively, are commutative.

Definition 1.4.2 Let V be an $\mathbb{F}$-vector space. The evaluation map for polynomials $p \in \mathbb{F}[X]$ on $\operatorname{Hom}(V, V)$ is defined to be compatible with the ring arithmetic of $\mathbb{F}[X]$ and $\operatorname{Hom}(V, V)$, based on the following:

$p(\varphi) = \mathbf{0}$ (null endomorphism)   for $p = 0$ (null polynomial)
$p(\varphi) = c \operatorname{id}_V$ for $p = c \in \mathbb{F}$ (a constant polynomial)
$p(\varphi) = \varphi$ for $p = X$.

Similarly, for the evaluation map on matrices in $\mathbb{F}^{(n,n)}$, based on

$p(A) = \mathbf{0}$ (null matrix)   for $p = 0$
$p(A) = cE_n$ for $p = c \in \mathbb{F}$
$p(A) = A$ for $p = X$.

Note that the extension to arbitrary polynomials $p \in \mathbb{F}[X]$ is uniquely determined by these stipulations and the requirement of compatibility with

the ring structure. For instance, consider $p = aX^3 + bX + c$ for $a, b, c \in \mathbb{F}$. If $\varphi \in \mathrm{Hom}(V, V)$, then $p(\varphi) = a\,\mathrm{id}_V \circ \varphi \circ \varphi \circ \varphi + b\,\mathrm{id}_V \circ \varphi + c\,\mathrm{id}_V = a\varphi^3 + b\varphi + c\,\mathrm{id}_V$; and if $A \in \mathbb{F}^{(n,n)}$, then $p(A) = aE_n AAA + bE_n A + cE_n = aA^3 + bA + cE_n$.

Moreover, if A happens to be the matrix representation of φ w.r.t. to some basis B of V, then $p(A)$ is the matrix representation, w.r.t. to B, of $p(\varphi)$.

Exercise 1.4.1 Check the last compatibility claim systematically by verifying it first for the basic cases (null and constant polynomials, and the polynomial X) and then showing that it is preserved under multiplication and addition of polynomials.

Before we proceed to the main theorem and its proof, we recall an important fact about matrices and determinants from last term, and convince ourselves that it lifts to the level of ring arithmetic in $\mathbb{F}[X]$ we need here.

Recall from chapter 4 in part I the matrices $A_{[ij]}$ obtained by deleting row i and column j in $A \in \mathbb{F}^{(n,n)}$. We used the determinants of these (up to a $+/-$ sign) as the coefficients in a matrix A' towards the construction of the inverse of A. More precisely, the coefficient in row i and column j of A' was

$$a'_{ij} = (-1)^{i+j}|A_{[ji]}|.$$

For this matrix we showed that

$$AA' = A'A = |A|E_n.$$

Looking at the proof, which was just a calculation based on multilinearity and antisymmetry of the determinant, we see that the same identity also obtains if the matrices involved have entries from some commutative ring rather than from a field. Up to that point we had had no occasion to use multiplicative inverses. (Of course, in order to obtain the inverse $A^{-1} = |A|^{-1}A'$ in the case of a regular matrix A we did work over a field.) We shall use the following in the proof of the Cayley–Hamilton Theorem below.

Observation 1.4.3 *Define the determinant function on $n \times n$ matrices over a commutative ring by the familiar $\sum_{\sigma \in S_n} \mathrm{sign}(\sigma) \cdots$ formula. Consider a matrix A over that ring with coefficients a_{ij}, $1 \leqslant i, j \leqslant n$, from that ring, and let, for $1 \leqslant i, j \leqslant n$,*

$$a'_{ij} = (-1)^{i+j}|A_{[ji]}|.$$

Then, for $1 \leqslant i, j \leqslant n$:

$$\sum_{k=1}^{n} a_{ik} a'_{ki} = |A| \qquad and \qquad \sum_{k=1}^{n} a_{ik} a'_{kj} = 0 \ \text{ for } i \neq j.$$

Exercise 1.4.2 Review the proof of the corresponding fact over a field $\mathbb{F}$, in section 4.2 of part I, and transfer it to the current situation. Discuss whether commutativity is essential.

Theorem 1.4.4 (Cayley–Hamilton Theorem) *Let p_φ be the characteristic polynomial of $\varphi \in \mathrm{Hom}(V, V)$. Then*

$$p_\varphi(\varphi) = \mathbf{0} \quad \text{(the null endomorphism).}$$

Similarly, for $A \in \mathbb{F}^{(n,n)}$ with characteristic polynomial p_A:

$$p_A(A) = \mathbf{0} \quad \text{(the null matrix).}$$

Proof.   The two statements are equivalent. We look at the matrix formulation. Let $p_A = \sum_{j=0}^{n} \alpha_j X^j$.

Recall that $p_A = |A - XE_n|$. Let $C = A - XE_n$ be the matrix with entries

$$a_{ij} - \delta_{ij} X$$

in row i and column j, where $\delta_{ij} = 1$ for $i = j$ and 0 otherwise. Note that C has entries in $\mathbb{F}[X]$. Let C' be the related matrix whose entry in row i and column j is

$$(-1)^{i+j} |C_{[ji]}|,$$

where the determinant is evaluated in the ring $\mathbb{F}[X]$ (!). By Observation 1.4.3 above,

$$(A - XE_n)C' = CC' = |C|E_n = p_A E_n.$$

Note that any entry $(-1)^{i+j} |C_{[ji]}|$ in C' is a polynomial of degree less than n, since $C_{[ji]}$ is an $(n-1) \times (n-1)$ matrix whose entries are polynomials of degree up to 1. One may therefore expand C' in terms of powers of X in the form

$$C' = \sum_{i=0}^{n-1} D_i X^i,$$

with matrices $D_i \in \mathbb{F}^{(n,n)}$ (all entries constant polynomials, hence in $\mathbb{F}$). Then

$$p_A E_n = (A - X E_n) C' = \sum_{i=0}^{n-1} (A D_i X^i - D_i X^{i+1}).$$

Comparing coefficients in respective powers of X in

$$p_A E_n = \sum_{j=0}^{n} \alpha_j E_n X^j = \sum_{i=0}^{n-1} (A D_i X^i - D_i X^{i+1}),$$

we find

order of X			
X^0	$\alpha_0 E_n$	$=$	$A D_0$
X^1	$\alpha_1 E_n$	$=$	$A D_1 - D_0$
X^2	$\alpha_2 E_n$	$=$	$A D_2 - D_1$
$\vdots$		$\vdots$	
X^n	$\alpha_n E_n$	$=$	$-D_{n-1}$

Multiplying the i-th equation in the table by A^i, summing them up and noting how the right-hand sides cancel, we get

$$\alpha_0 E_n + \alpha_1 A + \alpha_2 A^2 + \cdots + \alpha_n A^n = p_A(A) = \mathbf{0}.$$

$\square$

Corollary 1.4.5 *For any endomorphism $\varphi \in \mathrm{Hom}(V, V)$, V of dimension n: $\mathrm{id}_V, \varphi, \varphi^2, \ldots, \varphi^n$ are linearly dependent in $\mathrm{Hom}(V, V)$.*

Note that, by dimension comparison alone, we know that $\mathrm{id}_V, \varphi, \varphi^2, \ldots, \varphi^{n^2}$ cannot be linearly independent.

Proof.   Let $p = p_\varphi = \alpha_n X^n + \alpha_{n-1} X^{n-1} + \cdots + \alpha_1 X + \alpha_0$. Note that at least $\alpha_n = (-1)^n \neq 0$. Therefore the equation $p(\varphi) = \mathbf{0}$ provides a non-trivial linear combination of the null map, $\alpha_n \varphi^n + \alpha_{n-1} \varphi^{n-1} + \cdots + \alpha_1 \varphi + \alpha_0 \mathrm{id}_V = \mathbf{0}$. [Note that possibly $\mathrm{id}_V, \varphi, \varphi^2, \ldots, \varphi^n$ are not even pairwise distinct!]

$\square$

Example 1.4.6 Consider the matrix

$$A = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix},$$

or the associated endomorphism $\varphi = \varphi_A \colon \mathbb{R}^4 \to \mathbb{R}^4$ represented by A in terms of the standard basis.

The characteristic polynomial p_A is $p_A = (X-1)^3(X-2) = X^4 - 5X^3 + 9X^2 - 7X + 2$ with the two zeroes $\lambda_1 = 1$ and $\lambda_2 = 2$.

One checks that $p_A(A)$ is the null matrix. There is, however, also a polynomial of degree 3 that annihilates A in this manner: $q(A) = \mathbf{0}$ for $q = (X-1)^2(X-2)$. Note that q divides p, and one can also check that q is essentially the only proper divisor of p_A that annihilates A. For instance, $\big((X-1)(X-2)\big)(A) \neq \mathbf{0}$. [Below we shall understand such q as the so-called minimal polynomial.]

1.5   Minimal polynomial and diagonalisation

Whether or not φ has a representation by a diagonal matrix (whether there exists a basis of eigenvectors) is not directly determined by the characteristic polynomial. For this task we need to consider another polynomial associated with φ, the so-called minimal polynomial, which divides the characteristic polynomial.

In preparation for the definition of this polynomial, we check that for $\varphi \in \mathrm{Hom}(V, V)$,

$$I(\varphi) := \{p \in \mathbb{F}[X] \colon p(\varphi) = \mathbf{0}\} \subseteq \mathbb{F}[X]$$

is an ideal in $\mathbb{F}[X]$. $I(\varphi) \neq \emptyset$ as $0 \in I(\varphi)$, and $I(\varphi)$ is obviously closed under addition $(p, q \in I(\varphi) \Rightarrow p + q \in I(\varphi))$ and under multiplication with arbitrary $r \in \mathbb{F}[X]$ $(p \in I(\varphi), r \in \mathbb{F}[X] \Rightarrow rp \in I(\varphi))$. Therefore, by Proposition 1.2.17, $I(\varphi)$ is generated by a single polynomial.

Definition 1.5.1 Let $\varphi \in \mathrm{Hom}(V, V)$, V an n-dimensional $\mathbb{F}$-vector space. The *minimal polynomial* [Minimalpolynom] of φ is the normalised polynomial $q = q_\varphi \in \mathbb{F}[X]$ that generates the ideal $I(\varphi)$.

The minimal polynomial q_A for a matrix $A \in \mathbb{F}^{(n,n)}$ is analogously defined.

The following is essentially a corollary of the Cayley–Hamilton Theorem.

Proposition 1.5.2 *The minimal polynomial q_φ divides the characteristic polynomial p_φ. They have the same zeroes (roots) but possibly the algebraic multiplicity of some roots is smaller in q_φ.*

Proof. $q_\varphi | p_\varphi$ follows from the definition and the fact that $p_\varphi \in I$ by Cayley–Hamilton.

Any root in q_φ gives rise to an irreducible linear factor which must therefore also be a linear factor in p_φ, compare Proposition 1.2.20.

Conversely, a root λ of p_φ is an eigenvalue of φ. For a corresponding eigenvector $\mathbf{v}$ we have $\varphi^k(\mathbf{v}) = \lambda^k \mathbf{v}$. As $(q_\varphi(\varphi))(\mathbf{v}) = (q_\varphi(\lambda \operatorname{id}_V))(\mathbf{v}) = \mathbf{0}$, this implies $q_\varphi(\lambda) = 0$. Hence λ is a zero of q_φ as well.

$\square$

The following shows that, as far as upper triangle presentations are concerned, the minimal polynomial holds the same information as the characteristic polynomial, compare Proposition 1.3.1.

Proposition 1.5.3 *Suppose $\varphi \in \operatorname{Hom}(V, V)$ is such that its minimal polynomial q_φ splits into linear factors. Then φ is representable by an upper triangle matrix.*

Proof.   The proof is very similar to the corresponding part of the proof of Proposition 1.3.1, and we discuss only the crucial variations.

If q_φ splits into linear factors, then p_φ also has a linear factor, and hence an eigenvector $\mathbf{v}_1$. We split V as $V = U \oplus V'$, where $U = \operatorname{span}(\mathbf{v}_1)$. In order to piece together the desired representation, we want to find a suitable representation for φ over V', ignoring any components in U. We therefore consider the quotient V/U and the quotient map

$$\begin{aligned} \varphi' \colon V/U &\longrightarrow V/U \\ \mathbf{v} + U &\longmapsto \varphi(\mathbf{v}) + U. \end{aligned}$$

As the quotient $V/U = (U \oplus V')/U$ is isomorphic to V', we may therefore identify φ' with an endomorphism of V'. In order to apply the inductive hypothesis to $\varphi' \in \operatorname{Hom}(V', V')$ (or $\operatorname{Hom}(V/U, V/U)$), we need that also $q_{\varphi'}$ splits into linear factors.

This follows, as $q_{\varphi'}|q_\varphi$. One checks that for any polynomial $r \in \mathbb{F}[X]$ and $\mathbf{v} \in V$:

$$[r(\varphi')](\mathbf{v} + U) = [r(\varphi)](\mathbf{v}) + U.$$

In particular, $[q_\varphi(\varphi')](\mathbf{v}+U) = [q_\varphi(\varphi)](\mathbf{v})+U = \mathbf{0}+U$ shows that $q_\varphi(\varphi')$ is the null map. Therefore, $q_{\varphi'}|q_\varphi$, and $q_{\varphi'}$ must also split into linear factors.

The rest of the argument, i.e., the piecing together of an upper triangle representation of φ' for a suitable basis of V' with the extra basis vector $\mathbf{v}_1$, is exactly as in the proof of Proposition 1.3.1.

$\square$

Combining this with Proposition 1.3.1, we obtain the following.

Corollary 1.5.4 *For any $\varphi \in \mathrm{Hom}(V,V)$: p_φ splits into linear factors iff q_φ splits into linear factors.*

We link the minimal polynomial to *diagonalisability* of endomorphisms (matrices). This further clarifies the situation which was partially described in Proposition 1.1.15 above.

Lemma 1.5.5 *If $\varphi \in \mathrm{Hom}(V,V)$ is diagonalisable, i.e., if V has a basis consisting of eigenvectors of φ, then q_φ splits into linear factors of algebraic multiplicity 1: $q_\varphi = \prod_{i=1}^m (X - \lambda_i)$ if $\lambda_1, \ldots, \lambda_m$ are the pairwise distinct eigenvalues of φ.*

Proof.   It suffices to show that $q(\varphi) = \mathbf{0}$ for $q = \prod_{i=1}^m (X - \lambda_i)$, since the real q_φ divides any such q and must have these zeroes. (The point is that no higher multiplicities are necessary.)

Let B be a basis of eigenvectors, $\mathbf{b} \in B$ an eigenvector with eigenvalue λ_i. Writing $q = \prod_{j=1}^m (X - \lambda_j)$ as $q = (X - \lambda_i)q_i' = q_i'(X - \lambda_i)$, we see that $q(\varphi) = (q_i'(\varphi)) \circ (\varphi - \lambda_i \,\mathrm{id}_V)$. Hence $(q(\varphi))(\mathbf{b}) = (q_i'(\varphi))(\mathbf{0}) = \mathbf{0}$.

Since this argument carries through for every basis vector in B, $q(\varphi)$ is indeed the null map.

$\square$

Towards the converse of the lemma, we look at products of polynomials that are relatively prime.

Lemma 1.5.6 *Let $q_1, \ldots, q_m \in \mathbb{F}[X]$ be relatively prime. If $q = \prod_{i=1}^{m} q_i$ then*

$$\ker(q(\varphi)) = \bigoplus_{i=1}^{m} \ker(q_i(\varphi)).$$

Proof.　　We consider the case of two factors. Let $q = q_1 q_2$.

The inclusions $\ker(q_i(\varphi)) \subseteq \ker(q(\varphi))$ are easy. For instance, for $\mathbf{v} \in \ker(q_i(\varphi))$, we have that $(q(\varphi))(\mathbf{v}) = ((q_1 q_2)(\varphi))(\mathbf{v}) = ((q_2 q_1)(\varphi))(\mathbf{v}) = (q_2(\varphi) \circ q_1(\varphi))(\mathbf{v}) = (q_2(\varphi))(\mathbf{0}) = \mathbf{0}$.

So $\ker(q_1(\varphi)) + \ker(q_2(\varphi)) \subseteq \ker(q(\varphi))$.

For directness of the sum and for the opposite inclusion we use the fact that $1 = gq_1 + hq_2$ for suitable $g, h \in \mathbb{F}[X]$, since q_1 and q_2 are relatively prime (compare Lemma 1.2.19).

For $\mathbf{v} \in \ker(q_1(\varphi)) \cap \ker(q_2(\varphi))$ this implies that

$$\mathbf{v} = \mathrm{id}_V(\mathbf{v}) = (g(\varphi) \circ q_1(\varphi) + h(\varphi) \circ q_2(\varphi))(\mathbf{v}) = (g(\varphi))(\mathbf{0}) + (h(\varphi))(\mathbf{0}) = \mathbf{0}.$$

So the sum is direct.

For $\mathbf{v} \in \ker(q(\varphi))$, we may write $\mathbf{v}$ as $\mathbf{v} = (g(\varphi) \circ q_1(\varphi))(\mathbf{v}) + (h(\varphi) \circ q_2(\varphi))(\mathbf{v})$. Then $(g(\varphi) \circ q_1(\varphi))(\mathbf{v}) \in \ker(q_2(\varphi))$ and $(h(\varphi) \circ q_2(\varphi))(\mathbf{v}) \in \ker(q_1(\varphi))$ prove that $\mathbf{v} \in \ker(q_1(\varphi)) + \ker(q_2(\varphi))$.

Consider for instance the first of these: $(q_2(\varphi))((g(\varphi) \circ q_1(\varphi))(\mathbf{v})) = ((q_2 g q_1)(\varphi))(\mathbf{v}) = (gq(\varphi))(\mathbf{v}) = (g(\varphi) \circ q(\varphi))(\mathbf{v}) = (g(\varphi))(\mathbf{0}) = \mathbf{0}$.

It is now straightforward to extend the claim to any number of factors q_i by induction.

$\square$

Theorem 1.5.7 *For any $\varphi \in \mathrm{Hom}(V, V)$ with minimal polynomial q_φ, the following are equivalent*

(i) *φ is diagonalisable, i.e., V has a basis consisting of eigenvectors of φ.*

(ii) *q_φ splits into linear factors of algebraic multiplicity 1:*

$$q_\varphi = \prod_{i=1}^{m}(X - \lambda_i),$$

where $\lambda_1, \ldots, \lambda_m$ are the pairwise distinct eigenvalues of φ.

Proof.   The implication (i) $\Rightarrow$ (ii) was shown in Lemma 1.5.5.

For the converse consider now $q = q_\varphi$ splitting into distinct linear factors as in the theorem. These linear factors are irreducible and hence relatively prime. Therefore, by the last lemma,

$$V = \ker(q(\varphi)) = \bigoplus_{i=1}^{m} \ker(\varphi - \lambda_i \operatorname{id}_V) = \bigoplus_{i=1}^{m} V_{\lambda_i}.$$

Proposition 1.1.15 now tells us that φ can indeed be diagonalised.

$\square$

1.6   Jordan Normal Form

As far as nice representations for an arbitrary endomorphism $\varphi \in \operatorname{Hom}(V, V)$ of a finite dimensional $\mathbb{F}$-vector space V are concerned, we already know the following — using criteria in terms of the characteristic polynomial p_φ and the minimal polynomial q_φ:

φ may be diagonalised, i.e., V possesses a basis of eigenvectors, iff the minimal polynomial q_φ splits into distinct linear factors.

[That even p_φ splits up in this fashion is a sufficient condition for q_φ to do likewise, but not a necessary condition: φ may have eigenspaces of dimensions greater than 1.]

φ may be represented by an upper triangle matrix iff q_φ splits into (not necessarily distinct) linear factors iff p_φ splits into (not necessarily distinct) linear factors.

[This is always the case over $\mathbb{C}$, by the fundamental theorem of algebra.]

In terms of a matrix $A \in \mathbb{F}^{(n,n)}$, and its characteristic and minimal polynomials p_A and q_A, equivalently:

A is similar to a diagonal matrix iff the minimal polynomial q_A splits into distinct linear factors.

A is similar to an upper triangle matrix iff q_A (and p_A) split into (not necessarily distinct) linear factors.

Regarding the relationship between p_φ and q_φ (or p_A and q_A), we know that they have exactly the same zeroes, and hence the same linear factors, but the algebraic multiplicities may be lower in q_φ than in p_φ.

So in the case in which the characteristic polynomial p_φ splits into (not necessarily distinct) linear factors,

$$p_\varphi = (-1)^n \prod_{i=1}^{k} (X - \lambda_i)^{n_i},$$

where $\lambda_1, \ldots, \lambda_k$ are pairwise distinct, the minimal polynomial q_φ is of the form

$$q_\varphi = \prod_{i=1}^{k} (X - \lambda_i)^{n_i'},$$

with multiplicities $1 \leqslant n_i' \leqslant n_i$.

Only if the n_i' are all 1, may φ be diagonalised.

The question to be addressed in this section is the following. What more, apart from merely upper triangle form can be guaranteed if q_φ does split into linear factors, but if some of these occur with multiplicity greater than 1.

We shall proceed in stages towards the best possible representation of an endomorphism whose characteristic polynomial splits into linear factors (as is always the case over $\mathbb{C}$-vector spaces). These stages correspond to the identification of suitable invariant subspaces, into which V decomposes as a direct sum. We analyse φ further in restriction to these invariant subspaces, repeating the process as necessary. Putting the pieces together in the end, we obtain a representation of φ in Jordan normal form, which is as close to diagonal form as possible in this general situation.

Convention: For this entire section, again, V is a fixed finite-dimensional $\mathbb{F}$-vector space, $\dim(V) > 0$, and $\varphi \in \mathrm{Hom}(V, V)$ an endomorphism.

1.6.1   Block decomposition, part 1

Let $p = p_\varphi = (-1)^n \prod_{i=1}^{k} (X - \lambda_i)^{n_i} = (-1)^n \prod_{i=1}^{k} p_i$ with pairwise distinct λ_i, such that the polynomials

$$p_i = (X - \lambda_i)^{n_i}$$

are relatively prime. Put

$$V^{(i)} := \ker(p_i(\varphi)) = \ker((\varphi - \lambda_i \, \mathrm{id}_V)^{n_i}) \subseteq V.$$

Lemma 1.6.1 *In the above situation:*

(i) $V = \ker(p(\varphi)) = \bigoplus_{i=1}^{k} V^{(i)} = \bigoplus_{i=1}^{k} \ker(p_i(\varphi))$.

(ii) *each $V^{(i)} = \ker(p_i(\varphi))$ is an invariant subspace for φ.*

(iii) *for suitable bases $B_i = (\mathbf{b}_1^{(i)}, \ldots, \mathbf{b}_{n_i}^{(i)})$ of the $V^{(i)}$, we obtain a basis $B = (B_1, \ldots, B_k)$ of V such that w.r.t. to this basis, φ is represented by a block matrix*

$$A = \begin{pmatrix} A_1 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & A_2 & & & \mathbf{0} \\ \vdots & & & & \vdots \\ \mathbf{0} & \mathbf{0} & & & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & A_k \end{pmatrix},$$

where each $A_i \in \mathbb{F}^{(n_i, n_i)}$ is an upper triangle matrix with entries λ_i on the diagonal,

$$A_i = \begin{pmatrix} \lambda_i & * & \cdots & * & * \\ 0 & \lambda_i & & & * \\ \vdots & & \ddots & & \vdots \\ 0 & & & \lambda_i & * \\ 0 & 0 & \cdots & 0 & \lambda_i \end{pmatrix}.$$

Proof.   Assertion (i) follows from Lemma 1.5.6 above, as the p_i are relatively prime.

For (ii), we need to show that $\varphi(\mathbf{v}) \in V^{(i)}$ for $\mathbf{v} \in V^{(i)}$.

Let $\mathbf{v} \in V^{(i)} = \ker(p_i(\varphi))$, i.e., $(p_i(\varphi))(\mathbf{v}) = \mathbf{0}$. Then $[p_i(\varphi)](\varphi(v)) = [p_i(\varphi) \circ \varphi](\mathbf{v}) = [\varphi \circ p_i(\varphi)](\mathbf{v}) = \varphi((p_i(\varphi))(\mathbf{v})) = \varphi(\mathbf{0}) = \mathbf{0}$, so indeed $\varphi(\mathbf{v}) \in V^{(i)}$.

