

Linear Algebra II

Ulrich Pennig

November 25, 2019

Abstract

These are lecture notes I created for a one semester second year course about Linear Algebra at Cardiff University. Despite some careful proofreading, the document is probably still riddled with typos. If you want to help improve it, please send any comments or corrections to pennigu@cardiff.ac.uk.

Contents

1	The Story So Far...	3
2	Linear Transformations and Matrices	8
2.1	Composition and matrix multiplication	12
3	Invertibility of Linear Transformations	17
4	Systems of Linear Equations	26
4.1	Solving Systems of Linear Equations	30
4.2	Elementary Matrices	33
4.3	The Rank of a Matrix	36
5	Determinants	44
5.1	That horrible formula	48
5.2	Summary of the properties of the determinant	50
5.3	The determinant of a linear transformation $T: V \rightarrow V$	51
5.4	Minors and cofactor expansion – a nice formula for a change	52
6	Diagonalisation, Eigenvalues and Eigenvectors	58
6.1	Diagonalization: A Step by Step Instruction Manual	69
7	The Jordan Normal Form	70
7.1	Generalised eigenspaces	72
7.2	Nilpotent linear transformations	78

7.3	Low dimensional examples	85
7.3.1	Dimension 2, two distinct eigenvalues $\lambda_1 \neq \lambda_2$	85
7.3.2	Dimension 2, one eigenvalue $\lambda = \lambda_1 = \lambda_2$	86
7.3.3	Dimension 3, three distinct eigenvalues	87
7.3.4	Dimension 3, two distinct eigenvalues	87
7.3.5	Dimension 3, one eigenvalue	89
8	Duals, symmetric bilinear forms and inner product spaces	92
8.1	Dual spaces	93