For (iii), consider the map $\varphi_i \in \mathrm{Hom}(V^{(i)}, V^{(i)})$ which is the restriction of φ to the invariant subspace $V^{(i)}$.

Note first that φ_i cannot have an eigenvalue λ_j for $j \neq i$. Otherwise, φ_i and hence φ would have an eigenvector $\mathbf{v} \in V^{(i)}$ with $\varphi(\mathbf{v}) = \lambda_j \mathbf{v}$; but then $\mathbf{v}$ is an eigenvector with eigenvalue $\lambda_j - \lambda_i \neq 0$ of $(\varphi - \lambda_i \,\mathrm{id}_V)$, and hence $(p_i(\varphi))(\mathbf{v}) = (\varphi - \lambda_i \,\mathrm{id}_V)^{n_i}(\mathbf{v}) = (\lambda_i - \lambda_j)^{n_i} \mathbf{v} \neq \mathbf{0}$, contradicting $\mathbf{v} \in V^{(i)}$. So if p_{φ_i} splits into linear factors, we must have $p_{\varphi_i} = (-1)^{m_i}(X - \lambda_i)^{m_i}$, where m_i is the dimension of $V^{(i)}$.

That p_{φ_i} does split into linear factors, follows as $p(\varphi) = \prod_i p_{\varphi_i}$. This is a consequence of the observation that the $V^{(i)}$ split V into a direct sum

of invariant subspaces, see (i) and (ii). Therefore, for any choice of bases B_i for $V^{(i)}$, with corresponding representations A_i for φ_i, φ is represented w.r.t. the combined basis $B = (B_1, \ldots, B_k)$ by a block diagonal matrix of the form of A in (iii) (see Lemma 1.1.12 above; but here we do not know anything about the A_i yet). This implies that $p = p_\varphi = p_A = |A - XE_n| = \prod_i |A_i - XE_{m_i}| = \prod_i p_{\varphi_i}$, whence also the p_{φ_i} split into linear factors. It then follows that $m_i = n_i = \dim(V^{(i)})$.

Let finally then the basis B_i of $V^{(i)}$ be chosen according to Proposition 1.3.1. This gives the desired form for the A_i.

$\square$

Corollary 1.6.2 *If $p_\varphi = (-1)^n \prod_{i=1}^k (X - \lambda_i)^{n_i}$ for pairwise distinct λ_i, then $V^{(i)} = \ker\big((\varphi - \lambda_i \operatorname{id}_V)^{n_i}\big)$ is an invariant subspace of dimension n_i, and V is the direct sum of these $V^{(i)}$.*

1.6.2 Block decomposition, part 2

We turn to what emerged as the situation for φ_i above, i.e., for φ in restriction to one of the invariant subspaces $V^{(i)}$.

We therefore now assume that $p = p_\varphi = (-1)^n (X - \lambda)^n$ has a single linear factor corresponding to the eigenvalue λ, of multiplicity $n = \dim(V)$.

Note that the minimal polynomial q_φ must be of form $q_\varphi = (X - \lambda)^m$ for some $1 \leqslant m \leqslant n$. We know that φ is diagonalisable iff $m = 1$. So, what happens if $m > 1$? We study the situation with the help of the endomorphism

$$\psi = \varphi - \lambda \operatorname{id}_V.$$

Clearly $p_\varphi(\varphi) = (-1)^n \psi^n$ and hence $V = \ker(\psi^n)$ by the Cayley–Hamilton Theorem. Moreover, the degree m of the minimal polynomial q_φ is the minimal number m such that $V = \ker(\psi^m)$.

We therefore consider the subspaces

$$U_j := \ker(\psi^j) \subseteq V \quad \text{for } j = 0, \ldots, n.$$

Lemma 1.6.3 *For the U_j as just defined:*

$$\{\mathbf{0}\} = U_0 \subsetneq U_1 \subsetneq \cdots \subsetneq U_m = U_{m+1} = \ldots = U_n = V,$$

where m is the degree of the minimal polynomial $q_\varphi = (X - \lambda)^m$.

$$U_m = V$$

$$U_{m-1}$$

$$\vdots$$

$$U_2$$

$$U_1 = V_\lambda$$

Proof. $U_0 = \{\mathbf{0}\}$ as $\psi^0 = \mathrm{id}_V$.

$U_0 \subsetneqq U_1$ as $U_1 = \ker(\psi) = \ker(\varphi - \lambda\,\mathrm{id}_V) = V_\lambda$ is a non-trivial eigenspace because λ is an eigenvalue of φ.

Clearly $U_j \subseteq U_{j+1}$ for all j. Further, $U_j = U_{j+1} \;\Rightarrow\; U_{j+1} = U_{j+2}$, and hence the sequence of the U_j becomes constant as soon as it does not strictly increase once. This is because

$$\mathbf{v} \in U_{j+2} \setminus U_{j+1} \quad \text{iff} \quad \psi(\mathbf{v}) \in U_{j+1} \setminus U_j.$$

In particular, for the smallest m with $U_m = U_{m+1}$ it follows that also $U_m = V$. It remains to argue that this smallest m is the degree of q_φ. By the definition of the minimal polynomial q_φ, $\ker(q_\varphi(\varphi)) = V$ and $q_\varphi | q$ for any other q such that $\ker(q(\varphi)) = V$. Hence, as $q_\varphi = \psi^\ell$ for some ℓ, the first condition tells us that $\ell \geqslant m$, while the second implies that $\ell \leqslant m$.

$\square$

We can think now of V as stratified w.r.t. the chain of subspaces U_j. With a vector $\mathbf{v} \in V$ we associate its *height* $h(\mathbf{v})$ w.r.t. this stratification according to

$$h(\mathbf{v}) := j \quad \text{iff } \mathbf{v} \in U_j \setminus U_{j-1}.$$

The range of these heights is between 0 ($\mathbf{v} = \mathbf{0}$ only) and the degree m of q_φ. Note that for $h(\mathbf{v}) \geqslant 1$:

$$h(\psi(\mathbf{v})) = h(\mathbf{v}) - 1.$$

We first look at a vector $\mathbf{v}$ of height $\ell > 0$, and at its iterated ψ-images, $\psi(\mathbf{v}), \psi(\psi(\mathbf{v})), \ldots$, which create a sequence of vectors of decreasing heights $\ell, \ell - 1, \ldots, 0$.

Lemma 1.6.4 *Let $\mathbf{v} \in U_\ell \backslash U_{\ell-1}$, $\ell = h(\mathbf{v}) \geqslant 1$. Then the sequence of vectors*

$$B = (\psi^{\ell-1}(\mathbf{v}), \psi^{\ell-2}(\mathbf{v}), \ldots, \psi(\mathbf{v}), \mathbf{v}),$$

consisting of vectors of heights $1, 2, \ldots, \ell$, is linearly independent, and hence a labelled basis of the ℓ-dimensional subspace

$$[\![\mathbf{v}]\!] := \mathrm{span}(\psi^{\ell-1}(\mathbf{v}), \psi^{\ell-2}(\mathbf{v}), \ldots, \psi(\mathbf{v}), \mathbf{v}) \subseteq V.$$

$[\![\mathbf{v}]\!]$ is an invariant subspace of φ, and the restriction of φ to $[\![\mathbf{v}]\!]$ is represented w.r.t. B by the matrix

$$A_{\varphi,[\![\mathbf{v}]\!]} = \begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 \\ 0 & \lambda & 1 & & 0 \\ \vdots & & \ddots & \ddots & 0 \\ 0 & & & \lambda & 1 \\ 0 & \cdots & & 0 & \lambda \end{pmatrix}$$

with entries λ across the diagonal, entries 1 right above the diagonal, and 0 everywhere else.

Proof.   For linear independence, we argue by induction on $\ell = h(\mathbf{v})$.

The claim is obviously true for $\ell = 1$.

Suppose $h(\mathbf{v}) = \ell + 1$, and the claim is true for ℓ. In particular, since $h(\psi(\mathbf{v})) = \ell$, $(\psi(\mathbf{v}), \psi^2(\mathbf{v}), \ldots, \psi^\ell(\mathbf{v}))$ is linearly independent and forms a basis of $[\![\psi(\mathbf{v})]\!]$. $\mathbf{v} \notin [\![\psi(\mathbf{v})]\!]$, because $[\![\psi(\mathbf{v})]\!] \subseteq U_\ell = \ker(\psi^\ell)$ while $\mathbf{v} \in U_{\ell+1} \backslash U_\ell$. Therefore, $(\mathbf{v}, \psi(\mathbf{v}), \psi^2(\mathbf{v}), \ldots, \psi^\ell(\mathbf{v}))$ is also linearly independent and forms a basis of $[\![\mathbf{v}]\!]$.

For invariance of $[\![\mathbf{v}]\!]$ under φ, it suffices to show for each one of the basis vectors $\psi^i(\mathbf{v})$ of $[\![\mathbf{v}]\!]$, $0 \leqslant i < \ell$, that $\varphi(\psi^i(\mathbf{v})) \in [\![\mathbf{v}]\!]$. As $\varphi(\psi^i(\mathbf{v})) = \big(\varphi \circ (\varphi - \lambda\,\mathrm{id}_V)^i\big)(\mathbf{v})$, and since $\varphi \circ (\varphi - \lambda\,\mathrm{id}_V)^i = (\varphi - \lambda\,\mathrm{id}_V)^{i+1} + \lambda(\varphi - \lambda\,\mathrm{id}_V)^i$, we have

$$\varphi(\psi^i(\mathbf{v})) = \psi^{i+1}(\mathbf{v}) + \lambda\psi^i(\mathbf{v}) \in [\![\mathbf{v}]\!].$$

For the claim concerning the representational matrix $A_{\varphi,[\![\mathbf{v}]\!]}$, note that the i-th basis vector in $[\![\mathbf{v}]\!]$ is $\mathbf{b}_i = \psi^{\ell-i}(\mathbf{v})$, $i = 1, \ldots, \ell = h(\mathbf{v})$. The last equation translates into

$$\varphi(\mathbf{b}_i) = \varphi(\psi^{\ell-i}(\mathbf{v})) = \psi^{\ell-i+1}(\mathbf{v}) + \lambda\psi^{\ell-i}(\mathbf{v}) = \begin{cases} \mathbf{b}_{i-1} + \lambda\mathbf{b}_i & \text{for } i > 1 \\ \lambda\mathbf{b}_i & \text{for } i = 1 \end{cases}$$

which shows that $A_{\varphi,[\![\mathbf{v}]\!]}$ is as claimed.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad\square$
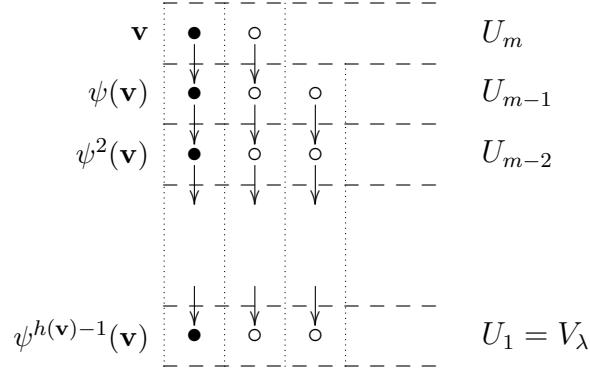
We now want to split V into a direct sum of invariant subspaces of the form $[\![\mathbf{v}]\!]$, in order to obtain a block diagonal matrix whose blocks are of the form of $A_{\varphi,[\![\mathbf{v}]\!]}$ above.

We proceed iteratively, extracting one suitable $[\![\mathbf{v}]\!]$ after another.

Choose $\mathbf{v} \in V$ of maximal height. We write $[\![\mathbf{v}]\!]_i$ for the subspace $[\![\mathbf{v}]\!] \cap U_i$, with basis $(\psi^{h(\mathbf{v})-i}(\mathbf{v}), \ldots, \psi^{h(\mathbf{v})-1}(\mathbf{v}))$. In particular, $[\![\mathbf{v}]\!]_1 = [\![\mathbf{v}]\!] \cap U_1$ is spanned by $\psi^{h(\mathbf{v})-1}(\mathbf{v})$. If $V \neq [\![\mathbf{v}]\!]$, let U_1' be any complement of $[\![\mathbf{v}]\!]_1$ in U_1:

$$U_1 = [\![\mathbf{v}]\!]_1 \oplus U_1'.$$

Starting with U_1', now inductively choose subspaces $U_i' \subseteq U_i$ such that

(i) $U_i' \subseteq U_{i+1}'$.

(ii) $U_i = [\![\mathbf{v}]\!]_i \oplus U_i'$.

(iii) $\psi(U_{i+1}') \subseteq U_i'$.

Then $V' := U_m'$ will be such that

$$V = [\![\mathbf{v}]\!] \oplus V'$$

is a decomposition of V into a direct sum of subspaces that are invariant under φ. Here $V = [\![\mathbf{v}]\!] \oplus V'$ is clear from (ii), as $V = U_m$ and $[\![\mathbf{v}]\!] = [\![\mathbf{v}]\!]_m$. Invariance of V' under φ follows from (iii) and (i), as $\varphi = \psi + \lambda\mathrm{id}$.

It remains to show that, for $i = 1, \ldots, m-1$, U_{i+1}' can be chosen as required in relation to the previously chosen U_i'. Let U_i' be as required, B_i' a basis of U_i'. Then the basis $(\psi^{h(\mathbf{v})-i}(\mathbf{v}), \ldots, \psi^{h(\mathbf{v})-1}(\mathbf{v}))$ of $[\![\mathbf{v}]\!]_i$ extends B_i' to a basis B_i of $U_i = [\![\mathbf{v}]\!]_i \oplus U_i'$ by (ii). As $\psi^{h(\mathbf{v})-(i+1)}(\mathbf{v})$ has height $i+1$, it is linearly independent from B_i. We now want to extend the linearly independent $B_i \cup$

$\{\psi^{h(\mathbf{v})-(i+1)}(\mathbf{v})\}$ to a basis B_{i+1} of U_{i+1} using new basis vectors $\mathbf{b}$ for which $\psi(\mathbf{b}) \in U_i'$. Then $B_{i+1}' := B_{i+1} \setminus \{\psi^{h(\mathbf{v})-(i+1)}(\mathbf{v}), \psi^{h(\mathbf{v})-i}(\mathbf{v}), \ldots, \psi^{h(\mathbf{v})-1}(\mathbf{v})\}$ can serve as a basis for the desired U_{i+1}'. Such an extension is possible due to the following claim.

Claim 1.6.5 $U_{i+1} \subseteq [\![\mathbf{v}]\!]_{i+1} + \{\mathbf{w} \in U_{i+1}: \psi(\mathbf{w}) \in U_i'\}$.

Proof.   Consider $\mathbf{w} \in U_{i+1} \setminus U_i$. Since $h(\mathbf{w}) = i+1$, $h(\psi(\mathbf{w})) = i$ and, by (ii) for U_i', $\psi(\mathbf{w}) = \mathbf{v}' + \mathbf{u}'$ for suitable $\mathbf{v}' \in [\![\mathbf{v}]\!]_i$ and $\mathbf{u}' \in U_i'$. Let $\mathbf{v}' = \sum_{1 \leqslant j \leqslant i} \lambda_j \psi^{h(\mathbf{v})-j}(\mathbf{v})$. Putting $\mathbf{v}'' := \sum_{1 \leqslant j \leqslant i} \lambda_j \psi^{h(\mathbf{v})-j-1}(\mathbf{v})$, we have $\mathbf{v}'' \in [\![\mathbf{v}]\!]_{i+1}$ and $\psi(\mathbf{v}'') = \mathbf{v}'$. Now $\mathbf{w} = \mathbf{v}'' + (\mathbf{w} - \mathbf{v}'')$ and $\psi(\mathbf{w} - \mathbf{v}'') = \psi(\mathbf{w}) - \psi(\mathbf{v}'') = \psi(\mathbf{w}) - \mathbf{v}' = \mathbf{u}' \in U_i'$ shows that $\mathbf{w} \in [\![\mathbf{v}]\!]_{i+1} + \{\mathbf{w} \in U_{i+1}: \psi(\mathbf{w}) \in U_i'\}$.

□

The process of eliminating one $[\![\mathbf{v}]\!]$ (of maximal remaining height in V') after the other can be iterated as long as there remains a non-trivial complement V'. We thus obtain the following.

Lemma 1.6.6 *If $p_\varphi = (-1)^n(X - \lambda)^n$ and $q_\varphi = (X - \lambda)^m$, then V can be decomposed into a direct sum of invariant subspaces of the form $[\![\mathbf{v}]\!] = \mathrm{span}(\mathbf{v}, \psi(\mathbf{v}), \ldots, \psi^{h(\mathbf{v})-1}(\mathbf{v}))$ for suitable $\mathbf{v}$, such that the dimensions of these subspaces are $m = \ell_1 \geqslant \cdots \geqslant \ell_s \geqslant 1$, with $\sum_j \ell_j = n$. Here s is the dimension of the eigenspace V_λ.*

W.r.t. a basis obtained from the bases $(\psi^{h(\mathbf{v})-1}(\mathbf{v}), \ldots, \mathbf{v})$ of these subspaces $[\![\mathbf{v}]\!]$, as described in Lemma 1.6.4, φ is represented by a block diagonal matrix with blocks of the form $A_{\varphi, [\![\mathbf{v}]\!]}$ described there.

Proof.   Suitable $[\![\mathbf{v}]\!]$ are successively obtained according to the above. Note that the first $\mathbf{v}$, of maximal height m in V, gives rise to an invariant subspace $[\![\mathbf{v}]\!]$ of dimension m. All consecutive iterations of the process produce contributions of weakly decreasing heights and dimensions, all greater than 0.

If $V = \bigoplus_{j=1}^s [\![\mathbf{v}_j]\!]$, then $\dim(V) = n = \sum_j \dim([\![\mathbf{v}_j]\!])$. As any $[\![\mathbf{v}_j]\!] \cap U_1 = [\![\mathbf{v}_j]\!] \cap V_\lambda$ has dimension 1, there are $s = \dim(V_\lambda)$ many successive contributions; each involving (the span of) just one eigenvector $\psi^{h(\mathbf{v}_j)-1}(\mathbf{v}_j)$.

The desired representation of φ follows with Lemma 1.1.12 (block decomposition in direct sums of invariant subspaces) and Lemma 1.6.4 (desired shape for the blocks).

□

Example 1.6.7 For a backward analysis of how we identify the crucial invariant subspaces and parameters look at an endomorphism φ of $\mathbb{R}^7$ that is already of the desired form w.r.t. the standard basis.

$$A_\varphi = \begin{pmatrix} 2 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 \end{pmatrix}.$$

For the characteristic polynomial we find $p_\varphi = -(X-2)^7$, while the minimal polynomial is $q_\varphi = (X-2)^3$. [Recall that these would be the same, for any representation of φ, no matter which basis.]

$\lambda = 2$ is the only eigenvalue, with eigenspace

$$V_\lambda = \ker(\psi) = \text{span}(\mathbf{e}_1, \mathbf{e}_4, \mathbf{e}_6, \mathbf{e}_7).$$

$\dim(V_\lambda) = 4$. i.e., the geometric multiplicity of the eigenvalue $\lambda = 2$ is 4 (while the algebraic multiplicity of the root $\lambda = 2$ in p_φ is 7). We expect four blocks; and as q_φ is of degree 3, at least one of these blocks will have dimension 3, others possibly less.

The auxiliary map $\psi = \varphi - \lambda\,\text{id}_V$ has the matrix representation

$$A_\psi = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Looking at the stratification of $V = \mathbb{R}^7$ w.r.t. to the $U_j = \ker(\psi^j)$, we obtain

$$\begin{aligned} U_1 = V_\lambda &= \text{span}(\mathbf{e}_1, \mathbf{e}_4, \mathbf{e}_6, \mathbf{e}_7), \\ U_2 &= \text{span}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_4, \mathbf{e}_5, \mathbf{e}_6, \mathbf{e}_7), \\ U_3 &= \text{span}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4, \mathbf{e}_5, \mathbf{e}_6, \mathbf{e}_7) = \mathbb{R}^7. \end{aligned}$$

So $\mathbf{e}_3 \in U_3 \setminus U_2$, is a vector of maximal height 3, which we may choose for $\mathbf{v}_1$. Note that $\psi(\mathbf{e}_3) = \mathbf{e}_2$ and $\psi(\mathbf{e}_2) = \mathbf{e}_1 \in U_1$. Therefore, the corresponding

invariant subspace is $[\![\mathbf{e}_3]\!] = \mathrm{span}(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ — giving rise to the first block,

$$A_{\varphi,[\![\mathbf{v}_1]\!]} = \begin{pmatrix} 2 & 1 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 2 \end{pmatrix}.$$

$[\![\mathbf{e}_3]\!] \cap U_1 = \mathrm{span}(\mathbf{e}_1)$. As a complement for $[\![\mathbf{e}_3]\!] \cap U_1$ we may choose $U_1' = \mathrm{span}(\mathbf{e}_4, \mathbf{e}_6, \mathbf{e}_7)$. The invariant subspace V' that is a complement of $[\![\mathbf{e}_3]\!]$ in V, is $V' = \mathrm{span}(\mathbf{e}_4, \mathbf{e}_5, \mathbf{e}_6, \mathbf{e}_7)$.

The restriction φ' of φ to V' is represented w.r.t. the standard basis $(\mathbf{e}_4, \mathbf{e}_5, \mathbf{e}_6, \mathbf{e}_7)$ for V', by

$$A_{\varphi'} = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 \end{pmatrix}.$$

Repeating the process for φ' in V', one can choose $\mathbf{v}_2 = \mathbf{e}_5$ as a vector of maximal height in V'; its height being 2: $\psi(\mathbf{e}_5) = \mathbf{e}_4 \in U_1$. [One can check that $p_{\varphi'} = (X - 2)^4$ and $q_{\varphi'} = (X - 2)^2$.]

A corresponding invariant subspace is $[\![\mathbf{e}_5]\!] = \mathrm{span}(\mathbf{e}_4, \mathbf{e}_5)$ — giving rise to the second block,

$$A_{\varphi,[\![\mathbf{v}_2]\!]} = \begin{pmatrix} 2 & 1 \\ 0 & 2 \end{pmatrix}.$$

A complement for $[\![\mathbf{v}_2]\!] \cap U_1 = \mathrm{span}(\mathbf{e}_4)$ in U_1' is $U_1'' = \mathrm{span}(\mathbf{e}_6, \mathbf{e}_7)$. Now $V'' = U_1''$ as there are no more vectors of height greater than 1. [The restriction φ'' to $V'' = \mathrm{span}(\mathbf{e}_6, \mathbf{e}_7)$ has $p_{\varphi''} = (X - 2)^2$ and $q_{\varphi''} = (x - 2)$.] Hence φ'' is diagonal, and the remaining two blocks are both of of size 1.

1.6.3   Jordan normal form

Retracing our steps to the beginning, and combining the results for the first and second block decompositions, we obtain the full statement of the Jordan normal form as follows.

Theorem 1.6.8 *Let $\varphi \in \mathrm{Hom}(V, V)$ such that p_φ splits into linear factors,*

$$p_\varphi = (-1)^n \prod_{i=1}^{k} (X - \lambda_i)^{n_i}.$$

Then $q_\varphi = \prod_{i=1}^{k}(X - \lambda_i)^{m_i}$ for some $1 \leqslant m_i \leqslant n_i$.

For a suitable basis of V, φ is represented by a matrix of (two-level) block diagonal form with

(i) on the first level: k blocks, of sizes $n_1, \ldots, n_k$, with entries λ_i on the diagonal and entries from $\{0, 1\}$ right above the diagonal.

(ii) the i-th block matrix is itself of block diagonal form, with blocks of sizes $\ell_j^{(i)}$, $\sum_{j=1}^{s_i} \ell_j^{(i)} = n_i$; $m_i = \ell_1^{(i)} \geqslant \ldots \geqslant \ell_{s_i}^{(i)} \geqslant 1$; each block has entries λ_i on the diagonal and entries 1 right above. Here $s_i = \dim(V_{\lambda_i})$ is the dimension of the corresponding eigenspace.

We consider in the following exercise an example closely related to the motivating Examples 1.0.1 and 1.0.2, to show that also Jordan normal form can have advantages close to those of diagonal form (which is important where only Jordan normal form can be achieved). These examples have to do with powers of a matrix and with the evaluation of the exponential function on a matrix — which has applications to differential equations as seen before.

Exercise 1.6.1 For simplicity consider a Jordan normal form matrix with a single block, of the form

$$
A = \begin{pmatrix}
\lambda & 1 & 0 & \cdots & 0 \\
0 & \lambda & 1 & & 0 \\
\vdots & & \ddots & \ddots & 0 \\
0 & & & \lambda & 1 \\
0 & \cdots & & 0 & \lambda
\end{pmatrix} \in \mathbb{R}^{(n,n)}.
$$

Writing A as $A = \lambda E_n + N$, we may evaluate powers A^i in terms of the binomial expansion in powers of $(\lambda E_n)^j = \lambda^j E_n$ and N^k. $N^0 = E_n$, $N^1 = N$. Show that

(i) for $1 \leqslant k \leqslant n-1$, N^k is the matrix with zero entries everywhere apart from the k-th tier above the diagonal which has entries 1.

(ii) $N^k = \mathbf{0}$ for exponents $k \geqslant n$.

In other word, each increase in the exponent shifts the line of ones one step towards the upper right-hand corner.

Correspondingly, for the exponential function

$$
e^A = \sum_{j=0}^{\infty} \frac{A^j}{j!},
$$

one finds that

$$
e^A = e^{\lambda E_n + N} = e^{\lambda E_n} e^N = e^\lambda E_n e^N = e^\lambda e^N = e^\lambda \sum_{j=0}^{\infty} \frac{N^j}{j!} = e^\lambda \sum_{j=0}^{n-1} \frac{N^j}{j!}
$$

reduces to a finite sum, which is easily evaluated. [The second equality relies on the fact that λE_n and N commute.]

(iii) Determine the coefficients of the matrix e^N.

(iv) Similarly expand the matrix e^{tN} in powers of $t \in \mathbb{R}$, and relate this to the function $t \mapsto e^{tA}$ using the equalities $e^{tA} = e^{t\lambda E_n + tN} = e^{\lambda t} e^{tN}$.

Chapter 2

Euclidean and Unitary Spaces

In this chapter we study vector spaces (over $\mathbb{R}$ and $\mathbb{C}$) with additional *metric structure* that allows us to speak of *lengths* [Länge, Betrag] of vectors and *angles* [Winkel] between vectors. The additional algebraic structure, over and above the underlying vector space structure, is that of a (real or complex) *scalar product* [Skalarprodukt] that associates a scalar with every pair of vectors.

It is clear that the vector space structure on its own does not support any geometric notion of length: after all, re-scalings of the form $\varphi\colon \mathbf{v} \mapsto \lambda\mathbf{v}$, for arbitrary $\lambda \in \mathbb{F} \setminus \{0\}$ are vector space automorphisms. Homomorphisms that preserve not just the vector space structure, but also the additional metric structure, will correspondingly be length-preserving or *isometric* homomorphisms, or *isometries*. We shall study these structure preserving maps in connection with the structure of real and complex vectors spaces equipped with a scalar product.

An $\mathbb{R}$-vector space with a real-valued scalar product is called a *euclidean space*, while a $\mathbb{C}$-vector space with a $\mathbb{C}$-valued scalar product is a *unitary space*. The main topics of this chapter are euclidean and unitary vector spaces; corresponding bases with special metric properties (orthonormal bases); lengths, angles and orthogonality in such spaces; isometries of euclidean and unitary spaces; corresponding matrix groups; the representations of vector space endomorphisms w.r.t. orthonormal bases; and diagonalisation of natural classes of homomorphisms and matrices that have special properties in relation to the given scalar product.

Column vectors and row vectors   We now have occasion to fix a convention with regard to row and column vectors. We make the following stipulation for this whole chapter: all vectors in the standard spaces $\mathbb{F}^n$ (here: $\mathbb{F} = \mathbb{R}$ or $\mathbb{F} = \mathbb{C}$) are written as column vectors. In other words, we identify $\mathbf{v} \in \mathbb{F}^n$ with an element of $\mathbb{F}^{(n,1)}$ (slim matrices: n rows, one column). The row vector corresponding to $\mathbf{v}$ is explicitly obtained as the *transpose* of $\mathbf{v}$:

$$\mathbf{v} = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix} \in \mathbb{F}^n = \mathbb{F}^{(n,1)} \qquad \text{versus} \qquad \mathbf{v}^t = (v_1, \ldots, v_n) \in \mathbb{F}^{(1,n)}.$$

Correspondingly, w.r.t. to some basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of the n-dimensional $\mathbb{F}$-vector space V, we think of the coefficients of a vector $\mathbf{v} = \sum_{i=1}^n \lambda_i \mathbf{b}_i$ as forming the column vector

$$\llbracket \mathbf{v} \rrbracket_B = \begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_n \end{pmatrix} \in \mathbb{F}^{(n,1)} \qquad \text{and write} \qquad \llbracket \mathbf{v} \rrbracket_B^t \text{ for } (\lambda_1, \ldots, \lambda_n).$$

2.1   Euclidean and unitary vector spaces

2.1.1   The standard scalar products in $\mathbb{R}^n$ and $\mathbb{C}^n$

$\mathbb{R}^n$ as a euclidean space

Definition 2.1.1 The *standard scalar product* [Standard Skalarprodukt] in $\mathbb{R}^n$ is the map

$$(\mathbf{v}, \mathbf{w}) \longmapsto \langle \mathbf{v}, \mathbf{w} \rangle = \sum_{i=1}^{n} v_i w_i = (v_1, \ldots, v_n) \begin{pmatrix} w_1 \\ \vdots \\ w_n \end{pmatrix} = \mathbf{v}^t \cdot \mathbf{w},$$

where the last presentation is in terms of matrix multiplication of $\mathbf{v}^t \in \mathbb{R}^{(1,n)}$ with $\mathbf{w} \in \mathbb{R}^{(n,1)}$ resulting in a matrix in $\mathbb{R}^{(1,1)} = \mathbb{R}$.

The $\mathbb{R}$-vector space $\mathbb{R}^n$ with this scalar product is referred to as the (standard) n-dimensional euclidean vector space $\mathbb{R}^n$.

The following definition collects some key properties of the standard scalar product in $\mathbb{R}^n$, viewed as a binary function

$$\begin{array}{rcl} V \times V & \longrightarrow & \mathbb{R} \\ (\mathbf{v}, \mathbf{w}) & \longmapsto & \langle \mathbf{v}, \mathbf{w} \rangle, \end{array}$$

over an $\mathbb{R}$-vector space V. These properties make sense as properties of arbitrary binary functions $\sigma \colon V \times V \to \mathbb{R}$; but we express them explicitly in the format of $\sigma(\mathbf{v}, \mathbf{w}) = \langle \mathbf{v}, \mathbf{w} \rangle$.

Definition 2.1.2 A map $\langle ., . \rangle \colon V \times V \to \mathbb{R}$ is
  (a) *bilinear* [bilinear] if it is linear in both arguments. This means that for all $\mathbf{v}, \mathbf{v}_1, \mathbf{v}_2, \mathbf{w}, \mathbf{w}_1, \mathbf{w}_2 \in V$ and $\lambda_1, \lambda_2 \in \mathbb{R}$:

$$\begin{array}{rcl} \langle \lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2, \mathbf{w} \rangle & = & \lambda_1 \langle \mathbf{v}_1, \mathbf{w} \rangle + \lambda_2 \langle \mathbf{v}_2, \mathbf{w} \rangle \\ \text{and} \ \ \langle \mathbf{v}, \lambda_1 \mathbf{w}_1 + \lambda_2 \mathbf{w}_2 \rangle & = & \lambda_1 \langle \mathbf{v}, \mathbf{w}_1 \rangle + \lambda_2 \langle \mathbf{v}, \mathbf{w}_2 \rangle. \end{array}$$

  (b) *symmetric* [symmetrisch] if for all $\mathbf{v}, \mathbf{w} \in V$:

$$\langle \mathbf{v}, \mathbf{w} \rangle = \langle \mathbf{w}, \mathbf{v} \rangle.$$

  (c) *positive definite* [positiv definit], if for all $\mathbf{v} \in V$,

$$\langle \mathbf{v}, \mathbf{v} \rangle \geqslant 0, \ \text{and} \ \langle \mathbf{v}, \mathbf{v} \rangle = 0 \ \text{iff} \ \mathbf{v} = \mathbf{0}.$$

Any map $\langle .,.\rangle \colon V \times V \to \mathbb{R}$ satisfying conditions (a), (b), (c) is a *scalar product* [Skalarprodukt] on V. One usually denotes a scalar product $\langle .,.\rangle$ (as for the standard scalar product in $\mathbb{R}^n$ above).

An $\mathbb{R}$-vector space with a scalar product $\langle .,.\rangle$, $(V, \langle .,.\rangle)$ is called a *euclidean vector space* [Euklidischer Vektorraum].

Note that bilinearity could equivalently be phrased in terms of compatibility with scalar multiplication and vector addition. This then implies compatibility with arbitrary linear combinations, as e.g. in $\langle \sum_{i=1}^k \lambda_i \mathbf{v}_i, \mathbf{w} \rangle = \sum_{i=1}^k \lambda_i \langle \mathbf{v}_i, \mathbf{w} \rangle$. Note also that linearity in the second argument in (a) follows from linearity in the first by symmetry (b).

We shall see later that any n-dimensional $\mathbb{R}$-vector space V with a scalar product is isomorphic (as a euclidean vector space) to $\mathbb{R}^n$ with the standard scalar product.

We now define lengths and angles in $\mathbb{R}^n$ with its standard scalar product. The definitions generalise to arbitrary euclidean vector spaces. In $\mathbb{R}^n$ and in particular for $\mathbb{R}^2$ and $\mathbb{R}^3$ these are the familiar definitions from elementary euclidean geometry.

Definition 2.1.3 In $\mathbb{R}^n$ with its standard scalar product $\langle \mathbf{v}, \mathbf{w} \rangle = \mathbf{v}^t \cdot \mathbf{w}$, define the *length* or *norm* [Länge, Betrag, Norm] of vector $\mathbf{v} \in \mathbb{R}^n$ as

$$\|\mathbf{v}\| := \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle} = \sqrt{\sum_{i=1}^n v_i^2}.$$

Vectors of length 1, $\|\mathbf{v}\| = 1$, are called *unit vectors* [Einheitsvektoren].

The *normalisation* of a vector $\mathbf{v} \neq \mathbf{0}$ is the unit vector $\hat{\mathbf{v}} := \frac{1}{\|\mathbf{v}\|}\mathbf{v}$.

For angles, we also appeal to the familiar connection between the value of the standard scalar product (in $\mathbb{R}^2$ say) and the angles and lengths of the vectors involved. For the standard scalar product and for an angle $\angle(\mathbf{v}, \mathbf{w})$ between $\mathbf{v}, \mathbf{w} \neq \mathbf{0}$:

$$\langle \mathbf{v}, \mathbf{w} \rangle = \|\mathbf{v}\| \cdot \|\mathbf{w}\| \cdot \cos(\angle(\mathbf{v}, \mathbf{w})).$$

Definition 2.1.4 The (pointed) *angle* between vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ is defined by

$$\angle(\mathbf{v}, \mathbf{w}) := \arccos\left(\frac{\langle \mathbf{v}, \mathbf{w} \rangle}{\|\mathbf{v}\|\|\mathbf{w}\|}\right) \in [0, \pi].$$

Two vectors $\mathbf{v}$ and $\mathbf{w}$ are said to be *orthogonal* [orthogonal], denoted $\mathbf{v} \perp \mathbf{w}$, if $\langle \mathbf{v}, \mathbf{w} \rangle = 0$.

Clearly for $\mathbf{v}, \mathbf{w} \neq \mathbf{0}$, $\mathbf{v} \perp \mathbf{w}$ iff $\angle(\mathbf{v}, \mathbf{w}) = \pi/2$.

2.1.2 $\mathbb{C}^n$ as a unitary space

For the complex analogue $\mathbb{C}^n$, the standard scalar product also needs to return a non-negative *real* number for the induced norm $\|\mathbf{v}\| = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}$, rather than an arbitrary complex number. Recall that the absolute value of a complex number $z = x + iy$ is $|z| = \sqrt{x^2 + y^2} = \sqrt{z\bar{z}}$. Complex conjugation is used to the same effect in the following definition of the standard scalar product on $\mathbb{C}^n$.

We adopt the notation $\bar{\mathbf{u}}$ for the vector obtained from $\mathbf{u} \in \mathbb{C}^n$ by component-wise complex conjugation:

$$\mathbf{u} = \begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix} \in \mathbb{C}^n \quad \longmapsto \quad \bar{\mathbf{u}} = \begin{pmatrix} \bar{u}_1 \\ \vdots \\ \bar{u}_n \end{pmatrix} \in \mathbb{C}^n.$$

Note that the map $\mathbf{u} \mapsto \bar{\mathbf{u}}$ is not linear over $\mathbb{C}^n$. While it is compatible with vector addition in $\mathbb{C}^n$, it is not compatible with scalar multiplication: instead of $\mathbb{C}$-linear w.r.t. scalars it is conjugated-linear, mapping $\lambda \mathbf{u}$ to $\bar{\lambda}\bar{\mathbf{u}}$ rather than to $\lambda \bar{\mathbf{u}}$.

We shall also often need the row vector corresponding to $\bar{\mathbf{u}}$, which is the transpose $\bar{\mathbf{u}}^t$. For the combined operation of (component-wise) complex conjugation and transposition we use the standard notation

$$\mathbf{u} = \begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix} \longmapsto \quad \mathbf{u}^+ := \bar{\mathbf{u}}^t = (\bar{u}_1, \ldots, \bar{u}_n),$$

which really is a special case of the adjoint operation (on matrices) to be defined in Definition 2.1.8 below.

Definition 2.1.5 The *standard scalar product* [1] in $\mathbb{C}^n$ is the map

$$(\mathbf{v}, \mathbf{w}) \longmapsto \langle \mathbf{v}, \mathbf{w} \rangle = \sum_{i=1}^{n} \bar{v}_i w_i = (\bar{v}_1, \ldots, \bar{v}_n) \begin{pmatrix} w_1 \\ \vdots \\ w_n \end{pmatrix} = \overline{\mathbf{v}}^t \cdot \mathbf{w} = \mathbf{v}^+ \mathbf{w}.$$

The $\mathbb{C}$-vector space $\mathbb{C}^n$ with this scalar product is referred to as the (standard) n-dimensional *unitary vector space* $\mathbb{C}^n$.

Again we collect properties of this standard scalar product, viewed as a binary function

$$\begin{array}{ccc} V \times V & \longrightarrow & \mathbb{C} \\ (\mathbf{v}, \mathbf{w}) & \longmapsto & \langle \mathbf{v}, \mathbf{w} \rangle, \end{array}$$

which could also be phrased for any binary function $\sigma \colon V \times V \to \mathbb{C}$.

Definition 2.1.6 A map $\langle ., . \rangle \colon V \times V \to \mathbb{C}$ is

(a) *semi-bilinear* (also: *sesquilinear*) [semi-bilinear, sesquilinear] if for all $\mathbf{v}, \mathbf{v}_1, \mathbf{v}_2, \mathbf{w}, \mathbf{w}_1, \mathbf{w}_2 \in V$ and $\lambda_1, \lambda_2 \in \mathbb{C}$

    (i) (linear in its second argument)

$$\langle \mathbf{v}, \lambda_1 \mathbf{w}_1 + \lambda_2 \mathbf{w}_2 \rangle = \lambda_1 \langle \mathbf{v}, \mathbf{w}_1 \rangle + \lambda_2 \langle \mathbf{v}, \mathbf{w}_2 \rangle.$$

    (ii) (conjugated-linear in the first argument)

$$\langle \lambda_1 \mathbf{v}_1 + \lambda_2 \mathbf{v}_2, \mathbf{w} \rangle = \bar{\lambda}_1 \langle \mathbf{v}_1, \mathbf{w} \rangle + \bar{\lambda}_2 \langle \mathbf{v}_2, \mathbf{w} \rangle.$$

(b) *hermitian* [hermitesch], if for all $\mathbf{v}, \mathbf{w} \in V$

$$\langle \mathbf{v}, \mathbf{w} \rangle = \overline{\langle \mathbf{w}, \mathbf{v} \rangle}.$$

(c) *positive definite* [positiv definit], if for all $\mathbf{v} \in V$

$$\langle \mathbf{v}, \mathbf{v} \rangle \geqslant 0 \text{ in } \mathbb{R}, \text{ and } \langle \mathbf{v}, \mathbf{v} \rangle = 0 \text{ iff } \mathbf{v} = \mathbf{0}.$$

Any σ satisfying conditions (a), (b), (c) is a (complex) *scalar product* on V, usually denoted $\langle ., . \rangle$.

A $\mathbb{C}$-vector space V with a scalar product $\langle ., . \rangle$, $(V, \langle ., . \rangle)$, is called a *unitary vector space* [unitärer Vektorraum].

---

[1] Beware that there are two conventions in the literature, as to whether a complex scalar product is linear in its first or in its second argument. We here follow the convention of linearity in the *second* argument.

Note that $\langle \mathbf{v}, \mathbf{v} \rangle \in \mathbb{R}$ follows from (b), as we have for all $z \in \mathbb{C}$ that $z \in \mathbb{R}$ iff $z = \bar{z}$. Also, (ii) in (a) follows from (i) together with (b).

Lengths and orthogonality in $\mathbb{C}^n$ are defined w.r.t. the standard scalar product in complete analogy with $\mathbb{R}^n$.

Definition 2.1.7 In $\mathbb{C}^n$ with its standard scalar product $\langle \mathbf{v}, \mathbf{w} \rangle = \mathbf{v}^+\mathbf{w}$, the *length* or *norm* of vector $\mathbf{v} \in \mathbb{C}^n$ is

$$\|\mathbf{v}\| := \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle} = \sqrt{\sum_{i=1}^{n} \bar{v}_i v_i} = \sqrt{\sum_{i=1}^{n} |v_i|^2}.$$

Vectors of length 1, $\|\mathbf{v}\| = 1$, are called *unit vectors* [Einheitsvektoren].

Vectors $\mathbf{v}, \mathbf{w} \in \mathbb{C}^n$ are called *orthogonal* iff $\langle \mathbf{v}, \mathbf{w} \rangle = 0$.

2.1.3   Bilinear and semi-bilinear forms

Before further exploration of scalar products in $\mathbb{R}$- and $\mathbb{C}$-vector spaces, we briefly consider bilinear and semi-bilinear forms in general and in particular link them to matrices for their representation w.r.t. to chosen bases. Keep in mind that scalar products are special bilinear/semi-bilinear forms.

Bilinear forms in $\mathbb{R}$-vector spaces

Let V be an n-dimensional $\mathbb{R}$-vector space. Recall that a map $\sigma \colon V \times V \to \mathbb{R}$ is called a *bilinear form* [Bilinearform] if it is linear in each argument.

With respect to a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of V, a bilinear form σ on V is uniquely determined by the values

$$a_{ij} := \sigma(\mathbf{b}_i, \mathbf{b}_j) \in \mathbb{R},$$

since, by bilinearity, for $\mathbf{v} = \sum_i \lambda_i \mathbf{b}_i$ and $\mathbf{w} = \sum_i \mu_i \mathbf{b}_i$, we get

$$\sigma(\mathbf{v}, \mathbf{w}) = \sigma(\textstyle\sum_i \lambda_i \mathbf{b}_i, \sum_i \mu_i \mathbf{b}_i) = \sum_{i,j} \lambda_i \mu_j \sigma(\mathbf{b}_i, \mathbf{b}_j) = \sum_{i,j} \lambda_i \mu_j a_{ij}.$$

Associate with σ the matrix $A = [\![\sigma]\!]^B = (a_{ij}) \in \mathbb{R}^{(n,n)}$ as its representation with respect to basis B,

$$[\![\sigma]\!]^B := (a_{ij}) \in \mathbb{R}^{(n,n)} \quad \text{where } a_{ij} = \sigma(\mathbf{b}_i, \mathbf{b}_j).$$

Then, in terms of coefficients w.r.t. basis B, and for vectors $\mathbf{v} = \sum_i \lambda_i \mathbf{b}_i$ and $\mathbf{w} = \sum_i \mu_i \mathbf{b}_i$:

$$\sigma(\mathbf{v}, \mathbf{w}) = \sum_{i,j} \lambda_i \mu_j a_{ij} = (\lambda_1, \ldots, \lambda_n) \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix}$$

$$= [\![\mathbf{v}]\!]_B^t \, [\![\sigma]\!]^B \, [\![\mathbf{w}]\!]_B.$$

Application of σ corresponds to a matrix product that multiplies the row and column vectors corresponding to $\mathbf{v}$ and $\mathbf{w}$, respectively, from the left and right into matrix $A = [\![\sigma]\!]^B$.

The matrix that represents the standard scalar product w.r.t. to the standard basis of $\mathbb{R}^n$, is just the unit matrix E_n.

Exercise 2.1.1 Show that a bilinear form σ on V is symmetric, i.e., that $\sigma(\mathbf{v}, \mathbf{w}) = \sigma(\mathbf{w}, \mathbf{v})$ for all $\mathbf{v}, \mathbf{w} \in V$, iff the matrix $A = [\![\sigma]\!]^B$ that represents σ w.r.t. some (any) basis B is symmetric in the sense that $A = A^t$.

A usefull alternative view of the representation

$$\sigma(\mathbf{v}, \mathbf{w}) = [\![\mathbf{v}]\!]_B^t [\![\sigma]\!]^B [\![\mathbf{w}]\!]_B = \sum_{i,j} \lambda_i \mu_j a_{ij} = (\lambda_1, \ldots, \lambda_n) \, A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix},$$

where $A[\![\sigma]\!]^B = (a_{ij}) = (\sigma(\mathbf{b}_i, \mathbf{b}_j))$ is the following. Regard the column vector $A(\mu_1, \ldots, \mu_n)$ as the representation of the image of $\mathbf{w}$ under the linear map φ for which $[\![\varphi]\!]_B^B = A$:

$$\mathbf{w} \longmapsto \varphi(\mathbf{w}) \quad \text{w.r.t. basis } B \colon [\![\mathbf{w}]\!]_B = \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix} \longmapsto [\![\varphi(\mathbf{w})]\!]_B = A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix}.$$

In these terms, $\sigma(\mathbf{v}, \mathbf{w})$ is the standard scalar product of $\mathbf{v}$ and $\varphi(\mathbf{w})$:

$$\sigma(\mathbf{v}, \mathbf{w}) = (\lambda_1, \ldots, \lambda_n) \, A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix} = \langle \mathbf{v}, \varphi(\mathbf{w}) \rangle = [\![\mathbf{v}]\!]_B^t [\![\varphi]\!]_B^B [\![\mathbf{w}]\!]_B.$$

Semi-bilinear forms in a $\mathbb{C}$-vector space

Let V be an n-dimensional $\mathbb{C}$-vector space, $\sigma \colon V \times V \to \mathbb{C}$ a semi-bilinear form (cf. Definition 2.1.6 (a)). With respect to a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of a $\mathbb{C}$-vector space V, a semi-bilinear form σ is represented by the matrix $A = [\![\sigma]\!]^B \in \mathbb{C}^{(n,n)}$ whose entries are

$$a_{ij} := \sigma(\mathbf{b}_i, \mathbf{b}_j) \in \mathbb{C}.$$

By semi-bilinearity, for $\mathbf{v} = \sum_i \lambda_i \mathbf{b}_i$ and $\mathbf{w} = \sum_i \mu_i \mathbf{b}_i$, we get

$$\sigma(\mathbf{v}, \mathbf{w}) \;=\; \sigma(\textstyle\sum_i \lambda_i \mathbf{b}_i, \sum_i \mu_i \mathbf{b}_i) = \sum_{i,j} \bar{\lambda}_i \mu_j \sigma(\mathbf{b}_i, \mathbf{b}_j) = \sum_{i,j} \bar{\lambda}_i \mu_j a_{ij}$$

$$= \; (\bar{\lambda}_1, \ldots, \bar{\lambda}_n) \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix} = \overline{(\lambda_1, \ldots, \lambda_n)} \, A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix}$$

$$= \; [\![\mathbf{v}]\!]_B^+ [\![\sigma]\!]^B [\![\mathbf{w}]\!]_B.$$

So, application of σ corresponds to a matrix product that multiplies the row and column vectors corresponding to the complex conjugate of $\mathbf{v}$ and $\mathbf{w}$, respectively, from the left and right into matrix $A = [\![\sigma]\!]^B$. Again, the unit matrix represents the standard scalar product (w.r.t. any basis).

Definition 2.1.8 For a matrix $A \in \mathbb{C}^{(n,n)}$ its *adjoint matrix* [Adjungierte] is defined as the matrix $A^+ = (\overline{A})^t = \overline{A^t}$, the result of taking the complex conjugate of all entries and applying transposition. For $A = (a_{ij})$, the entries a_{ij}^+ of A^+ are $a_{ij}^+ = \bar{a}_{ji}$.

A matrix $A \in \mathbb{C}^{(n,n)}$ is called *hermitian* or *self-adjoint* [hermitisch, selbstadjungiert] iff $A = A^+$.

Exercise 2.1.2 Show that a semi-bilinear form σ on V is hermitian iff the matrix $A = [\![\sigma]\!]^B$ that represents σ w.r.t. some (any) basis B is hermitian.

Observation 2.1.9 *The following rules apply w.r.t. the operations of complex conjugation and transposition of matrices in $\mathbb{C}^{(n,n)}$:*
  (i) $\overline{A + B} = \overline{A} + \overline{B}$ *and* $\overline{AB} = \overline{A}\,\overline{B}$.
  (ii) $(A + B)^t = A^t + B^t$ *and* $(AB)^t = B^t A^t$ *(order inverted!).*
  (iii) $(A + B)^+ = A^+ + B^+$ *and* $(AB)^+ = B^+ A^+$ *(order inverted!).*

Exercise 2.1.3 Find the transformation pattern for the representational matrices $A = [\![\sigma]\!]^B$ and $\hat{A} = [\![\sigma]\!]^{B'}$ of a bilinear form σ w.r.t. two different bases $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ and $B' = (\mathbf{b}'_1, \ldots, \mathbf{b}'_n)$ of an n-dimensional $\mathbb{R}$-vector space V. Similarly find the transformation pattern for the representational matrices of a semi-bilinear form σ of an n-dimensional $\mathbb{C}$-vector space V.

In both cases, use the basis transformation matrices $C = [\![\mathrm{id}_V]\!]^B_{B'}$ and $C^{-1} = [\![\mathrm{id}_V]\!]^{B'}_B$ as well as their transposes and adjoints as appropriate.

2.1.4 Scalar products in euclidean and unitary spaces

The analogy between scalar products in $\mathbb{R}^n$ and $\mathbb{C}^n$ suggests to treat $\mathbb{R}$ and $\mathbb{C}$-vector spaces together as far as possible. In fact, the definition of a real scalar product in Definition 2.1.2 is just the specialisation of the complex case in Definition 2.1.6. Viewing $\mathbb{R} \subsetneq \mathbb{C}$, as characterised by the condition that $\lambda \in \mathbb{R}$ iff $\bar{\lambda} = \lambda$, and noting that $\bar{\mathbf{u}} = \mathbf{u}$ for $\mathbf{u} \in \mathbb{R}^n$, we see that a (real-valued) positive definite hermitian semi-bilinear form *is* symmetric and bilinear over $\mathbb{R}^n$. ([2])

------

[2]Similarly, for the associated matrices, a matrix $A \in \mathbb{R}^{(n,n)} \subseteq \mathbb{C}^{(n,n)}$ is hermitian iff it is symmetric, as complex conjugation does not affect its entries. Also see section 2.4.2

Proposition 2.1.10 (Cauchy-Schwarz inequality) *Let $(V, \langle ., . \rangle)$ be a euclidean or unitary vector space, with induced norm $\|\mathbf{v}\| = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}$. Then, for all $\mathbf{v}, \mathbf{w} \in V$:*

$$|\langle \mathbf{v}, \mathbf{w} \rangle| \leqslant \|\mathbf{v}\| \cdot \|\mathbf{w}\|.$$

Proof.   The claimed inequality is equivalent to $|\langle \mathbf{v}, \mathbf{w} \rangle|^2 \leqslant \|\mathbf{v}\|^2 \cdot \|\mathbf{w}\|^2$, or

$$\langle \mathbf{v}, \mathbf{w} \rangle \overline{\langle \mathbf{v}, \mathbf{w} \rangle} \leqslant \langle \mathbf{v}, \mathbf{v} \rangle \langle \mathbf{w}, \mathbf{w} \rangle.$$

Clearly the inequality is true if either $\mathbf{v}$ or $\mathbf{w}$ is $\mathbf{0}$. Assuming therefore that $\mathbf{v}, \mathbf{w} \neq \mathbf{0}$, and hence $\langle \mathbf{v}, \mathbf{v} \rangle, \langle \mathbf{w}, \mathbf{w} \rangle > 0$, we consider

$$\mathbf{u} := \mathbf{w} - \frac{\langle \mathbf{v}, \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \mathbf{v}.$$

Let $\lambda := \frac{\langle \mathbf{v}, \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle}$. Geometrically, $\frac{\langle \mathbf{v}, \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle} \mathbf{v}$ is the orthogonal projection of $\mathbf{w}$ onto $\mathbf{v}$ (as indicated in the diagram).

Now

$$
\begin{aligned}
0 \leqslant \langle \mathbf{u}, \mathbf{u} \rangle &= \langle \mathbf{w} - \lambda \mathbf{v}, \mathbf{w} - \lambda \mathbf{v} \rangle \\
&= \langle \mathbf{w}, \mathbf{w} \rangle - \bar{\lambda}\langle \mathbf{v}, \mathbf{w} \rangle - \lambda \langle \mathbf{w}, \mathbf{v} \rangle + \lambda\bar{\lambda}\langle \mathbf{v}, \mathbf{v} \rangle \\
&= \langle \mathbf{w}, \mathbf{w} \rangle - \bar{\lambda}\langle \mathbf{v}, \mathbf{w} \rangle - \lambda \langle \mathbf{w}, \mathbf{v} \rangle + \bar{\lambda}\langle \mathbf{v}, \mathbf{w} \rangle \\
&= \langle \mathbf{w}, \mathbf{w} \rangle - \lambda \langle \mathbf{w}, \mathbf{v} \rangle \\
&= \langle \mathbf{w}, \mathbf{w} \rangle - \frac{\langle \mathbf{w}, \mathbf{v} \rangle \langle \mathbf{v}, \mathbf{w} \rangle}{\langle \mathbf{v}, \mathbf{v} \rangle},
\end{aligned}
$$

whence $\langle \mathbf{v}, \mathbf{w} \rangle \overline{\langle \mathbf{v}, \mathbf{w} \rangle} \leqslant \langle \mathbf{v}, \mathbf{v} \rangle \langle \mathbf{w}, \mathbf{w} \rangle$, as claimed.

$\square$

The norm function $\mathbf{v} \mapsto \|\mathbf{v}\| = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}$ associated with the scalar product $\langle ., . \rangle$ can be used to introduce a *metric* on the euclidean or unitary vector space V. This is the familiar euclidean distance $d(\mathbf{v}, \mathbf{w}) = \sqrt{\sum_{i=1}^{n} |v_i - w_i|^2}$ for the standard scalar products over $\mathbb{R}^n$ or $\mathbb{C}^n$ with their norms.

---

below.

Definition 2.1.11 The *metric* [Metrik] induced on V by the norm $\mathbf{v} \mapsto \|\mathbf{v}\|$ is the distance function $d \colon V \times V \to \mathbb{R}$,

$$d(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\|.$$

This distance function satisfies the axioms of a metric. For all $\mathbf{u}, \mathbf{v}, \mathbf{w}$:

  (i) (symmetry) $d(\mathbf{v}, \mathbf{w}) = d(\mathbf{w}, \mathbf{v})$.

  (ii) (positivity) $d(\mathbf{v}, \mathbf{w}) \geqslant 0$, with $d(\mathbf{v}, \mathbf{w}) = 0$ iff $\mathbf{v} = \mathbf{w}$.

  (iii) (triangle inequality) $d(\mathbf{u}, \mathbf{w}) \leqslant d(\mathbf{u}, \mathbf{v}) + d(\mathbf{v}, \mathbf{w})$.

The first two properties are immediate from the axioms for scalar products; for the third one, one uses the Cauchy-Schwarz inequality:

Lemma 2.1.12 *The distance function* $d(\mathbf{v}, \mathbf{w}) = \|\mathbf{v} - \mathbf{w}\| = \sqrt{\langle \mathbf{v} - \mathbf{w}, \mathbf{v} - \mathbf{w} \rangle}$ *induced by the norm from a scalar product* $\langle ., . \rangle$ *on* V *satisfies the triangle inequality: for all* $\mathbf{u}, \mathbf{v}, \mathbf{w}$:

$$d(\mathbf{u}, \mathbf{w}) \leqslant d(\mathbf{u}, \mathbf{v}) + d(\mathbf{v}, \mathbf{w}).$$

Proof.   Let $\mathbf{a} := \mathbf{w} - \mathbf{v}$, $\mathbf{b} := \mathbf{v} - \mathbf{u}$ and $\mathbf{c} = \mathbf{w} - \mathbf{u}$. Then

$$
\begin{aligned}
(d(\mathbf{u}, \mathbf{w}))^2 &= \langle \mathbf{c}, \mathbf{c} \rangle \\
&= \langle \mathbf{a} + \mathbf{b}, \mathbf{a} + \mathbf{b} \rangle \\
&= \langle \mathbf{a}, \mathbf{a} \rangle + \langle \mathbf{a}, \mathbf{b} \rangle + \langle \mathbf{b}, \mathbf{a} \rangle + \langle \mathbf{b}, \mathbf{b} \rangle \\
&\leqslant \|\mathbf{a}\|^2 + 2|\langle \mathbf{a}, \mathbf{b} \rangle| + \|\mathbf{b}\|^2 \\
&\leqslant \|\mathbf{a}\|^2 + 2\|\mathbf{a}\| \cdot \|\mathbf{b}\| + \|\mathbf{b}\|^2 \qquad \text{(Cauchy-Schwarz)} \\
&= (\|\mathbf{a}\| + \|\mathbf{b}\|)^2 \\
&= (d(\mathbf{u}, \mathbf{v}) + d(\mathbf{v}, \mathbf{w}))^2.
\end{aligned}
$$

$\square$

Exercise 2.1.4 Show that equality holds in the Cauchy-Schwarz inequality, i.e., $|\langle \mathbf{v}, \mathbf{w} \rangle| = \|\mathbf{v}\| \cdot \|\mathbf{w}\|$, if and only if $\mathbf{v}$ and $\mathbf{w}$ are linearly dependent.

Exercise 2.1.5 Let $\mathbf{u}, \mathbf{v}, \mathbf{w}$ be pairwise distinct, $\mathbf{a} := \mathbf{v} - \mathbf{u}$, $\mathbf{b} := \mathbf{w} - \mathbf{v}$. Show that equality holds in the triangle inequality

$$d(\mathbf{u}, \mathbf{w}) = d(\mathbf{u}, \mathbf{v}) + d(\mathbf{v}, \mathbf{w}), \quad \text{or, equivalently } \|\mathbf{a} + \mathbf{b}\| = \|\mathbf{a}\| + \|\mathbf{b}\|,$$

if and only if $\mathbf{a}$ and $\mathbf{b}$ are *positive real* scalar multiples of each other (geometrically: $\mathbf{v} = \mathbf{u} + s(\mathbf{w} - \mathbf{u})$ for some $s \in (0, 1) \subseteq \mathbb{R}$).

The following *polarisation equalities* express the scalar product in terms of its induced norm, showing that any such norm determines 'its' scalar product.

Proposition 2.1.13 *Let* $\mathbf{v} \mapsto \|\mathbf{v}\| = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}$ *be the induced norm in a euclidean or unitary vector space* V, *respectively. Then* $\langle \mathbf{v}, \mathbf{w} \rangle$ *can be expressed in terms of norms in the following ways.*

In the euclidean case:

$$\langle \mathbf{v}, \mathbf{w} \rangle = \frac{1}{4} \big(\|\mathbf{v} + \mathbf{w}\|^2 - \|\mathbf{v} - \mathbf{w}\|^2 \big),$$

and in the unitary case:

$$\langle \mathbf{v}, \mathbf{w} \rangle = \frac{1}{4} \big(\|\mathbf{v} + \mathbf{w}\|^2 - \|\mathbf{v} - \mathbf{w}\|^2 - i\|\mathbf{v} + i\mathbf{w}\|^2 + i\|\mathbf{v} - i\mathbf{w}\|^2 \big).$$

Proof.   Do the calculations as an exercise.

□

Exercise 2.1.6 Show that for any linear map $\rho\colon V \to V$ from a euclidean or unitary vector space $(V, \langle ., . \rangle)$ into itself the following are equivalent:

 (i) $\rho(\mathbf{0}) = \mathbf{0}$ and ρ preserves distances:
      $d(\rho(\mathbf{v}), \rho(\mathbf{w})) = d(\mathbf{v}, \mathbf{w})$ for all $\mathbf{v}, \mathbf{w} \in V$.
 (ii) ρ preserves norms: $\|\rho(\mathbf{v})\| = \|\mathbf{v}\|$ for all $\mathbf{v} \in V$.
 (iii) ρ preserves scalar products: $\langle \rho(\mathbf{v}), \rho(\mathbf{w}) \rangle = \langle \mathbf{v}, \mathbf{w} \rangle$ for all $\mathbf{v}, \mathbf{w} \in V$.

For the crucial implication, (ii) $\Rightarrow$ (iii), use the above proposition.
[In the euclidean case, (i) even implies linearity, for any function ρ, whence (i) also characterises euclidean isometries, as defined in Definition 2.1.14 below.]

Exercise 2.1.7 Prove Pythagoras' theorem in the following form for orthogonal vectors in a euclidean or unitary space V:

$$\mathbf{v} \perp \mathbf{w} \quad \text{implies that} \quad \|\mathbf{v} + \mathbf{w}\|^2 = \|\mathbf{v}\|^2 + \|\mathbf{w}\|^2.$$

Isometries of euclidean and unitary spaces

We are particularly interested in linear maps between vector spaces that preserve the metric structure along with the linear structure.

Definition 2.1.14 Let $(V, \langle ., . \rangle^V)$ and $(W, \langle ., . \rangle^W)$ be two euclidean vector spaces (or two unitary vector spaces) with corresponding scalar products indexed to indicate where they belong. A linear map $\varphi \in \mathrm{Hom}(V, W)$ is called an *isometry* [Isometrie] if for all $\mathbf{v}_1, \mathbf{v}_2 \in V$:

$$\langle \varphi(\mathbf{v}_1), \varphi(\mathbf{v}_2) \rangle^W = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle^V.$$

Note that an isometry is necessarily injective: $\varphi(\mathbf{v}) = \mathbf{0}$ iff $\langle \varphi(\mathbf{v}), \varphi(\mathbf{v}) \rangle = 0$ iff $\langle \mathbf{v}, \mathbf{v} \rangle = 0$ iff $\mathbf{v} = \mathbf{0}$ shows that $\ker(\varphi) = \{\mathbf{0}\}$.

Clearly, as an isometry preserves the scalar product, it in particular also preserves the induced norms and distances. Exercise 2.1.6 based on Proposition 2.1.13 shows that isometries could equivalently be defined in terms of distance preservation or norm preservation.

Observation 2.1.15 *For a linear map between euclidean or unitary vectors spaces $(V, \langle ., . \rangle^V)$ and $(W, \langle ., . \rangle^W)$, the following are equivalent:*

(i) *φ preserves distances:*
$d^W(\varphi(\mathbf{v}_1), \varphi(\mathbf{v}_2)) = d^V(\mathbf{v}_1, \mathbf{v}_2)$ for all $\mathbf{v}_1, \mathbf{v}_2 \in V$.

(ii) *φ preserves norms: $\|\varphi(\mathbf{v})\|^W = \|\mathbf{v}\|^V$ for all $\mathbf{v} \in V$.*

(iii) *φ preserves scalar products:*
$\langle \varphi(\mathbf{v}_1), \varphi(\mathbf{v}_2) \rangle^W = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle^V$ for all $\mathbf{v}_1, \mathbf{v}_2 \in V$.

We shall further examine isometries and automorphisms of euclidean and unitary spaces in section 2.3.3, after an exploration of orthogonality and orthonormal bases in section 2.3.

2.2   Further examples

Real and complex function spaces are often naturally equipped with scalar products. They provide natural examples of infinite-dimensional euclidean and unitary spaces.

Let $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$. Recall the $\mathbb{F}$-vector spaces $\mathcal{F}(A, \mathbb{F})$ consisting of all $\mathbb{F}$-valued functions $f \colon A \to \mathbb{F}$ on A, with point-wise addition and scalar multiplication. Special cases we considered were those with $A = \mathbb{N}$ (sequences) and $A = \mathbb{F}$ (total functions on $\mathbb{F}$).

Consider $C([0, 1], \mathbb{F}) \subseteq \mathcal{F}([0, 1], \mathbb{F})$, the subspace of continuous $\mathbb{F}$-valued functions on the unit interval in $\mathbb{R}$, $f \colon [0, 1] \to \mathbb{F}$, for either $\mathbb{F} = \mathbb{R}$ or $\mathbb{C}$. One verifies that this is indeed a linear subspace, and hence an $\mathbb{F}$-vector space.

Now consider the following (semi-)bilinear form over $C([0,1], \mathbb{F})$:

$$\langle f, g \rangle := \int_0^1 \bar{f}(x)g(x)\mathrm{d}x,$$

where $\bar{f}(x)$ is the function obtained from f by complex conjugation.

(Semi-)bilinearity is easily checked (integration itself is a linear operation!); symmetric or hermitian behaviour, respectively, are obvious. Also positivity, $\langle f, f \rangle = \int_0^1 |f|^2 \mathrm{d}x \geqslant 0$ and > 0 unless $f = 0$ are clear for continuous f. Hence we have a scalar product that turns $\big(C([0,1], \mathbb{F}), \langle ., . \rangle\big)$ into a euclidean or unitary space, respectively.

The induced norm is

$$\|f\| = \Big(\int_0^1 |f(x)|^2 \mathrm{d}x \Big)^{1/2}.$$

Exercise 2.2.1 Spell out the Cauchy-Schwarz and triangle inequalities for these spaces. Find families of pairwise orthogonal functions in $C([0,1], \mathbb{F})$, e.g., based on the functions $\sin(2\pi n x)$.

In the sequence space $\mathcal{F}(\mathbb{N}, \mathbb{F})$ we could try to generalise the standard scalar product by putting $\langle \mathbf{a}, \mathbf{b} \rangle := \sum_{i=0}^{\infty} \bar{a}_i b_i$. This works in the subspace of those sequences $\mathbf{a} \in \mathcal{F}(\mathbb{N}, \mathbb{F})$ for which $\sum_{i=0}^{\infty} |a_i|^2$ converges. For any two such sequences $\mathbf{a}$ and $\mathbf{b}$ one can show that also $\sum_{i=0}^{\infty} \bar{a}_i b_i$ is absolutely convergent.

Exercise 2.2.2 Show for $\mathbf{a}, \mathbf{b} \in \mathcal{F}(\mathbb{N}, \mathbb{C})$ for which $\sum_{i=0}^{\infty} |a_i|^2$ and $\sum_{i=0}^{\infty} |b_i|^2$ are convergent, $\sum_{i=0}^{\infty} \bar{a}_i b_i$ is absolutely convergent.
[Hint: relate the claim to an application of Cauchy-Schwarz in $\mathbb{R}^k$.]
Check that the set of sequences $\mathbf{a} \in \mathcal{F}(\mathbb{N}, \mathbb{C})$ for which $\sum_{i=0}^{\infty} |a_i|^2$ converges, forms a subspace of the $\mathbb{F}$-vector space $\mathcal{F}(\mathbb{N}, \mathbb{C})$. Show that this subspace is unitary, with the scalar product

$$\langle \mathbf{a}, \mathbf{b} \rangle := \sum_{i=0}^{\infty} \bar{a}_i b_i.$$

Show that the "generalised standard basis vectors" consisting of the sequences with a single 1 and zeroes otherwise, form an infinite family of pairwise orthogonal unit vectors; but they do *not* form a basis (fail to span the entire subspace, why?).

2.3 Orthogonality and orthonormal bases

2.3.1 Orthonormal bases

Definition 2.3.1 A set (or family) S of vectors in a euclidean or unitary vector space V is said to be an *orthonormal system* [Orthonormalsystem] if it consist of pairwise orthogonal unit vectors: $\|\mathbf{u}\| = 1$ for all $\mathbf{u} \in S$ and $\mathbf{u} \perp \mathbf{u}'$ for all $\mathbf{u} \neq \mathbf{u}'$ in S.

An orthonormal system that is a basis (of V or of some subspace $U \subseteq V$) is called an *orthonormal basis* [Orthonormalbasis].

Clearly the standard bases of $\mathbb{R}^n$ and $\mathbb{C}^n$ are orthonormal bases for the standard scalar products.

The coefficients of vectors w.r.t. an orthonormal basis can be expressed in terms of the scalar product very easily.

Lemma 2.3.2 *Let* $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ *be an orthonormal basis of* V. *For all* $\mathbf{v} \in V$:

$$\mathbf{v} = \sum_{i=1}^{n} \langle \mathbf{b}_i, \mathbf{v} \rangle \mathbf{b}_i, \qquad i.e., \ \ [\![\mathbf{v}]\!]_B = \begin{pmatrix} \langle \mathbf{b}_1, \mathbf{v} \rangle \\ \vdots \\ \langle \mathbf{b}_n, \mathbf{v} \rangle \end{pmatrix}.$$

Proof.   Simply express $\mathbf{v}$ as a linear combination $\mathbf{v} = \sum_i \lambda_i \mathbf{b}_i$ and apply scalar products with $\mathbf{b}_j$ to determine the coefficients λ_j:

$$\langle \mathbf{b}_j, \mathbf{v} \rangle = \langle \mathbf{b}_j, \sum_i \lambda_i \mathbf{b}_i \rangle = \sum_i \lambda_i \langle \mathbf{b}_j, \mathbf{b}_i \rangle = \lambda_j.$$

$\square$

Corollary 2.3.3 *A basis* $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ *of* V *is orthonormal iff the matrix representation of the scalar product* $\langle ., . \rangle$ *as a (semi-)bilinear form w.r.t.* B *is the unit matrix:* $[\![\sigma]\!]^B = E_n$ *for* $\sigma(\mathbf{v}, \mathbf{w}) := \langle \mathbf{v}, \mathbf{w} \rangle$.

In other words: in terms of coefficients w.r.t. an orthonormal basis, the scalar product $\langle ., . \rangle$ of V is computed just like the standard scalar product in the corresponding standard space $\mathbb{R}^n$ or $\mathbb{C}^n$.

Gram-Schmidt orthonormalisation

The following is an existence theorem for orthonormal bases. We restrict attention to finite-dimensional euclidean or unitary vector spaces $(V, \langle ., . \rangle)$. The constructive proof provides a method to obtain an orthonormal basis from an arbitrary given basis.

Theorem 2.3.4 *Let $(V, \langle ., . \rangle)$ be an n-dimensional euclidean or unitary vector space with basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$.*
Then there is an orthonormal basis $\hat{B} = (\hat{\mathbf{b}}_1, \ldots, \hat{\mathbf{b}}_n)$ for V such that for $1 \leqslant k \leqslant n$, $\operatorname{span}(\mathbf{b}_1, \ldots, \mathbf{b}_k) = \operatorname{span}(\hat{\mathbf{b}}_1, \ldots, \hat{\mathbf{b}}_k)$.

Proof. The proof is by the Gram-Schmidt procedure that successively constructs the new basis vectors form the given ones. Formally we proceed by induction on k, constructing a sequence $\hat{B}_1 \subseteq \hat{B}_1 \subseteq \cdots \subseteq \hat{B}_k$ such that $\hat{B}_k = (\hat{\mathbf{b}}_1, \ldots, \hat{\mathbf{b}}_k)$ is an orthonormal basis of $U_k := \operatorname{span}(\mathbf{b}_1, \ldots, \mathbf{b}_k) \subseteq V$.

Base case: $k = 1$. Let $\hat{\mathbf{b}}_1 := \frac{1}{\|\mathbf{b}_1\|} \mathbf{b}_1$. Clearly $\hat{B}_1 := \{\hat{\mathbf{b}}_1\}$ is as desired.

In the induction step we assume that $\hat{B}_k = (\hat{\mathbf{b}}_1, \ldots, \hat{\mathbf{b}}_k)$ is an orthonormal basis of U_k and need to find a next unit vector $\hat{\mathbf{b}}_{k+1}$ such that

  (i) $\hat{\mathbf{b}}_{k+1} \in U_{k+1} = U_k + \operatorname{span}(\mathbf{b}_{k+1}) = \operatorname{span}(\hat{B}_k) + \operatorname{span}(\mathbf{b}_{k+1})$.

  (ii) $\|\hat{\mathbf{b}}_{k+1}\| = 1$ and $\hat{\mathbf{b}}_{k+1} \perp \operatorname{span}(\hat{B}_k)$, i.e., $\hat{\mathbf{b}}_{k+1} \perp \hat{\mathbf{b}}_i$ for $i = 1, \ldots, k$.

Note that (i) and (ii) together imply that $\hat{B}_{k+1} := (\hat{B}_k, \hat{\mathbf{b}}_{k+1})$ spans U_{k+1} and that it is an orthonormal basis of this subspace.

Put

$$\mathbf{u} := \sum_{i=1}^{k} \langle \hat{\mathbf{b}}_i, \mathbf{b}_{k+1} \rangle \hat{\mathbf{b}}_i.$$

Geometrically, $\mathbf{u}$ is the orthogonal projection of $\mathbf{b}_{k+1}$ onto $U_k = \operatorname{span}(\hat{B}_k)$ (see section 2.3.2 below). Clearly $\mathbf{u} \in U_k$ and $(\mathbf{b}_{k+1} - \mathbf{u}) \perp U_k$, as for every basis vector $\hat{\mathbf{b}}_j$ of U_k:

$$\langle \hat{\mathbf{b}}_j, \mathbf{b}_{k+1} - \mathbf{u} \rangle = \langle \hat{\mathbf{b}}_j, \mathbf{b}_{k+1} \rangle - \sum_{i=1}^{k} \langle \hat{\mathbf{b}}_i, \mathbf{b}_{k+1} \rangle \langle \hat{\mathbf{b}}_j, \hat{\mathbf{b}}_i \rangle = \langle \hat{\mathbf{b}}_j, \mathbf{b}_{k+1} \rangle - \langle \hat{\mathbf{b}}_j, \mathbf{b}_{k+1} \rangle = 0.$$

It follows that

$$\hat{\mathbf{b}}_{k+1} := \frac{\mathbf{b}_{k+1} - \mathbf{u}}{\|\mathbf{b}_{k+1} - \mathbf{u}\|}$$

is as desired. (Note that $\mathbf{b}_{k+1} - \mathbf{u}$ cannot be $\mathbf{0}$ because $\mathbf{b}_{k+1}$ is not a linear combination over $\hat{B}_k$.)

$\square$

The above technique adapts to give corresponding basis extension constructions, like the following. Its use is demonstrated below in connection with orthogonal complements and projections.

Corollary 2.3.5 *Let $U \subseteq V$ be a subspace of the finite-dimensional euclidean or unitary vector space V. Then any orthonormal basis B_0 for U can be extended to an orthonormal basis B of V.*

In the language of Definition 2.3.6 below, we shall see that the basis vectors in $B_1 := B \setminus B_0$ form an orthonormal basis of the *orthogonal complement* [orthogonales Komplement] $U^\perp$ of U.

Proof.   Let $B_0 = (\mathbf{b}_1, \dots, \mathbf{b}_m)$ and let $B' = (\mathbf{b}_1, \dots, \mathbf{b}_m, \mathbf{b}_{m+1}, \dots, \mathbf{b}_n)$ be any extension of B_0 to a basis of V. Applying Gram-Schmidt to B' we firstly reproduce $\hat{\mathbf{b}}_i = \mathbf{b}_i$ for $i = 1, \dots, m$ as these are already orthonormal; the vectors $\hat{\mathbf{b}}_{m+1}, \dots, \hat{\mathbf{b}}_n$ then extend B_0 as required.

$\square$

2.3.2   Orthogonality and orthogonal complements

Again, we restrict attention to finite-dimensional euclidean or unitary spaces.

Definition 2.3.6 Let $U, U_1, U_2 \subseteq V$ be subspaces of the finite-dimensional euclidean or unitary vector space $(V, \langle ., . \rangle)$.

(i) U_1 and U_2 are *orthogonal* [orthogonal], denoted $U_1 \perp U_2$, if for all $\mathbf{u}_1 \in U_1$ and $\mathbf{u}_2 \in U_2$, $\langle \mathbf{u}_1, \mathbf{u}_2 \rangle = 0$.

(ii) The *orthogonal complement* [orthogonales Komplement] of the subspace $U \subseteq V$ is the subspace $U^\perp := \{ \mathbf{v} \in V \colon \mathbf{v} \perp \mathbf{u} \text{ for all } \mathbf{u} \in U \}$.

We also write $\mathbf{u} \perp U$ to indicate that the single vector $\mathbf{u}$ (or the subspace spanned by it) is orthogonal on the subspace U.

Exercise 2.3.1 Show that $\mathbf{u} \perp U$ iff $\mathbf{u} \perp \mathbf{b}$ for every basis vector $\mathbf{b}$ from some basis B of U. Check that $U^\perp = \{ \mathbf{v} \in V \colon \mathbf{v} \perp \mathbf{u} \text{ for all } \mathbf{u} \in U \}$ is a subspace of V.

Lemma 2.3.7 *In the situation of the definition:*

  (i) *the sum of orthogonal subspaces is direct, i.e., $U_1 \cap U_2 = \{\mathbf{0}\}$ whenever $U_1 \perp U_2$.*

  (ii) *the subspace $U^\perp \subseteq V$ (the orthogonal complement of U) is in particular a linear complement of the subspace U (cf. section 2.6.2 of part I): $V = U^\perp \oplus U$.*

Proof. For (i) let $\mathbf{v} \in U_1 \cap U_2$: then $\langle \mathbf{v}, \mathbf{v} \rangle = 0$ implies that $\mathbf{v} = \mathbf{0}$.

For (ii), one easily checks the subspace criterion (cf. Exercise 2.3.1). It remains to show that $U + U^\perp = V$. We may extend an orthonormal basis B_0 of U to an orthonormal basis $B \supseteq B_0$ of V. Let $B_1 = B \setminus B_0$. We claim that $U^\perp = \mathrm{span}(B_1)$. Clearly $B_1 \subseteq U^\perp$, as all basis vectors in B_1 are orthogonal with those in B_0, since B is orthonormal. But since $V = \mathrm{span}(B_0) \oplus \mathrm{span}(B_1) = U \oplus \mathrm{span}(B_1)$ we have that $\dim(\mathrm{span}(B_1)) = \dim(V) - \dim(U) = \dim(U^\perp)$ as the sum of U and $U^\perp$ is also direct by (i). Therefore $B_1 \subseteq U^\perp$ implies $\mathrm{span}(B_1) = U^\perp$, and $V = \mathrm{span}(B_0) \oplus \mathrm{span}(B_1)$ implies $V = U \oplus U^\perp$.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

For $U \subseteq V$ with orthogonal complement $U^\perp$, $V = U \oplus U^\perp$ implies that every vector $\mathbf{v} \in V$ has a unique decomposition as

$$\mathbf{v} = \mathbf{v}_0 + \mathbf{v}_1 \text{ where } \mathbf{v}_0 \in U \text{ and } \mathbf{v}_1 \in U^\perp.$$

The corresponding map
$$\pi \colon V \longrightarrow U$$
$$\mathbf{v} \longmapsto \pi(\mathbf{u}) := \mathbf{v}_0$$

is linear and surjective with kernel $U^\perp$ (check these!). It is called the *orthogonal projection* [Orthogonalprojektion] onto U.

Exercise 2.3.2 Show that the orthogonal projection π onto U is a projection, i.e., that $\pi \circ \pi = \pi$, and that it is uniquely determined by this condition together with the requirement that its image is U and that $(\mathbf{v} - \pi(\mathbf{v})) \perp U$ for all $\mathbf{v}$.

Exercise 2.3.3 Show that the orthogonal projections of an n-dimensional euclidean or unitary space V are precisely those endomorphisms of V that are represented w.r.t. suitable orthonormal basis by block diagonal matrices with blocks E_k $(0 \leqslant k \leqslant n)$ and $\mathbf{0}$.

Exercise 2.3.4 Consider two subspaces $U, W \subseteq V$ of a euclidean or unitary space V, with orthogonal projections π_U and π_W onto U and W, respectively. Under which conditions is $\pi_U \circ \pi_W$ an orthogonal projection? Under which conditions do π_U and π_W commute?

Observation 2.3.8 *An alternative characterisation of orthogonal projections is in terms of "best approximations". Let $U \subseteq V$ be a subspace of a finite-dimensional euclidean or unitary space V, π the orthogonal projection onto U. Then $\pi(\mathbf{v})$ is uniquely determined as that vector $\mathbf{u}$ in U for which $\|\mathbf{v} - \mathbf{u}\|$ is minimal.*

Proof.   Let $\pi(\mathbf{v}) = \mathbf{v}_0 \in U$, $\mathbf{v} - \pi(\mathbf{v}) = \mathbf{v}_1 \in U^\perp$. We need to show that for any $\mathbf{u} \in U$, $\|\mathbf{v} - \mathbf{u}\| \geqslant \|\mathbf{v}_1\|$ with equality only for $\mathbf{u} = \mathbf{v}_0$. Express $\mathbf{u}$ as $\mathbf{u} = \mathbf{v}_0 + \mathbf{w}$ for $\mathbf{w} \in U$. Then $\|\mathbf{v} - \mathbf{u}\|^2 = \langle \mathbf{v} - \mathbf{u}, \mathbf{v} - \mathbf{u} \rangle = \langle \mathbf{v} - \mathbf{v}_0 - \mathbf{w}, \mathbf{v} - \mathbf{v}_0 - \mathbf{w} \rangle = \langle \mathbf{v}_1 - \mathbf{w}, \mathbf{v}_1 - \mathbf{w} \rangle = \langle \mathbf{v}_1, \mathbf{v}_1 \rangle + \langle \mathbf{w}, \mathbf{w} \rangle$ by orthogonality of $\mathbf{v}_1 \in U^\perp$ and $\mathbf{w} \in U$. Hence $\|\mathbf{v} - \mathbf{u}\|^2 = \|\mathbf{v}_1\|^2 + \|\mathbf{w}\|^2 \geqslant \|\mathbf{v}_1\|^2$ with equality only for $\mathbf{w} = \mathbf{0}$.

$\square$

Exercise 2.3.5 Show the following identities in a finite-dimensional euclidean or unitary vector space V, for arbitrary subspaces $U, U_1, U_2 \subseteq V$:

   (i) $\left(U^\perp\right)^\perp = U$.
   (ii) $(U_1 \cap U_2)^\perp = U_1^\perp + U_2^\perp$.

Exercise 2.3.6 Show that any linear map $\eta\colon V \to \mathbb{F}$, i.e., any linear form (member of the dual space $V^* = \mathrm{Hom}(V, \mathbb{F})$), for a finite-dimensional euclidean or unitary space $(V, \langle .,. \rangle)$, can be represented in the form

$$\eta\colon \mathbf{v} \longmapsto \langle \mathbf{a}, \mathbf{v} \rangle$$

for a suitable $\mathbf{a} \in V$. This correspondence induces a canonical isomorphism between V^* and V for finite-dimensional euclidean or unitary spaces.

2.3.3 Orthogonal and unitary maps

Orthogonal and unitary maps are the automorphisms of euclidean and unitary spaces, respectively. So they are vector space automorphisms that are also isometries of the scalar product (or of the induced norm and metric).

Let us first consider isomorphisms between two euclidean or unitary spaces, i.e., isometric vector space isomorphisms.

Lemma 2.3.9 *Let $(V, \langle .,. \rangle^V)$ and $(W, \langle .,. \rangle^W)$ be two euclidean or two unitary spaces. Then the following are equivalent for $\varphi \in \mathrm{Hom}(V, W)$:*

(i) *φ is an isomorphism of the euclidean/unitary spaces, i.e., an isometric vector space isomorphism.*

(ii) *For some (any) orthonormal basis B of V, $\varphi(B) = (\varphi(\mathbf{b}))_{\mathbf{b} \in B}$ is an orthonormal basis of W.*

Proof. Recall that φ is a vector space isomorphism iff it transforms a basis of V into a basis of W. It remains to deal with the metric part, according to Definition 2.1.14. By definition, an isometry preserves the scalar product, and hence norms and orthogonality. It therefore maps orthonormal systems into orthonormal systems. Conversely, the values of the scalar product are uniquely determined by its values $\langle \mathbf{b}_i, \mathbf{b}_j \rangle$ on pairs of basis vectors. If, therefore, an orthonormal basis is mapped into an orthonormal basis, then $\delta_{ij} = \langle \varphi(\mathbf{b}_i), \varphi(\mathbf{b}_j) \rangle^W = \langle \mathbf{b}_i, \mathbf{b}_j \rangle^V = \delta_{ij}$ guarantee that φ preserves scalar products. Compare also Corollary 2.3.3.

$\square$

Corollary 2.3.10 *Every n-dimensional euclidean vector space V is isomorphic (isometrically isomorphic) to the n-dimensional standard euclidean space $\mathbb{R}^n$, and similarly for unitary V and $\mathbb{C}^n$.*

Proof.   Choose an orthonormal basis in the given euclidean or unitary space. The vector space isomorphism that associates these basis vectors with the standard basis vectors in $\mathbb{R}^n$ or $\mathbb{C}^n$, respectively, is an isometry.

<div align="right">□</div>

Definition 2.3.11 Let $(V, \langle ., . \rangle)$ be a euclidean or unitary space, respectively. A linear map $\varphi \in \operatorname{Hom}(V, V)$ is an *orthogonal map* [orthogonale Abbildung] or a *unitary map* [unitäre Abbildung], respectively, if

(i) it is bijective, i.e., a vector space automorphism of V, and

(ii) an isometry of $(V, \langle ., . \rangle)$, i.e., preserves the scalar product.

See Observation 2.1.15 for equivalent characterisations in terms of distance or norm preservation in place of (ii). Orthogonal and unitary maps also in particular preserve orthogonality, and in fact can also be characterised by the property in the lemma above, that they transform orthonormal bases into orthonormal bases.

Also note that while injectivity of φ follows from (ii), surjectivity only follows in the finite-dimensional case. An infinite-dimensional euclidean or unitary space with an orthonormal basis $B = (\mathbf{b}_0, \mathbf{b}_1, \mathbf{b}_2, \ldots)$ for instance, has an isometry $\varphi \colon \mathbf{b}_i \mapsto \mathbf{b}_{i+1}$ for which $\mathbf{b}_0 \notin \operatorname{image}(\varphi)$.

Example 2.3.12 The following are very common orthogonal maps in the euclidean spaces $\mathbb{R}^2$ and $\mathbb{R}^3$:

(i) all rotations about the origin in $\mathbb{R}^2$ or about an axis through the origin in $\mathbb{R}^3$.

(ii) reflections in a line through the origin in $\mathbb{R}^2$ or in a plane through the origin in $\mathbb{R}^3$.

Despite their name, orthogonal projections are not orthogonal maps! (Why?)

Exercise 2.3.7 Clearly the identity id_V is orthogonal (unitary). Show that the composition of any two orthogonal (unitary) maps of $(V. \langle ., . \rangle)$ is orthogonal (unitary), and that the inverse of an orthogonal (unitary) map is orthogonal (unitary). Hence, these sets of maps form groups with composition.

Definition 2.3.13 The subgroups of the group $\operatorname{Aut}(V)$ of all vector space automorphisms consisting of just the orthogonal or unitary maps of $(V, \langle ., . \rangle)$ are called the *orthogonal group* [orthogonale Gruppe] $\operatorname{O}(V)$ (in the euclidean case, over $\mathbb{R}$) and the *unitary group* $\operatorname{U}(V)$ [unitäre Gruppe] (in the unitary case, over $\mathbb{C}$).

Orthogonal and unitary matrices

For the unitary case recall that we write A^+ for the matrix obtained from A by simultaneous complex conjugation and transposition. If $A = (a_{ij})$, then the entries in A^+ are $a_{ij}^+ = \bar{a}_{ji}$.

Definition 2.3.14 A regular matrix $A \in \mathbb{R}^{(n,n)}$ is called *orthogonal* [orthogonal] iff $A^t = A^{-1}$ (its transpose is its inverse: $AA^t = A^t A = E_n$).
   A regular matrix $A \in \mathbb{C}^{(n,n)}$ is called *unitary* [unitär] iff $A^+ = A^{-1}$ (its adjoint is its inverse: $AA^+ = A^+ A = E_n$).

Exercise 2.3.8 Show that a matrix $A \in \mathbb{R}^{(n,n)}$ is orthogonal iff its column vectors form an orthonormal basis w.r.t. the standard scalar product, iff its row vectors form an orthonormal basis w.r.t. the standard scalar product.

Exercise 2.3.9 Show that any change of basis transformation between orthonormal bases in the standard unitary or euclidean spaces $\mathbb{C}^n$ or $\mathbb{R}^n$, respectively, are described by unitary or orthogonal matrices, respectively.

   Orthogonal (unitary) matrices are precisely the matrix representations of orthogonal (unitary) maps w.r.t. orthonormal bases. They are in particular the representations of orthogonal (unitary) maps in the standard euclidean (unitary) spaces w.r.t. their standard bases.

Proposition 2.3.15 *Let $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ be an orthonormal basis of the euclidean (unitary) space $(V, \langle ., . \rangle)$. Let $\varphi \in \mathrm{Aut}(V)$ a vector space automorphism of V, represented by the matrix $A = [\![\varphi]\!]_B^B \in \mathbb{R}^{(n,n)}$ (respectively in $\mathbb{C}^{(n,n)}$). Then the following are equivalent*

   (i) *A is orthogonal (unitary).*

   (ii) *φ is orthogonal (unitary).*

Proof.   We treat the unitary case (its specialisation to real matrices gives the euclidean analogue). Let $A = (a_{ij})$ be a regular matrix, $a_{ij}^+ = \bar{a}_{ji}$ be the entries in A^+ and recall that we write δ_{ij} to describe the entries in E_n, $\delta_{ij} = 0$ for $i \neq j$ and $\delta_{ii} = 1$.
   Now φ is an isometry iff $\varphi(B) = (\varphi(\mathbf{b}_1), \ldots, \varphi(\mathbf{b}_n))$ is orthonormal, i.e., iff $\langle \varphi(\mathbf{b}_i), \varphi(\mathbf{b}_j) \rangle = \delta_{ij}$. As $[\![\varphi]\!]_B^B = A$, $\varphi(\mathbf{b}_i) = \sum_k a_{ki}\mathbf{b}_k$ and therefore

$$
\begin{aligned}
\langle \varphi(\mathbf{b}_i), \varphi(\mathbf{b}_j) \rangle &= \langle \textstyle\sum_k a_{ki}\mathbf{b}_k, \sum_\ell a_{\ell j}\mathbf{b}_\ell \rangle \\
&= \textstyle\sum_{k,\ell} \bar{a}_{ki} a_{\ell j} \langle \mathbf{b}_k, \mathbf{b}_\ell \rangle \\
&= \textstyle\sum_{k,\ell} \bar{a}_{ki} a_{\ell j} \delta_{k\ell} = \sum_k \bar{a}_{ki} a_{kj} \\
&= \textstyle\sum_k a_{ik}^+ a_{kj}
\end{aligned}
$$

is the entry in position i, j of A^+A. Hence, $\varphi = \varphi_A$ is an isometry iff $\langle \varphi(\mathbf{b}_i), \varphi(\mathbf{b}_j) \rangle = \delta_{ij}$ iff $A^+A = E_n$.

$\square$

It follows that the orthogonal, repectively unitary, matrices form groups, in fact subgroups of the general linear groups $\mathrm{GL}_n(\mathbb{R})$ and $\mathrm{GL}_n(\mathbb{C})$, respectively.

Definition 2.3.16 The matrix groups of orthogonal and unitary matrices are denoted $\mathrm{O}(n) = \mathrm{O}(\mathbb{R}^n) \subseteq \mathrm{GL}_n(\mathbb{R})$ and $\mathrm{U}(n) = \mathrm{U}(\mathbb{C}^n) \subseteq \mathrm{GL}_n(\mathbb{C})$.

Isometries of $\mathbb{R}^n$

Isometries of the standard euclidean space $\mathbb{R}^n$ are necessarily bijective, and hence orthogonal maps. W.r.t. the standard basis (which is an orthonormal basis) an isometry φ is represented by an orthogonal matrix $A = A_\varphi \in \mathrm{O}(n)$. We now explore the geometric possibilities.

Consider any invariant subspace $U \subseteq V$ of φ. By injectivity, φ must map U onto U (dimension!), and as φ preserves orthogonality, $U^\perp$ is also an invariant subspace of φ. It follows that any invariant subspace of an orthogonal (or unitary) map gives rise to a decomposition into invariant subspaces that are orthogonal complements and hence allow a diagonal block decomposition w.r.t. orthogonal subspaces.

It remains to find invariant subspaces.

(A) If φ has an eigenvalue λ, then $|\lambda| = 1$ (as φ preserves norms). So $\lambda \in \{1, -1\}$. If $\mathbf{b}$ is a corresponding eigenvector, then $U = \mathrm{span}(\mathbf{b})$ and $U^\perp$ are invariant subspaces. For $\lambda = 1$, φ is trivial in the direction of $\mathbf{b}$. For $\lambda = -1$, φ operates as a reflection in the hyperplane $U^\perp$, combined with the isometric operation of φ inside $U^\perp$.

(B) Suppose that φ has no (real) eigenvalue. This means that the characteristic polynomial $p = p_\varphi = p_A$ has no real zeroes. We then know that n must be even, and that p must have an irreducible factor which is a product of two complex conjugate linear factors of p in $\mathbb{C}[X]$. Let this factor be $(X - \lambda)(X - \bar{\lambda})$ for $\lambda \in \mathbb{C} \setminus \mathbb{R}$, $\lambda = x + iy$, $y \neq 0$. Consider the map

$$\varphi^{\mathbb{C}} \colon \mathbb{C}^n \longrightarrow \mathbb{C}^n$$
$$\mathbf{v} \longmapsto A\mathbf{v}$$

As $A^+ = A^t$ (A has real entries), A is unitary, and hence $\varphi^{\mathbb{C}}$ is a unitary map. As λ and $\bar{\lambda}$ are eigenvalues, $|\lambda| = |\bar{\lambda}| = \sqrt{x^2 + y^2} = 1$ and $\lambda = e^{i\alpha}, \bar{\lambda} = e^{-i\alpha}$ for some $\alpha \in [0, 2\pi)$. Equivalently, $\lambda = \cos(\alpha) + i\sin(\alpha)$ and $\bar{\lambda} = \cos(\alpha) - i\sin(\alpha)$.

If $\mathbf{v} \in \mathbb{C}^n$ is an eigenvector with eigenvalue λ, then $A\mathbf{v} = \lambda\mathbf{v}$ implies (by complex conjugation) that $A\bar{\mathbf{v}} = \bar{\lambda}\bar{\mathbf{v}}$, i.e., $\bar{\mathbf{v}}$ is an eigenvector with eigenvalue $\bar{\lambda}$. But $\mathbf{b}_1 := 1/2(\mathbf{v} + \bar{\mathbf{v}})$ and $\mathbf{b}_2 := i/2(\mathbf{v} - \bar{\mathbf{v}})$ are both in $\mathbb{R}^n$. One checks that

$$\begin{array}{rclcl} \varphi(\mathbf{b}_1) & = & \varphi^{\mathbb{C}}(\mathbf{b}_1) & = & \cos(\alpha)\mathbf{b}_1 + \sin(\alpha)\mathbf{b}_2, \\ \varphi(\mathbf{b}_2) & = & \varphi^{\mathbb{C}}(\mathbf{b}_2) & = & -\sin(\alpha)\mathbf{b}_1 + \cos(\alpha)\mathbf{b}_2. \end{array}$$

So $U = \text{span}(\mathbf{b}_1, \mathbf{b}_2)$ is an invariant subspace of φ, and in restriction to U, φ operates as a rotation through angle α: one checks that, if $\mathbf{v}$ is normalised, then $(\mathbf{b}_1, \mathbf{b}_2)$ is orthonormal (this uses that A is unitary and that $\lambda \notin \mathbb{R}$).

Inductively decomposing V into orthogonal invariant subspaces we thus find the following characterisation of orthogonal maps.

Proposition 2.3.17 *For any isometry of the standard euclidean space $\mathbb{R}^n$ there is a decomposition of $\mathbb{R}^n$ into a direct sum of pairwise orthogonal subspaces U_i of dimensions 1 or 2 such that φ is the composition of*

(a) *reflections in hyperplanes $(U_i)^\perp$ for some i with $\dim(U_i) = 1$.*

(b) *rotations in the plane U_j for some j with $\dim(U_j) = 2$.*

Corollary 2.3.18 *For any orthogonal map of an n-dimensional euclidean space $(V, \langle., .\rangle)$, there are $k, \ell, r \in \mathbb{N}$ such that $k + \ell + 2r = n$, angles $\alpha_j \in [0, 2\pi)$ for $1 \leqslant j \leqslant r$, and an orthonormal basis B of V, such that φ is represented w.r.t. B by the block diagonal matrix*

$$A = \begin{pmatrix} E_k & & & & \\ & -E_\ell & & \mathbf{0} & \\ & & R_{\alpha_1} & & \\ & \mathbf{0} & & \ddots & \\ & & & & R_{\alpha_r} \end{pmatrix},$$

where R_α is the 2×2 rotation matrix $R_\alpha = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{pmatrix}$.

Corollary 2.3.19 *Any orthogonal matrix $A \in \mathbb{R}^{(n,n)}$ is similar to a matrix of the above kind via a similarity transformation induced by an orthogonal matrix C.*

A familiar geometric distinction between a rotation on the one hand, and a reflection in a line of $\mathbb{R}^2$ or in a plane of $\mathbb{R}^3$ on the other hand, concerns their behaviour w.r.t. *orientation*[Orientierung]. An orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $\mathbb{R}^n$ is said to be *positively oriented* if

$$\det(\mathbf{b}_1, \ldots, \mathbf{b}_n) = 1,$$

negatively oriented if this determinant is -1.

Exercise 2.3.10 Show that

  (i) $|A| \in \{1, -1\}$ for any orthogonal matrix $A \in \mathbb{R}^{(n,n)}$.
  (ii) $\det(\mathbf{b}_1, \ldots, \mathbf{b}_n) \in \{-1, 1\}$ for any orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of the standard euclidean space $\mathbb{R}^n$.

Definition 2.3.20 An orthogonal map of the standard space $\mathbb{R}^n$ is *orientation preserving* [orientierungserhaltend] if it maps the standard basis (or any other positively oriented orthonormal basis) into a positively oriented orthonormal basis.

An orthogonal matrix $A \in \mathrm{O}(n) \subseteq \mathrm{GL}_n(\mathbb{R})$ is called a special orthogonal matrix if $|A| = 1$. The special orthogonal matrices form a subgroup of $\mathrm{O}(n)$, called the *special orthogonal group*[spezielle orthogonale Gruppe], denoted $\mathrm{SO}(n)$.

Exercise 2.3.11 Show that the special orthogonal matrices are precisely the representations of the orientation preserving orthogonal maps w.r.t. orthonormal bases. (Compare Proposition 2.3.15.)

Exercise 2.3.12 Classify all orthogonal maps of $\mathbb{R}^2$ and $\mathbb{R}^3$ according to Proposition 2.3.17 and orientation preservation.

2.4 Endomorphisms in euclidean or unitary spaces

In this section we study the diagonalisability of endomorphisms of a finite-dimensional euclidean or unitary space. We are interested in matrix representations of φ w.r.t. orthonormal bases and corresponding similarity transformations. Recall that a similarity transformation $A \mapsto \hat{A}$ that represents a change from one orthonormal basis of V to another is based on an orthogonal or unitary matrix C for which $\hat{A} = CAC^{-1}$.

2.4.1 The adjoint map

Definition 2.4.1 Let $(V, \langle ., . \rangle)$ be a euclidean or unitary space, $\varphi \colon V \to V$ an endomorphism of V. The endomorphism $\varphi^+ \colon V \to V$ is called an *adjoint* of φ [adjungierte Abbildung] if for all $\mathbf{v}, \mathbf{w} \in V$:

$$\langle \varphi(\mathbf{v}), \mathbf{w} \rangle = \langle \mathbf{v}, \varphi^+(\mathbf{w}) \rangle.$$

$\varphi \in \mathrm{Hom}(V, V)$ is called *self-adjoint* (or, in the real case, also *symmetric*) [selbstadjungiert, reell: symmetrisch] if φ is its own adjoint.

Observation 2.4.2 *For any orthogonal or unitary φ, its inverse is its adjoint.*

Proof. Let $\varphi \in \mathrm{Hom}(V, V)$ be an isometry, $\mathbf{v}, \mathbf{w} \in V$. Let $\mathbf{w}' := \varphi^{-1}(\mathbf{w})$, then the isometry condition that $\langle \varphi(\mathbf{v}), \varphi(\mathbf{w}') \rangle = \langle \mathbf{v}, \mathbf{w}' \rangle$ implies that $\langle \varphi(\mathbf{v}), \mathbf{w} \rangle = \langle \mathbf{v}, \varphi^{-1}(\mathbf{w}) \rangle$.

$\square$

Example 2.4.3 Let V be a finite-dimensional euclidean or unitary space, $U \subseteq V$ a subspace, and consider the orthogonal projection π of V onto U. For an adjoint π^+ we need
   (i) $\mathrm{image}(\pi^+) \subseteq (U^\perp)^\perp = U$.
       For $\mathbf{v} \in U^\perp$ and any $\mathbf{w} \in V$: $\langle \mathbf{v}, \pi^+(\mathbf{w}) \rangle = \langle \pi(\mathbf{v}), \mathbf{w} \rangle = \langle \mathbf{0}, \mathbf{w} \rangle = 0$.
  (ii) $U^\perp \subseteq \ker(\pi^+)$.
       For $\mathbf{w} \in U^\perp$ and any $\mathbf{v} \in V$: $\langle \mathbf{v}, \pi^+(\mathbf{w}) \rangle = \langle \pi(\mathbf{v}), \mathbf{w} \rangle = 0$.
 (iii) $\pi^+ \circ \pi^+ = \pi^+$ (i.e., π^+ needs to be a projection too).
       Fix $\mathbf{w} \in W$; then $\langle \mathbf{v}, \pi^+(\pi^+(\mathbf{w})) \rangle = \langle \pi(\mathbf{v}), \pi^+(\mathbf{w}) \rangle = \langle \pi(\pi(\mathbf{v})), \mathbf{w} \rangle = \langle \pi(\mathbf{v}), \mathbf{w} \rangle = \langle \mathbf{v}, \pi^+(\mathbf{w}) \rangle$ for all $\mathbf{v} \in V$ implies that $\pi^+(\pi^+(\mathbf{w})) = \pi^+(\mathbf{w})$.
One checks that indeed $\pi^+ = \pi$ is an adjoint, and that orthogonal projections are self-adjoint.

Exercise 2.4.1 Show that φ^+ is an adjoint of φ iff φ is an adjoint of φ^+.

Recall Definition 2.1.8, of self-adjoint (or hermitian) complex matrices $A = A^+ \in \mathbb{C}^{(n,n)}$; their real counterparts, or specialisations, are the symmetric real matrices $A = A^t \in \mathbb{R}^{(n,n)}$ (note that $A^+ = A^t$ for real A, whence a real self-adjoint matrix *is* symmetric).

Proposition 2.4.4 *For finite-dimensional V: any $\varphi \in \mathrm{Hom}(V, V)$ possesses an adjoint and this adjoint is uniquely determined.*

In terms of an orthonormal basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ of V, the adjoint φ^+ of φ is represented by the adjoint matrix of the matrix that represents φ:

$$\llbracket \varphi^+ \rrbracket_B^B = (\llbracket \varphi \rrbracket_B^B)^+.$$

W.r.t. orthonormal bases, self-adjoint/symmetric endomorphisms are represented by self-adjoint (hermitian)/symmetric matrices.

Proof.   We first show uniqueness. Let ψ, ψ' both be adjoint maps for φ. Then

$$\langle \varphi(\mathbf{v}), \mathbf{w} \rangle = \langle \mathbf{v}, \psi(\mathbf{w}) \rangle = \langle \mathbf{v}, \psi'(\mathbf{w}) \rangle \text{ for all } \mathbf{v} \in V$$

implies that $(\psi(\mathbf{w}) - \psi'(\mathbf{w})) \perp V$, hence $\psi(\mathbf{w}) - \psi'(\mathbf{w}) = \mathbf{0}$ and $\psi = \psi'$.

For existence, let $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ be an orthonormal basis of V, $A = \llbracket \varphi \rrbracket_B^B$ the matrix representing φ. Then the endomorphism ψ whose representation w.r.t. B is A^+, is easily seen to satisfy the conditions for an adjoint. Consider $\mathbf{v} = \sum_i \lambda_i \mathbf{b}_i$ and $\mathbf{w} = \sum_j \mu_j \mathbf{b}_j$. Then $\varphi(\mathbf{v}) = \sum_{ik} a_{ki} \lambda_i \mathbf{b}_k$ and $\psi(\mathbf{w}) = \sum_{j\ell} a_{\ell j}^+ \mu_j \mathbf{b}_\ell = \sum_{j\ell} \bar{a}_{j\ell} \mu_j \mathbf{b}_\ell$. Therefore

$$\langle \varphi(\mathbf{v}), \mathbf{w} \rangle = \langle \sum_{ik} a_{ki} \lambda_i \mathbf{b}_k, \sum_j \mu_j \mathbf{b}_j \rangle = \sum_{ikj} \bar{\lambda}_i \bar{a}_{ki} \mu_j \delta_{kj} = \sum_{ij} \bar{\lambda}_i \bar{a}_{ji} \mu_j$$

and

$$\langle \mathbf{v}, \psi(\mathbf{w}) \rangle = \langle \sum_i \lambda_i \mathbf{b}_i, \sum_{j\ell} \bar{a}_{j\ell} \mu_j \mathbf{b}_\ell \rangle = \sum_{ij\ell} \bar{\lambda}_i \bar{a}_{j\ell} \mu_j \delta_{i\ell} = \sum_{ij} \bar{\lambda}_i \bar{a}_{ji} \mu_j.$$

So the equality $\langle \varphi(\mathbf{v}), \mathbf{w} \rangle = \langle \mathbf{v}, \psi(\mathbf{w}) \rangle$ proves that ψ is an adjoint (and hence *the* adjoint of φ), as well as the claim about its matrix representation.

$\square$

2.4.2   Diagonalisation of self-adjoint maps and matrices

Proposition 2.4.5 *Let $\varphi = \varphi^+$ be a self-adjoint endomorphism of the finite-dimensional unitary $\mathbb{C}$-vector space $(V, \langle ., . \rangle)$. Then all eigenvalues of φ are real, and φ can be diagonalised w.r.t. an orthonormal basis of V.*

Similarly, any self-adjoint (i.e., symmetric) endomorphism of a finite-dimensional euclidean $\mathbb{R}$-vector space is diagonalisable w.r.t. an orthonormal basis.

Proof.   Let $\mathbf{v}$ be an eigenvector with eigenvalue $\lambda \in \mathbb{C}$ of φ. Then $\langle \varphi(\mathbf{v}), \mathbf{v} \rangle = \bar{\lambda} \|\mathbf{v}\|^2 = \langle \mathbf{v}, \varphi(\mathbf{v}) \rangle = \lambda \|\mathbf{v}\|^2$ shows that $\lambda = \bar{\lambda}$, whence $\lambda \in \mathbb{R}$.

The characteristic polynomial p_φ always splits into linear factors in $\mathbb{C}[X]$. As we now see that all these linear factors are of the form $(X - \lambda)$ for $\lambda \in \mathbb{R}$, p_φ splits into linear factors in $\mathbb{R}[X]$ also in the symmetric (real self-adjoint) case.

We claim that (in the euclidean as well as in the unitary case) the orthogonal complement of an invariant subspace of φ is again an invariant subspace. Let $U \subseteq V$ be invariant under φ, and let $\mathbf{v} \in U^\perp$. Then

$$\langle \varphi(\mathbf{v}), \mathbf{u} \rangle = \langle \mathbf{v}, \varphi(\mathbf{u}) \rangle = 0 \text{ for all } \mathbf{u} \in U,$$

since $\varphi(\mathbf{u}) \in U$ and $\mathbf{v} \perp U$. Hence also $\varphi(\mathbf{v}) \in U^\perp$.

If $\mathbf{v}$ is an eigenvector of φ, therefore, $V = \mathrm{span}(\mathbf{v}) \oplus (\mathrm{span}(\mathbf{v}))^\perp$ is a decomposition into orthogonal complements that are invariant. W.l.o.g. $\|\mathbf{v}\| = 1$ and we may choose $\mathbf{b}_1 := \mathbf{v}$ as the first basis vector for the desired orthonormal basis that diagonalises φ. Putting $V' := (\mathrm{span}(\mathbf{v}))^\perp$ and letting φ' be the restriction of φ to V', we see that φ' is again self-adjoint, and that we can proceed inductively to select new normalised eigenvectors, each orthogonal on the subspace spanned by the previous. This gives an orthonormal basis w.r.t. which φ is represented by a *real* diagonal matrix (which in particular is self-adjoint of course).

$\square$

Corollary 2.4.6 *Any self-adjoint (hermitian) matrix $A \in \mathbb{C}^{(n,n)}$ is similar to a real diagonal matrix, by means of a similarity transformation based on a unitary matrix: for suitable unitary C, the matrix CAC^{-1} is diagonal and in $\mathbb{R}^{(n,n)}$.*

Similarly, any symmetric $A \in \mathbb{R}^{(n,n)}$ is similar to a diagonal matrix, by means of a similarity transformation based on an orthogonal matrix.

Proof.   Let $A = A^+$ be given. Consider $\mathbb{C}^n$ with its standard scalar product as a unitary space, and let $\varphi = \varphi_A$ be the endomorphism that corresponds to A w.r.t. the standard basis B_0. Then φ is self-adjoint, since A is self-adjoint and since the standard basis is orthonormal. By the proposition, there is another orthonormal basis B of $\mathbb{C}^n$ such that $[\![\varphi]\!]_B^B$ is diagonal with real entries. The basis transformation between the standard basis and the new B is unitary, as both bases are orthonormal. Hence

$$[\![\varphi]\!]_B^B = CAC^{-1},$$

where $C = [\![\text{id}]\!]_B^{B_0}$ is the unitary matrix whose column vectors consist of the coefficients of the standard basis vectors w.r.t. the new orthonormal basis B, $[\![\mathbf{e}_j]\!]_B$.

In the real case, the corresponding basis transformation similarly is described by an orthogonal matrix C.

$\square$

2.4.3 Normal maps and matrices

Propositions 2.4.5 and 2.3.17 both essentially used the fact that the orthogonal complement of the invariant subspace spanned by an eigenvector was again an invariant subspace. The following classes of *normal* endomorphisms and matrices provide a common generalisation that includes unitary/orthogonal as well as self-adjoint/symmetric maps or matrices.

Definition 2.4.7 (i) An endomorphism $\varphi \in \text{Hom}(V, V)$ of a euclidean or unitary space $(V, \langle ., . \rangle)$ is called *normal* [normal] if φ and its adjoint φ^+ commute:

$$\varphi^+ \circ \varphi = \varphi \circ \varphi^+.$$

(ii) A matrix $A \in \mathbb{C}^{(n,n)}$ is called *normal* if it commutes with its adjoint:

$$A^+ A = A A^+.$$

Lemma 2.4.8 *Let $(V, \langle ., . \rangle)$ be euclidean or unitary with orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$. Let $\varphi \in \text{Hom}(V, V)$ be represented w.r.t. B by the matrix $A = [\![\varphi]\!]_B^B$ Then the following are equivalent:*

(i) *φ is normal.*

(ii) *A is normal.*

Proof. Straightforward calculation, similar to the corresponding part in the proof of Proposition 2.4.4.

$\square$

The following combines observations about the correspondence of unitary/orthogonal or self-adjoint/symmetric maps with the corresponding matrices. Also observe that trivially, A commutes with A and, for regular A, also with A^{-1}.

Observation 2.4.9 *The following classes of matrices in particular are normal:*

(i) *unitary (in the euclidean case, for $\mathbb{F} = \mathbb{R}$: orthogonal) matrices.*

(ii) *self-adjoint (in the euclidean case, for $\mathbb{F} = \mathbb{R}$: symmetric) matrices.*

Analogously, the following classes of endomorphisms are normal:

(i) *unitary and orthogonal maps, respectively.*

(ii) *self-adjoint and symmetric maps, respectively.*

Diagonalisation over $\mathbb{C}$

Theorem 2.4.10 *Let $(V, \langle ., . \rangle)$ be a finite-dimensional unitary space, $\varphi \in \mathrm{Hom}(V, V)$ normal. Then φ is diagonalisable w.r.t. an orthonormal basis for V.*

Proof. Note that, as we are working over $\mathbb{C}$, φ does have an eigenvalue λ. We claim that the (non-trivial) eigenspace $V_\lambda \subseteq V$ of φ is an invariant subspace also for φ^+. Let $\mathbf{v} \in V_\lambda$, i.e., $\varphi(\mathbf{v}) = \lambda \mathbf{v}$. We need to show that also $\varphi^+(\mathbf{v}) \in V_\lambda$. By normality, $\varphi(\varphi^+(\mathbf{v})) = \varphi^+(\varphi(\mathbf{v})) = \varphi^+(\lambda \mathbf{v}) = \lambda \varphi^+(\mathbf{v})$ shows that indeed $\varphi^+(\mathbf{v}) \in V_\lambda$.

Therefore, φ^+ will also have an eigenvector $\mathbf{w} \in V_\lambda$ (look at the restriction of φ^+ to V_λ). Any such $\mathbf{w}$ therefore is a simultaneous eigenvector of both φ and φ^+: $\varphi(\mathbf{w}) = \lambda \mathbf{w}$ and $\varphi^+(\mathbf{w}) = \gamma \mathbf{w}$. As $\gamma \|\mathbf{w}\|^2 = \langle \mathbf{w}, \varphi^+(\mathbf{w}) \rangle = \langle \varphi(\mathbf{w}), \mathbf{w} \rangle = \lambda \|\mathbf{w}\|^2$, it follows that $\gamma = \lambda$.

We next argue that for any such simultaneous eigenvector $\mathbf{w}$, both $W = \mathrm{span}(\mathbf{w})$ and its orthogonal complement $V' := W^\perp$ are invariant subspaces for φ and φ^+. This is clear for W; so consider $W^\perp$, and for instance invariance under φ. For $\mathbf{a} \in W^\perp$ we need to show that $\varphi(\mathbf{a}) \in W^\perp$, i.e., that $\varphi(\mathbf{a}) \perp W$, or that $\langle \varphi(\mathbf{a}), \mathbf{w} \rangle = 0$. But $\langle \varphi(\mathbf{a}), \mathbf{w} \rangle = \langle \mathbf{a}, \varphi^+(\mathbf{w}) \rangle = \gamma \langle \mathbf{a}, \mathbf{w} \rangle = 0$ since $\mathbf{a} \perp \mathbf{w}$. Similarly, for invariance under φ^+, use that $\langle \varphi^+(\mathbf{a}), \mathbf{w} \rangle = \langle \mathbf{a}, \varphi(\mathbf{w}) \rangle = \lambda \langle \mathbf{a}, \mathbf{w} \rangle$.

Hence, the desired orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ is found inductively, as follows. For a first basis vector $\mathbf{b}_1$ pick any normalised simultaneous eigenvector $\mathbf{w}$ of φ and φ^+. Let $W = \mathrm{span}(\mathbf{b}_1)$ and $V' := W^\perp$ so that $V = W \oplus V'$, $W \perp V'$ are orthogonal complements.

As W and $W^\perp$ are both invariant under φ and φ^+, the restrictions of φ and φ^+ to V' are adjoints of each other, and normal. Any extension of $\mathbf{b}_1$ to an (orthonormal) basis of V by an (orthonormal) basis of V' will lead to a

representation of φ in the block diagonal form

$$A = \begin{pmatrix} \lambda & 0 & \cdots & 0 \\ 0 & & & \\ \vdots & & A' & \\ 0 & & & \end{pmatrix},$$

where A' is the representation of the restriction of φ to V'. Iteration of this selection procedure, in subspaces of decreasing dimension, produces the desired orthonormal basis.

□

Corollary 2.4.11 *Any normal matrix* $A \in \mathbb{C}^{(n,n)}$ *is similar to a diagonal matrix by means of a similarity transformation based on a unitary matrix: for suitable unitary* C, *the matrix* CAC^{-1} *diagonal.*

Proof.   Let A be given. Consider $\mathbb{C}^n$ as a unitary space with the standard scalar product. Let $\varphi = \varphi_A$ be the endomorphism that corresponds to A w.r.t. the standard basis B_0. Then φ is normal, since A is normal and since the standard basis is orthonormal. By the theorem, there is another orthonormal basis B of $\mathbb{C}^n$ such that $[\![\varphi]\!]_B^B$ is diagonal. The basis transformation between the standard basis and the new B is unitary, as both bases are orthonormal. Hence $[\![\varphi]\!]_B^B = CAC^{-1}$ is diagonal and C is unitary.

□

Similarly, one obtains the following.

Corollary 2.4.12 *Any unitary matrix* $A \in \mathbb{C}^{(n,n)}$ *is similar to a diagonal matrix, by means of a similarity transformation based on a unitary matrix: for suitable unitary* C, *the matrix* CAC^{-1} *is diagonal, with diagonal entries of the form* $\lambda_j = e^{i\alpha_j}$ *for suitable* $\alpha_j \in [0, 2\pi)$.

Chapter 3

Bilinear and Quadratic Forms

In this chapter we restrict attention to finite-dimensional $\mathbb{R}$-vector spaces and euclidean spaces. Unless stated otherwise, let $(V, \langle ., . \rangle)$ be an n-dimensional euclidean space.

3.1 Matrix representations of bilinear forms

Recall from section 2.1.3 how we represent a bilinear form $\sigma \colon V \times V \to \mathbb{R}$ over an n-dimensional $\mathbb{R}$-vector space V w.r.t. a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ by the matrix $A = [\![\sigma]\!]^B = (a_{ij})$ with entries

$$a_{ij} = \sigma(\mathbf{b}_i, \mathbf{b}_j).$$

If $\mathbf{v} = \sum_i \lambda_i \mathbf{b}_i$ and $\mathbf{w} = \sum_j \mu_j \mathbf{b}_j$, then σ is evaluated on these arguments according to

$$\sigma(\mathbf{v}, \mathbf{w}) = \sum_{ij} \lambda_i a_{ij} \mu_j = (\lambda_1, \ldots, \lambda_n) \, A \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_n \end{pmatrix} = [\![\mathbf{v}]\!]_B^t [\![\sigma]\!]^B [\![\mathbf{w}]\!]_B.$$

If $B' = (\mathbf{b}'_1, \ldots, \mathbf{b}'_n)$ is another basis, and if the corresponding basis transformation from B' to B is described by the matrix $C = (c_{ij}) = [\![\mathrm{id}_V]\!]_B^{B'}$ such that $\mathbf{b}'_i = \sum_k c_{ki} \mathbf{b}_k$, then the matrix $A' = [\![\sigma]\!]^{B'} = (a'_{ij})$ that represents σ

w.r.t. B' has entries

$$
\begin{aligned}
a'_{ij} &= \sigma(\mathbf{b}'_i, \mathbf{b}'_j) = \sigma(\textstyle\sum_k c_{ki}\mathbf{b}_k, \sum_\ell c_{\ell j}\mathbf{b}_\ell) \\
&= \textstyle\sum_{k\ell} c_{ki}c_{\ell j}\sigma(\mathbf{b}_k, \mathbf{b}_\ell) = \sum_{k\ell} c_{ki}a_{k\ell}c_{\ell j} \\
&= \left(C^t A C\right)_{ij},
\end{aligned}
$$

$$
\text{or } \ [\![\sigma]\!]^{B'} = ([\![\mathrm{id}_V]\!]^{B'}_B)^t [\![\sigma]\!]^B [\![\mathrm{id}_V]\!]^{B'}_B.
$$

Observation 3.1.1 *As representations of bilinear forms matrices $A \in \mathbb{R}^{(n,n)}$ transform under changes of basis according to*

$$
A \longmapsto C^t A C,
$$

for $C \in \mathrm{GL}_n(\mathbb{R})$ (regular matrices).

Note how this differs from the transformation pattern for matrices as representations of endomorphisms, which is $A \mapsto C^{-1}AC$. If C, however, is orthogonal ($C \in \mathrm{O}(n)$, cf. Definitions 2.3.14 and 2.3.16: $C^t = C^{-1}$), then these two transformations coincide.

Exercise 3.1.1 Show that the relation $\approx$ on $\mathbb{R}^{(n,n)}$ defined as $A \approx A'$ iff $A' = C^t A C$ for some $C \in \mathrm{GL}_n(\mathbb{R})$ is an equivalence relation. What are sufficient criteria for $A \not\approx A'$?

3.2 Simultaneous diagonalisation

We have seen in the previous chapter that certain endomorphisms of a euclidean space $(V, \langle ., . \rangle)$ can be diagonalised w.r.t. an orthonormal basis. Another way of saying that is that the underlying scalar product $\langle ., . \rangle$ and the given endomorphism $\varphi \colon V \to V$ can simultaneously be diagonalised: an orthonormal basis (by definition) leads to a diagonal matrix for the representation of the scalar product, namely the unit matrix E_n; and at the same time provides a diagonal matrix for the representation of φ. In particular that was seen to be possible for self-adjoint endomorphisms, see Proposition 2.4.5.

A similar question can be asked about the given scalar product and a second symmetric bilinear form σ on V: when is it possible to find one basis that simultaneously represents both, $\langle ., . \rangle$ and σ by diagonal matrices? Equivalently, when can σ be represented by a diagonal matrix w.r.t. an orthonormal basis of $(V, \langle ., . \rangle)$? These questions in fact reduce to what we already know about self-adjoint diagonalisation of endomorphism.

3.2.1   Symmetric bilinear forms vs. self-adjoint maps

Let $\varphi \in \mathrm{Hom}(V, V)$, $(V, \langle ., . \rangle)$ euclidean of dimension n. From φ and $\langle ., . \rangle$ we may construct a new bilinear form

$$
\begin{aligned}
\langle ., . \rangle_\varphi \colon V \times V &\longrightarrow \mathbb{R} \\
(\mathbf{v}, \mathbf{w}) &\longmapsto \langle \mathbf{v}, \mathbf{w} \rangle_\varphi := \langle \mathbf{v}, \varphi(\mathbf{w}) \rangle.
\end{aligned}
$$

It is easily checked that $\langle ., . \rangle_\varphi$ is bilinear.

In terms of matrix representations over a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$, one checks that if $\langle ., . \rangle$ is represented by the matrix $G = (g_{ij})$ with $g_{ij} = \langle \mathbf{b}_i, \mathbf{b}_j \rangle$, and φ is represented by the matrix $A = [\![\varphi]\!]_B^B$, then $\langle ., . \rangle_\varphi$ is represented by the matrix GA:

$$
\begin{aligned}
\langle \mathbf{b}_i, \mathbf{b}_j \rangle_\varphi &= \langle \mathbf{b}_i, \varphi(\mathbf{b}_j) \rangle = \langle \mathbf{b}_i, \textstyle\sum_k a_{kj} \mathbf{b}_k \rangle \\
&= \textstyle\sum_k a_{kj} \langle \mathbf{b}_i, \mathbf{b}_k \rangle = \sum_k a_{kj} g_{ik} = \sum_k g_{ik} a_{kj} \\
&= (GA)_{ij}.
\end{aligned}
$$

If B is an orthonormal basis w.r.t. $\langle ., . \rangle$, then $G = E_n$, and $\langle ., . \rangle_\varphi$ is represented by the matrix A. Note that the same matrix A occurs as a representation of an endomorphism and as a representation of an induced bilinear form.

Lemma 3.2.1 *Any bilinear form $\sigma \colon V \times V \to \mathbb{R}$ over a finite-dimensional euclidean $(V, \langle ., . \rangle)$ is of the form $\sigma = \langle ., . \rangle_\varphi$ for suitable $\varphi \in \mathrm{Hom}(V, V)$.*

The endomorphism φ is uniquely determined by σ (and vice versa).

Proof.   Let $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ be an orthonormal basis of V and let $[\![\sigma]\!]^B = A$ be represented by the matrix $A = (a_{ij})$ where $a_{ij} = \sigma(\mathbf{b}_i, \mathbf{b}_j)$. Let φ be the endomorphism whose matrix representation w.r.t. B is A, such that also $[\![\varphi]\!]_B^B = [\![\sigma]\!]^B = A$. Then $\langle ., . \rangle_\varphi$ is represented by the matrix A, and hence $\sigma = \langle ., . \rangle_\varphi$.

For uniqueness: $\langle \mathbf{v}, \mathbf{w} \rangle_\varphi = \langle \mathbf{v}, \mathbf{w} \rangle_\psi$ iff $\langle \mathbf{v}, \varphi(\mathbf{w}) - \psi(\mathbf{w}) \rangle = 0$. So $\langle \mathbf{v}, \mathbf{w} \rangle_\varphi = \langle \mathbf{v}, \mathbf{w} \rangle_\psi$ for all $\mathbf{v}$, implies that $\varphi(\mathbf{w}) = \psi(\mathbf{w})$.

$\square$

Recall the definitions of symmetry and positive definiteness of a bilinear form from Definition 2.1.2.

Definition 3.2.2 A bilinear form $\sigma \colon V \times V \to \mathbb{R}$ is *non-degenerate* [nicht ausgeartet] if for all $\mathbf{v} \in V$: $\sigma(\mathbf{v}, \mathbf{w}) = 0$ for all $\mathbf{w} \in V$ implies $\mathbf{v} = \mathbf{0}$.

Note that positive definiteness implies non-degeneracy, but non-degeneracy is strictly weaker and in particular makes sense also for forms that produce negative values for some $\sigma(\mathbf{v}, \mathbf{v})$.

Lemma 3.2.3 *Over a finite dimensional euclidean space $(V, \langle ., .\rangle)$ and for any $\varphi \in \mathrm{Hom}(V, V)$:*

(i) *$\langle ., .\rangle_\varphi$ is non-degenerate iff φ is regular ($\mathrm{rank}(\varphi) = \dim(V)$).*

(ii) *$\langle ., .\rangle_\varphi$ is symmetric iff φ is self-adjoint.*

Proof.   For (i): $\langle \mathbf{v}, \mathbf{w}\rangle_\varphi = 0$ iff $\langle \mathbf{v}, \varphi(\mathbf{w})\rangle = 0$ iff $\mathbf{v} \perp \varphi(\mathbf{w})$. Hence $\langle \mathbf{v}, \mathbf{w}\rangle_\varphi = 0$ for all $\mathbf{w}$ iff $\mathbf{v} \perp \mathrm{image}(\varphi)$.

If φ is regular, then $\mathrm{image}(\varphi) = V$ and $\mathbf{v} \perp \mathrm{image}(\varphi)$ implies $\mathbf{v} = \mathbf{0}$.

Otherwise, any $\mathbf{v} \in \mathrm{image}(\varphi)^\perp \setminus \{\mathbf{0}\}$ demonstrates that $\langle ., .\rangle_\varphi$ is degenerate.

For (ii), recall that the adjoint φ^+ of φ is defined by the condition that

$$\langle \mathbf{v}, \varphi^+(\mathbf{w})\rangle = \langle \varphi(\mathbf{v}), \mathbf{w}\rangle,$$

which now gives

$$\langle \mathbf{w}, \mathbf{v}\rangle_\varphi = \langle \mathbf{w}, \varphi(\mathbf{v})\rangle = \langle \varphi(\mathbf{v}), \mathbf{w}\rangle = \langle \mathbf{v}, \varphi^+(\mathbf{w})\rangle = \langle \mathbf{v}, \mathbf{w}\rangle_{\varphi^+}.$$

Uniqueness in Lemma 3.2.1 shows that $\langle ., .\rangle_\varphi$ is symmetric iff $\varphi = \varphi^+$. Alternatively, one may also consider the matrix representations w.r.t. an orthonormal basis and use the fact that symmetry of a bilinear form corresponds to symmetry of its matrix representation (w.r.t. arbitrary bases); and that self-adjoint endomorphisms are represented by self-adjoint (i.e., symmetric) matrices w.r.t. orthonormal bases.

□

Exercise 3.2.1 Show that φ is regular iff φ^+ is regular.

3.2.2   Principal axes

Definition 3.2.4 Let $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ be a basis of V, and σ a bilinear form which is represented by a diagonal matrix w.r.t. basis B: $\sigma(\mathbf{b}_i, \mathbf{b}_j) = \lambda_i \delta_{ij}$ for suitable $\lambda_j \in \mathbb{R}$. Then the one-dimensional subspaces spanned by basis vectors $\mathbf{b}_i$ are called *principal axes* [Hauptachsen] for σ.

Note that principal axes of σ are characterised by pairwise 'orthogonality w.r.t. σ'. It turns out that any symmetric bilinear form σ has such principal axes. Moreover these can always be chosen orthogonal w.r.t. a given scalar product in V.

Theorem 3.2.5 *Any symmetric bilinear form σ over a finite-dimensional euclidean space $(V, \langle ., . \rangle)$ can be diagonalised w.r.t. an orthonormal basis of $(V, \langle ., . \rangle)$. I.e., there is an orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $(V, \langle ., . \rangle)$ such that σ is represented by a diagonal matrix w.r.t. B: $\sigma(\mathbf{b}_i, \mathbf{b}_j) = \lambda_i \delta_{ij}$ for suitable λ_i.*

Proof.   The claim reduces to our results on diagonalisation of self-adjoint (symmetric) endomorphisms, if we go via a presentation of σ as $\sigma = \langle ., . \rangle_\varphi$ for a self-adjoint φ.

By Proposition 2.4.5, we have an orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $(V, \langle ., . \rangle)$ such that the endomorphism φ is represented by a diagonal matrix $D = [\![\varphi]\!]_B^B$:

$$D = \begin{pmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{pmatrix}.$$

Since $\langle ., . \rangle$ is represented by E_n w.r.t. B, $\sigma = \langle ., . \rangle_\varphi$ is also represented by D w.r.t. B.

$\square$

If we consider diagonalisation w.r.t. bases that are not necessarily orthonormal, then the actual diagonal entries λ_i are clearly not determined by σ. A rescaling of basis vector $\mathbf{b}_i$ to $\mathbf{b}_i' = \alpha \mathbf{b}_i$, $\alpha \neq 0$, changes λ_i into $\alpha^2 \lambda_i$. However, the sign distribution of the λ_i is an invariant.

Theorem 3.2.6 (Sylvester) *Let σ be a symmetric bilinear form over the finite-dimensional $\mathbb{R}$-vector space V. For any basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of V such that σ is represented by a diagonal matrix with entries $\sigma(\mathbf{b}_i, \mathbf{b}_j) = \lambda_i \delta_{ij}$ for suitable λ_i, the numbers of λ_i that are positive, negative, and equal to 0, respectively, are uniquely determined by σ (independent of B).*

Proof.   Let s be the number of positive λ_i, t the number of negative λ_i, and consequently $d := n - s - t$ the number of i for which $\lambda_i = 0$. W.l.o.g.

assume that $\sigma(\mathbf{b}_i, \mathbf{b}_i) > 0$ for $i = 1, \ldots, s$; $\sigma(\mathbf{b}_i, \mathbf{b}_i) < 0$ for $i = s+1, \ldots, s+t$; $\sigma(\mathbf{b}_i, \mathbf{b}_i) = 0$ for $i = s + t + 1, \ldots, n$.

One checks that d, the number of $\lambda_i = 0$, is the dimension of the subspace

$$U_0 = \{\mathbf{u} \in V : \sigma(\mathbf{u}, \mathbf{v}) = 0 \text{ for all } \mathbf{v} \in V\},$$

hence independent of B.

We claim that the number s of positive λ_i is the dimension of a maximal subspace $U_+ \subseteq V$ on which the restriction of σ is positive definite. Clearly V has an s-dimensional subspace on which σ is positive definite, namely $\mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_s)$. It therefore suffices to show that if σ is positive definite in restriction to $U \subseteq V$, then $\dim(U) \leqslant s$. Otherwise $U \cap \mathrm{span}(\mathbf{b}_{s+1}, \ldots, \mathbf{b}_n) \neq \{\mathbf{0}\}$. But then there is a $\mathbf{u} \in U \setminus \{\mathbf{0}\}$ for which $\sigma(\mathbf{u}, \mathbf{u}) \leqslant 0$, contradicting positive definiteness.

Since we have found basis independent characterisations for the numbers $n - s - t$ and s, and since t is uniquely determined by n and these, the claim is proved.

$\square$

Definition 3.2.7 The *signature* [Signatur] of a symmetric bilinear form is

$$(+^s, -^t, 0^d) = (\underbrace{+, \ldots, +}_{s}, \underbrace{-, \ldots, -}_{t}, \underbrace{0, \ldots, 0}_{d=n-s-t})$$

where the numbers s, t, d are determined according to the theorem as the number of λ_i in any diagonal representation of σ that are positive, negative, or equal to 0, respectively.

Exercise 3.2.2 Show that σ is non-degenerate iff its signature consists of 1 and -1 only, without 0.

Corollary 3.2.8 *For a symmetric bilinear form σ of signature $(+^s, -^t, 0^d)$ over $(V, \langle ., . \rangle)$: there is a basis consisting of pairwise orthogonal basis vectors such that σ is represented by the diagonal matrix*

$$A = \begin{pmatrix} E_s & & \\ & -E_t & \\ & & \mathbf{0} \end{pmatrix}.$$

Proof.   Diagonalise σ w.r.t. an orthonormal basis B, according to Theorem 3.2.5. A further rescaling of these basis vectors according to

$$
\hat{\mathbf{b}}_i := \begin{cases} \frac{1}{\sqrt{\sigma(\mathbf{b}_i, \mathbf{b}_i)}} \mathbf{b}_i & \text{if } \sigma(\mathbf{b}_i, \mathbf{b}_i) > 0 \\[2mm] \frac{1}{\sqrt{-\sigma(\mathbf{b}_i, \mathbf{b}_i)}} \mathbf{b}_i & \text{if } \sigma(\mathbf{b}_i, \mathbf{b}_i) < 0 \\[2mm] \mathbf{b}_i & \text{if } \sigma(\mathbf{b}_i, \mathbf{b}_i) = 0 \end{cases}
$$

and permutation of the $\hat{\mathbf{b}}_i$ if necessary, produces the desired result.

$\square$

In terms of matrices, the above considerations correspond to the following.

Proposition 3.2.9 *Let $A \in \mathbb{R}^{(n,n)}$ be symmetric ($A^t = A$). Then there is an orthogonal matrix $C \in \mathrm{O}(n) \subseteq \mathrm{GL}_n(\mathbb{R})$ such that $C^t A C = C^{-1} A C$ is diagonal.*

With $C \in \mathrm{GL}_n(\mathbb{R})$ one correspondingly achieves diagonal $C^t A C$ with entries from $\{1, -1, 0\}$. The numbers of entries 1, -1 and 0, respectively, in any such representation are uniquely determined.

Proof.   Regard A as the representation of a corresponding symmetric bilinear form $\sigma_A \colon (\mathbf{v}, \mathbf{w}) \mapsto \mathbf{v}^t A \mathbf{w}$. This means that A represents σ_A w.r.t. the standard basis of $\mathbb{R}^n$, which is an orthonormal basis of the standard scalar product. By Theorem 3.2.5, there is also an orthonormal basis such that σ_A is represented by a diagonal matrix D. Since the change of basis involved maps one orthonormal basis into another, the transformation is effected by an orthogonal matrix $C \in \mathrm{O}(n)$. We therefore have $D = C^t A C = C^{-1} A C$ diagonal as claimed.

The second point is similarly obtained with the help of Corollary 3.2.8 and Theorem 3.2.6.

$\square$

Exercise 3.2.3 Compare Definition 3.2.10 below for positive definiteness of a symmetric matrix $G \in \mathbb{R}^{(n,n)}$.

Show that for two symmetric matrices $A, G \in \mathbb{R}^{(n,n)}$ such that $G \in \mathbb{R}^{(n,n)}$ is also positive definite, there is some $C \in \mathrm{GL}_n(\mathbb{R})$ such that $C^t A C$ and $C^t G C$ are both diagonal.

[Hint: regard G as the representation of a scalar product over $\mathbb{R}^n$, and then proceed as above.]

3.2.3   Positive definiteness

Recall that a bilinear form σ on V is positive definite if $\sigma(\mathbf{v}, \mathbf{v}) \geqslant 0$ for all $\mathbf{v} \in V$ and if $\sigma(\mathbf{v}, \mathbf{v}) = 0$ implies $\mathbf{v} = \mathbf{0}$.

Definition 3.2.10 For a symmetric matrix $A \in \mathbb{R}^{(n,n)}$, $A = (a_{ij})_{1 \leqslant i,j \leqslant n}$:

  (i) A is *positive definite* if the bilinear form $(\mathbf{v}, \mathbf{w}) \mapsto \mathbf{v}^t A \mathbf{w}$ is positive definite on $\mathbb{R}^n$, i.e., if $\mathbf{v}^t A \mathbf{v} \geqslant 0$ for all $\mathbf{v} \in \mathbb{R}^n$ and $\mathbf{v}^t A \mathbf{v} = 0$ only for $\mathbf{v} = \mathbf{0}$.

 (ii) The k-th *principal minor* [Hauptminor] of A, for $k = 1, \ldots, n$, is the matrix
$$A_k = (a_{ij})_{1 \leqslant i,j \leqslant k} \in \mathbb{R}^{(k,k)},$$
the restriction of A to the first k rows and columns.

Part (i) of the definition is just such that a symmetric bilinear σ that is represented by the symmetric matrix A w.r.t. some basis of V is positive definite iff A is positive definite.

The following establishes criteria that can be useful in determining whether a given bilinear form or symmetric matrix is positive definite.

Proposition 3.2.11 *The following are equivalent for any symmetric matrix* $A \in \mathbb{R}^{(n,n)}$:

  (i) *A is positive definite.*

 (ii) *There is an orthogonal matrix C such that $C^t A C (= C^{-1} A C)$ is diagonal with positive entries on the diagonal.*

(iii) *All eigenvalues of A are positive (A viewed as the representation of an endomorphism of $\mathbb{R}^n$).*

(iv) *All principal minors of A have positive determinant: $|A_k| > 0$ for $k = 1, \ldots, n$.*

Proof.   Clearly the symmetric bilinear form represented by some diagonal matrix is positive definite iff the diagonal entries are all positive.

As any symmetric bilinear form and symmetric matrix can be diagonalised by means of an orthogonal change of basis transformation, the equivalence between (i),(ii) and (iii) is clear. [Note that $C^t A C = C^{-1} A C$ means that A is diagonalised in both its roles, as a representation of a symmetric bilinear form and as a representation of an endomorphism!]

(i) $\Rightarrow$ (iv). Let σ be the symmetric bilinear form on $\mathbb{R}^n$ defined by $\sigma(\mathbf{v}, \mathbf{w}) = \mathbf{v}^t A \mathbf{w}$ (the symmetric bilinear form over $\mathbb{R}^n$ represented by A w.r.t. the standard basis). Then A_k represents the restriction of σ to the k-dimensional subspace spanned by the first k standard basis vectors of $\mathbb{R}^n$. If A and hence σ are positive definite, so are these restrictions. Now A_k can be diagonalised by means of an orthogonal $C \in \mathrm{O}(k)$. Let $C^t A_k C = C^{-1} A_k C$ be diagonal with diagonal entries $\lambda_1, \ldots, \lambda_k$. By positive definiteness, all $\lambda_i > 0$. Hence $0 < \prod_{i=1}^k \lambda_i = |C^{-1} A_k C| = |A_k|$.

(iv) $\Rightarrow$ (i) is shown by induction on n. Base case, $n = 1$. In this case, $A = a \in \mathbb{R}$ and (iv) says that $a > 0$, hence positive definite.

Induction step $n \to n + 1$. Let $A \in \mathbb{R}^{(n+1,n+1)}$ be symmetric and let σ be the symmetric bilinear form over $\mathbb{R}^{n+1}$ represented by A w.r.t. the standard basis. Consider the n-th principal minor A_n of A, which represents the restriction of σ to the n-dimensional subspace of $\mathbb{R}^{n+1}$ spanned by the first n standard basis vectors. Since $|A_k| > 0$ for $k = 1, \ldots, n$, the induction hypothesis implies that this restriction of σ is positive definite.

Let $U_+ \subseteq \mathbb{R}^{n+1}$ be a maximal subspace on which σ is positive definite and containing the first n standard basis vectors. Suppose there is also a non-trivial subspace U_0 such that $\sigma(\mathbf{u}, \mathbf{u}) \leqslant 0$ for all $\mathbf{u} \in U_0$. It follows that $U_0 \cap U_+ = \{\mathbf{0}\}$, as any $\mathbf{v} \in (U_0 \cap U_+) \setminus \{\mathbf{0}\}$ would have to have $\sigma(\mathbf{v}, \mathbf{v}) > 0$ and $\sigma(\mathbf{v}, \mathbf{v}) \leqslant 0$. So $\dim(U_0) \leqslant 1$.

Therefore diagonalising A with an orthogonal $C \in \mathrm{O}(n + 1)$, we get a diagonal matrix $A' = C^t A C = C^{-1} A C$ with at most one diagonal entry that is non-positive. If A' had indeed one diagonal entry $\lambda \leqslant 0$ then $|A'| \leqslant 0$. But $|A'| = |C^{-1} A C| = |A| = |A_n| > 0$ by assumption. Therefore all the diagonal entries of A' are positive, and σ is positive definite. $\qquad\square$

3.3 Quadratic forms and quadrics

Let $\sigma \colon V \times V \to \mathbb{R}$ be a symmetric bilinear form. With it associate the induced function

$$
\begin{aligned}
Q \colon V &\longrightarrow \mathbb{R} \\
\mathbf{v} &\longmapsto Q(\mathbf{v}) := \sigma(\mathbf{v}, \mathbf{v}).
\end{aligned}
$$

Such a function is called a quadratic form, as defined below.

Note that for positive definite σ, this is just the square of the positive definite norm associated with σ.

Definition 3.3.1 A function $Q\colon V \to \mathbb{R}$ over the $\mathbb{R}$-vector space V is a *quadratic form* [quadratische Form] if

  (i) for all $\mathbf{v} \in V$ and $\lambda \in \mathbb{R}$: $Q(\lambda\mathbf{v}) = \lambda^2 Q(\mathbf{v})$.

  (ii) the map

$$\begin{aligned}
\sigma_Q\colon V \times V &\longrightarrow \mathbb{R}\\
(\mathbf{v}, \mathbf{w}) &\longmapsto \sigma_Q(\mathbf{v}, \mathbf{w}) := 1/2(Q(\mathbf{v} + \mathbf{w}) - Q(\mathbf{v}) - Q(\mathbf{w}))
\end{aligned}$$

    is (symmetric and) bilinear.

Clearly there is a one-to-one correspondence between quadratic forms and symmetric bilinear forms: according to (ii) in the definition, a quadratic form Q has an associated symmetric bilinear form σ_Q.

Observation 3.3.2 *For a quadratic form Q there is precisely one symmetric bilinear form σ such that $Q(\mathbf{v}) = \sigma(\mathbf{v}, \mathbf{v})$ for all $\mathbf{v} \in V$.*

Proof.   Clearly $\sigma := \sigma_Q$ as in (ii) of the definition is as desired. For uniqueness observe that $Q(\mathbf{v}) = \sigma(\mathbf{v}, \mathbf{v})$ for symmetric σ implies that

$$\begin{aligned}
Q(\mathbf{v} + \mathbf{w}) &= \sigma(\mathbf{v} + \mathbf{w}, \mathbf{v} + \mathbf{w}) = \sigma(\mathbf{v}, \mathbf{v}) + \sigma(\mathbf{w}, \mathbf{w}) + \sigma(\mathbf{v}, \mathbf{w}) + \sigma(\mathbf{w}, \mathbf{v})\\
&= Q(\mathbf{v}) + Q(\mathbf{w}) + 2\sigma(\mathbf{v}, \mathbf{w}).
\end{aligned}$$

Hence $\sigma(\mathbf{v}, \mathbf{w}) = 1/2(Q(\mathbf{v} + \mathbf{w}) - Q(\mathbf{v}) - Q(\mathbf{w}))$ is recovered from Q. $\qquad\square$

The matrix representation of the quadratic form Q w.r.t. basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ is the symmetric matrix that represents σ_Q, with entries

$$q_{ij} = \sigma_Q(\mathbf{b}_i, \mathbf{b}_j) = \frac{1}{2}(Q(\mathbf{b}_i + \mathbf{b}_j) - Q(\mathbf{b}_i) - Q(\mathbf{b}_j)).$$

Example 3.3.3 The quadratic form associated with the standard scalar product in $\mathbb{R}^n$ is $Q((x_1, \ldots, x_n)) = \sum_i x_i^2$, the square of the standard euclidean norm.

For $c \in \mathbb{R}$, the quadratic equation

$$Q(\mathbf{x}) = \sum_i x_i^2 = c$$

defines

(i) for $c > 0$: a sphere of radius $\sqrt{c}$ in $\mathbb{R}^n$.
   Especially, for $c = 1$, the unit sphere $S^{n-1} \subseteq \mathbb{R}^n$.

(ii) for $c < 0$: the empty set.

(iii) for $c = 0$: the singleton set $\{\mathbf{0}\}$.

Example 3.3.4 For the symmetric bilinear form $\sigma((x,y),(x',y')) = xx' - yy'$ over $\mathbb{R}^2$, the associated quadratic form is $Q(x,y) = x^2 - y^2$.

For $c \in \mathbb{R}$, the quadratic equation

$$Q(\mathbf{x}) = x^2 - y^2 = c$$

defines

(i) for $c \neq 0$: a *hyperbola* (with two branches), symmetric w.r.t. the x- and y-axes:

$$\mathbb{X}_c = \{(x,y) \in \mathbb{R}^2 : x^2 - y^2 = c\} = \{(x,y) \in \mathbb{R}^2 : x + y = \frac{c}{x - y}\},$$

with asymptotic lines $x + y = 0$ and $x - y = 0$.

(ii) for $c = 0$: the union of the two lines $x + y = 0$ and $x - y = 0$.

We now use the principal axes of the associated bilinear form σ_Q in order to analyse the geometry of Q. W.l.o.g. we consider $\mathbb{R}^n$ with the standard scalar product.

Proposition 3.3.5 *Let Q be a quadratic form on $\mathbb{R}^n$. Then there are an orthonormal basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $(\mathbb{R}^n, \langle .,.\rangle)$ and numbers $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ such that for $\mathbf{x} = \sum_i x_i \mathbf{b}_i$:*

$$Q(\mathbf{x}) = \sum_{i=1}^{n} \lambda_i \, x_i^2.$$

Proof.   Direct from Theorem 3.2.5, for $\sigma := \sigma_Q$.

$\square$

Geometrically, a rescaling of the basis vectors from some orthonormal basis (with positive real factors) corresponds to distortions in the direction of the orthogonal axes given by these basis vectors. Just as in Corollary 3.2.8, we thus find that, up to such a rescaling and with a permutation of the basis vectors, any quadratic form is representable as follows.

Corollary 3.3.6 *For every quadratic form Q over $\mathbb{R}^n$ there are a basis of pairwise orthogonal (but not necessarily normalised) vectors $\mathbf{b}_i$ and numbers $s, t \in \mathbb{N}$ such that $k := s+t \leqslant n$, such that in terms of coordinates w.r.t. this basis, Q is given by*

$$Q(x_1, \ldots, x_n) = \sum_{i=1}^{s} x_i^2 - \sum_{i=s+1}^{k} x_i^2.$$

The *signature* of the quadratic form Q is defined, in accordance with the definition for symmetric bilinear forms, to be $(+^s, -^t, 0^d)$, for s, t and $d = n - s - t$ as in the corollary.

3.3.1   Quadrics in $\mathbb{R}^n$

Definition 3.3.7 A *quadric* [Quadrik] $\mathbb{X} \subseteq \mathbb{R}^n$ is a subset defined by a quadratic equation

$$\mathbb{X} = \big\{ \mathbf{x} \in \mathbb{R}^n \colon \mathbf{x}^t A \mathbf{x} + \mathbf{b}^t \mathbf{x} + c = 0 \big\},$$

where $A \in \mathbb{R}^{(n,n)}$ is a symmetric matrix, $\mathbf{b} \in \mathbb{R}^n$ and $c \in \mathbb{R}$.

We may think of a quadric $\mathbb{X}$ as being defined in terms of an associated quadratic form $Q(\mathbf{x}) = \mathbf{x}^t A \mathbf{x}$ with matrix representation A, a linear form $\eta(\mathbf{x}) = \mathbf{b}^t \mathbf{x} = \langle \mathbf{b}, \mathbf{x} \rangle$, and a constant term $c \in \mathbb{R}$.

Exercise 3.3.1 Show that any quadratic equation over $\mathbb{R}^n$, of the form

$$\sum_{1 \leqslant i,j \leqslant n} \alpha_{ij} x_i x_j + \sum_{1 \leqslant i \leqslant n} \beta_i x_i + \gamma = 0$$

for arbitrary $\alpha_{ij}, \beta_i, \gamma \in \mathbb{R}$, can be represented in the form $\mathbf{x}^t A \mathbf{x} + \mathbf{b}^t \mathbf{x} + c = 0$ for a suitable *symmetric* matrix A and suitable $\mathbf{b}$ and c.

Example 3.3.8 We classify the quadrics without linear terms ($\mathbf{b} = \mathbf{0}$)

$$\mathbb{X} = \big\{ \mathbf{x} \in \mathbb{R}^3 \colon Q(\mathbf{x}) = c \big\} \subseteq \mathbb{R}^3,$$

up to bijective linear transformations. It turns out that the signature of the quadratic form and a case distinction w.r.t. $c = 0$, $c > 0$, or $c < 0$ determines the type of $\mathbb{X}$. Since a simultaneous change of sign in c and in the signature,

corresponds to just a central reflection of $\mathbb{X}$ in $\mathbf{0}$, we need consider only the signatures $(+, +, +)$, $(+, +, -)$, $(+, +, 0)$, $(+, -, 0)$, $(+, 0, 0)$ and $(0, 0, 0)$. We thus look at these cases:

Signature $(+, +, +)$. σ_Q is positive definite. In a suitable orthonormal basis $Q(x, y, z) = \alpha x^2 + \beta y^2 + \gamma z^2$, where $\alpha, \beta, \gamma > 0$.
    For $c < 0$ the quadric is empty; for $c = 0$ it consists of just the null vector. For $c > 0$ the quadric

$$\mathbb{X} = \{(x, y, z) \in \mathbb{R}^3 \colon \alpha x^2 + \beta y^2 + \gamma z^2 = c\}$$

is an *ellipsoid*, which may be regarded as the image of the unit sphere $S^2 \subseteq \mathbb{R}^3$ under the transformation

$$(x, y, z) \longmapsto (\sqrt{c/\alpha}\, x, \sqrt{c/\beta}\, y, \sqrt{c/\gamma}\, z).$$

Up to linear transformations, $\mathbb{X}$ is the unit sphere.

Signature $(+, +, -)$. In a suitable orthonormal basis $Q(x, y, z) = \alpha x^2 + \beta y^2 - \gamma z^2$, where $\alpha, \beta, \gamma > 0$.
    Through normalisation with factors $\sqrt{|c|/\alpha}$, $\sqrt{|c|/\beta}$, $\sqrt{|c|/\gamma}$, the quadric transforms into

$$\mathbb{X} = \{(x, y, z) \colon x^2 + y^2 - z^2 = c\},$$

where $c \in \{0, 1, -1\}$. Note that $\mathbb{X}$ is rotation symmetric about the z-axis; note that $x^2 + y^2$ is the square of the radial distance of (x, y, z) from the z-axis.
    For $c = 0$, $\mathbb{X}$ is the *cone* $\{\mathbf{x} = (x, y, z) \colon x^2 + y^2 = z^2\}$.
    For $c = 1$, $\mathbb{X} = \{(x, y, z) \colon x^2 + y^2 = 1 + z^2\}$ is the surface generated by rotation about the z-axis from the hyperbola [Hyperbel] in the x-z-plane

$$\{(x, z) \in \mathbb{R}^2 \colon x^2 = 1 + z^2\} = \{(x, z) \in \mathbb{R}^2 \colon x + z = \frac{1}{x - z}\}.$$

This is called a *single-sheet hyperboloid* ('single-sheet' because it has one connected component). Interestingly, the same surface is generated by rotation of the line

$$\{(1, 0, 0) + \lambda(0, 1, 1) \colon \lambda \in \mathbb{R}\}$$

about the z-axis.

For $c = -1$, $\mathbb{X} = \{(x, y, z)\colon x^2 + y^2 = z^2 - 1\}$ is the surface generated by rotation about the z-axis from the hyperbola in the x-z-plane

$$\{(x, z) \in \mathbb{R}^2\colon z^2 - x^2 = 1\} = \{(x, z) \in \mathbb{R}^2\colon z + x = \frac{1}{z - x}\},$$

a *two-sheet hyperboloid* ('two-sheet' because of the two connected components).

Signature $(+, +, 0)$. In a suitable orthonormal basis $Q(x, y, z) = \alpha x^2 + \beta y^2$, where $\alpha, \beta > 0$.

For $c < 0$, the quadric $\mathbb{X}$ is empty, for $c = 0$ it consists of the z-axis.

For $c > 0$, the quadric is an *elliptic cylinder*, erected in the direction of the z-axis over the ellipse $\{(x, y)\colon \alpha x^2 + \beta y^2 = 0\}$ in the x-y-plane. Up to linear transformations it is equivalent to the standard cylinder

$$\mathbb{X} = \{(x, y, z)\colon x^2 + y^2 = 1\}.$$

Signature $(+, -, 0)$. In a suitable orthonormal basis $Q(x, y, z) = \alpha x^2 - \beta y^2$, where $\alpha, \beta > 0$, and after transformation with corresponding factors,

$$\mathbb{X} = \{(x, y, z)\colon x^2 - y^2 = c\},$$

where $c \in \{0, -1, 1\}$.

For $c = 0$, $\mathbb{X}$ is the union of the two planes defined by $x = y$ and $x = -y$, respectively.

For $c = 1$ or -1, $\mathbb{X}$ is a *cylinder* erected in the z-direction over suitable hyperbolas in the x-y-plane.

Signature $(+, 0, 0)$. In a suitable orthonormal basis $Q(x, y, z) = \alpha x^2$, $\alpha > 0$.

For $c = 0$ the quadric $\mathbb{X}$ consists of the y-z-plane, defined by $x = 0$.

For $c < 0$, $\mathbb{X} = \emptyset$.

For $c > 0$, $\mathbb{X}$ consists of two parallel planes, defined by $x = \pm\sqrt{c/\alpha}$.

Signature $(0, 0, 0)$. $Q(x, y, z) = 0$; for $c = 0$ the quadric $\mathbb{X}$ is the whole space, otherwise it is empty.

We shall further classify also all quadrics with linear terms over $\mathbb{R}^2$ below. For that however, it will be convenient to regard them in the context of projective spaces, which will simplify the analysis.

Consider the general quadric

$$\mathbb{X} = \{\mathbf{x} \in \mathbb{R}^n \colon \mathbf{x}^t A \mathbf{x} + \mathbf{b}^t \mathbf{x} + c = 0\}.$$

For the geometric study of $\mathbb{X}$ as a subset of $\mathbb{R}^n$, we now consider as equivalence transformations all linear transformations (arbitrary choice of basis) and translations (arbitrary choice of an origin). In other words, we work up to affine rather than just up to linear transformations for the sake of a more uniform analysis.

Let $Q(\mathbf{x}) = \mathbf{x}^t A \mathbf{x}$ be the associated quadratic form, $(+^s, -^t, 0^{n-s-t})$ the signature of Q. Put $k := s + t$. By a suitable choice of basis (orthogonal but not necessarily orthonormal) in $\mathbb{R}^n$, we may assume that A is diagonal with non-zero entries $a_{ii} = 1$ for $i = 1, \ldots, s$ and $a_{ii} = -1$ for $i = s+1, \ldots, k$. So w.o.l.g.

$$\mathbb{X} = \{\mathbf{x} \in \mathbb{R}^n \colon \textstyle\sum_{i=1}^s x_i^2 - \sum_{i=s+1}^k x_i^2 + \sum_{i=1}^n b_i x_i + c = 0\}.$$

We now use a translation to further simplify the linear term $\sum_i b_i x_i$. A translation through vector $-\mathbf{u}$ for $\mathbf{u} = (u_1, \ldots, u_n) \in \mathbb{R}^n$ transforms $\mathbb{X}$ into $\{\mathbf{x} \colon \mathbf{x} + \mathbf{u} \in \mathbb{X}\}$, defined by the equation

$$\textstyle\sum_{i=1}^s (x_i + u_i)^2 - \sum_{i=s+1}^k (x_i + u_i)^2 + \sum_{i=1}^n b_i(x_i + u_i) + c = 0,$$

which is equivalent to

$$
\begin{array}{lll}
& \sum_{i=1}^s x_i^2 - \sum_{i=s+1}^k x_i^2 & \text{[quadratic]} \\
+ & \sum_{i=1}^s (b_i + 2u_i)x_i + \sum_{i=s+1}^k (b_i - 2u_i)x_i + \sum_{i=k+1}^n b_i x_i & \text{[linear]} \\
+ & \sum_{i=1}^s u_i^2 - \sum_{i=s+1}^k u_i^2 + \sum_{i=1}^n b_i u_i + c & \text{[constant]} \\
= & 0. &
\end{array}
$$

So we see that we can eliminate linear terms in x_i for $i = 1, \ldots, k$ by choosing $u_i := -b_i/2$ for $i = 1, \ldots, s$ and $u_i := b_i/2$ for $i = s+1, \ldots, k$. This puts $\mathbb{X}$ into the form

$$\{\mathbf{x} \in \mathbb{R}^n \colon \textstyle\sum_{i=1}^s x_i^2 - \sum_{i=s+1}^k x_i^2 + \sum_{i=k+1}^n b_i x_i + c = 0\},$$

with new b_i and c.

A change in the basis vectors $\mathbf{b}_{k+1}, \ldots, \mathbf{b}_n$ does not affect the diagonalisation of Q. Note that the vector $\mathbf{b}$ responsible for the new linear term is in $\mathrm{span}(\mathbf{b}_{k+1}, \ldots, \mathbf{b}_n)$. If $\mathbf{b}_{k+1}, \ldots, \mathbf{b}_n$ are replaced by new pairwise orthogonal

basis vectors starting with $\mathbf{b}'_{k+1} := \mathbf{b}$, if $\mathbf{b} \neq \mathbf{0}$ then we obtain the following standardised form of a general quadric

$$\mathbb{X} = \left\{ \mathbf{x} \in \mathbb{R}^n \colon \sum_{i=1}^{s} x_i^2 - \sum_{i=s+1}^{k} x_i^2 + \delta\, x_{k+1} + c = 0 \right\}.$$

Up to an affine transformation, any quadric with quadratic form of signature $(+^s, -^t, 0^{n-k})$ can be brought into this form, for a $\delta \in \{0, 1\}$. If the signature has no 0, i.e., if $k = s + t = n$ then the linear part is completely eliminated.

3.3.2 Projective space $\mathbb{P}^n$

Definition 3.3.9 With an $\mathbb{R}$-vector space V, associate the set $\mathbb{P}(V)$ of all its 1-dimensional subspaces (lines through $\mathbf{0}$). $\mathbb{P}(V)$ is called the *projective space* [projektiver Raum] associated to V. If V is of dimension $n + 1$, $\mathbb{P}(V)$ is said to have dimension n.

$\mathbb{P}^n := \mathbb{P}(\mathbb{R}^{n+1})$ is the standard n-dimensional projective space.

If $U \subseteq V$ is a linear subspace of V, then $\mathbb{P}(U) \subseteq \mathbb{P}(V)$ is the *projective subspace* of $\mathbb{P}(V)$ consisting of all the 1-dimensional subspaces (lines through $\mathbf{0}$) that are contained in U.

One may think of the elements of $\mathbb{P}^n$ as equivalence classes of vectors in $\mathbb{R}^{n+1} \setminus \{\mathbf{0}\}$, where $\mathbf{v}, \mathbf{w} \neq \mathbf{0}$ are equivalent if they are non-zero multiples of each other, or as pairs of antipodal points of the unit sphere $S^n \subseteq \mathbb{R}^{n+1}$. Projective geometry is useful in connection with central perspective and imaging techniques. As we shall see below it also gives rise to a unified view of certain notions in affine geometry, especially related to quadrics.

We write $[\mathbf{x}]$ for the equivalence class of $\mathbf{x} \neq \mathbf{0}$:

$$[\mathbf{x}] = [(x_1, \ldots, x_n, x_{n+1})] = \{\lambda \cdot (x_1, \ldots, x_n, x_{n+1}) \colon \lambda \in \mathbb{R} \setminus \{0\}\},$$

and think of $\mathbf{x} = (x_1, \ldots, x_n, x_{n+1})$ as coordinates of $[\mathbf{x}]$, which are determined only up to non-zero scalar multiples. Such coordinates are called *homogeneous coordinates* for points in $\mathbb{P}^n$.

For $x_{n+1} \neq 0$, $[\mathbf{x}]$ possesses a unique representative with $x_{n+1} = 1$ in the *affine hyperplane*

$$\mathbb{A} = \{(\mathbf{x}', 1) \colon \mathbf{x}' \in \mathbb{R}^n\} \subseteq \mathbb{R}^{n+1}.$$

The remaining points in $\mathbb{P}^n$ correspond to lines through $\mathbf{0}$ that do not intersect $\mathbb{A}$ (parallel to $\mathbb{A}$), or inside the linear subspace U corresponding to $\mathbb{A}$. These points form a projective subspace $\mathbb{P}(U) \subseteq \mathbb{P}^n$ which is isomorphic to $\mathbb{P}^{n-1} = \mathbb{P}(\mathbb{R}^n)$ if we identify U with $\mathbb{R}^n$ according to $U = \{(\mathbf{x}', 0) \colon \mathbf{x}' \in \mathbb{R}^n\}$.

The projective space $\mathbb{P}^n = \mathbb{P}(\mathbb{R}^{n+1})$ may thus be visualised as the disjoint union of the n-dimensional affine hyperplane $\mathbb{A}$ and a "*projective hyperplane at infinity*" [unendlich ferne Ebene]. In the case of the projective plane $\mathbb{P}^2 = \mathbb{P}(\mathbb{R}^3)$, this corresponds to its representation as the disjoint union of an affine plane and a "*line at infinity*." [Just like any other 1-dimensional projective subspace, the line at infinity has the topological type of the circle S^1, rather than that of a real line.]

Quadrics, affine and projective

A general quadric over $\mathbb{R}^{n+1}$, with linear and constant terms, does not define a subset of $\mathbb{P}^n$ in a natural way, because the defining equation is not homogeneous. It therefore does not in general define a set of lines in $\mathbb{R}^{n+1}$, or of points in $\mathbb{P}^n$. More generally, a system of equations over $\mathbb{R}^{n+1}$ defines a subset in projective space, if each equation is homogeneous in the sense of having the same degree in the coefficients x_i in all terms. For a quadric in

$\mathbb{R}^{n+1}$ to define a subset of $\mathbb{P}^n$, its linear and constant terms must be 0. We therefore consider quadrics of the form

$$\mathbb{X} = \{\mathbf{x} \in \mathbb{R}^{n+1} \colon Q(\mathbf{x}) = 0\},$$

which we take to define the projective set of points

$$\mathbb{P}(\mathbb{X}) = \{[\mathbf{x}] \in \mathbb{P}^n \colon \mathbf{x} \in \mathbb{X} \setminus \{\mathbf{0}\}\} \subseteq \mathbb{P}^n.$$

A general quadric, with linear and constant terms over $\mathbb{R}^n$, though, can always be extended to a homogeneous quadric over $\mathbb{R}^{n+1}$. We simply use the extra coordinate x_{n+1} to pad all terms that are of sub-quadratic degree. From the general quadric

$$\mathbb{X} = \left\{\mathbf{x} \in \mathbb{R}^n \colon \ \sum_{i=1}^{s} x_i^2 - \sum_{i=s+1}^{k} x_i^2 + \delta\, x_{k+1} + c = 0\right\} \subseteq \mathbb{R}^n$$

we pass to the homogeneous, projective quadric

$$\overline{\mathbb{X}} := \left\{(\mathbf{x}, x_{n+1}) \in \mathbb{R}^{n+1} \colon \ \sum_{i=1}^{s} x_i^2 - \sum_{i=s+1}^{k} x_i^2 + \delta\, x_{k+1}x_{n+1} + cx_{n+1}^2 = 0\right\}$$

in R^{n+1}, which defines a subset of $\mathbb{P}^n$ in terms of homogeneous coordinates.

The effect of this is that the intersection of the new quadric in $\mathbb{R}^{n+1}$ with the affine hyperplane $\mathbb{A}$ defined by $x_{n+1} = 1$, exactly corresponds to the given (inhomogeneous) quadric:

$$\mathbb{X} = \overline{\mathbb{X}} \cap \mathbb{A}.$$

General quadrics are thus seen to be the affine parts of projective quadrics.

Conic sections

Returning to the general quadrics in two-dimensional (projective and affine) planes, we are thus led to analyse them in terms of the projective quadrics, or the homogeneous quadrics in $\mathbb{R}^3$ which are of the form

$$Q(x, y, z) = 0.$$

There are only just a few types, depending on the signature of the quadratic form: $(+, +, +)$, $(+, +, -)$, $(+, +, 0)$, $(+, -, 0)$, $(+, 0, 0)$ and $(0, 0, 0)$. We may resort to our classification above, looking just at the cases $c = 0$.

For non-degenerate Q, the only non-trivial case to be considered is that of signature $(+, +, -)$. Up to linear transformations as detailed above,

$$\overline{\mathbb{X}} = \left\{(x, y, z) \colon x^2 + y^2 = z^2\right\}$$

defines a cone, which is clearly a union of lines through $\mathbf{0}$ in $\mathbb{R}^3$. All non-degenerate general quadrics in the affine plane are thus seen to be *conic sections* [Kegelschnitte] in the sense of being (linearly equivalent to) the intersections of the image of the standard cone $\overline{\mathbb{X}}$ under some linear transformation with the affine hyperplane $\mathbb{A} = \{(x, y, z) \colon z = 1\} \subseteq \mathbb{R}^3$.

The only distinguishing feature between the different affine incarnations is determined by where the projective line at infinity intersects $\overline{\mathbb{X}} \setminus \{\mathbf{0}\}$. We let U be the x-y-plane, parallel to $\mathbb{A}$; $\mathbb{P}(U)$ the line at infinity. There are the following cases:

- no intersection: $\mathbb{P}(U) \cap \overline{\mathbb{X}} = \emptyset$. U cuts the cone in the origin only. The affine part is an ellipse (or, up to affine transformations within $\mathbb{A}$, equivalent to the unit circle S^1).

- one point of intersection: $|\mathbb{P}(U) \cap \overline{\mathbb{X}}| = 1$. U cuts the cone in precisely one line on its mantle. The affine part is a parabola (or, up to affine equivalence the standard *parabola* $y = x^2$).

- two points of intersection: $|\mathbb{P}(U) \cap \overline{\mathbb{X}}| = 2$. U cuts the cone in precisely two lines on its mantle. The affine part is a hyperbola (or, up to affine equivalence the standard *hyperbola* $x + y = \frac{1}{x-y}$).

Remark 3.3.10 In a similar way, and one dimension up, qualitatively different three-dimensional affine quadrics as for instance the single-sheet hyperboloid and its relative, the saddle surface $\{(x, y, z) \colon x^2 - y^2 = z\}$ are seen to be different affine sections of the projective quadric $\overline{\mathbb{X}} = \{[\mathbf{x}] \colon x_4 x_3 - x_1 x_2 = 0\}$ of signature $(+, +, -, -)$. The three other prominent three-dimensional affine quadrics, the two-sheet hyperboloid, the ellipsoid and the paraboloid are also just distinct affine views of the same projective point set, viz., of the projective quadric with signature $(+, +, +, -)$.

Exercise 3.3.2 Determine the projective hull $\overline{\mathbb{X}}$ of the above-mentioned three-dimensional affine quadrics and determine the underlying signatures. In order to demonstrate the distinctions between the two mentioned families – single-sheet hyperboloids and saddle surfaces on one hand, and two-sheet hyperboloids, the ellipsoids and paraboloids on the other – show that the former are generated by families of affine lines, while the latter do not contain any lines. It follows that members of the second family cannot occur as projective images of members from the first family (why?). Within each family, describe the different affine views in terms of different choices for the plane at infinity.

Index

Index (deutsch)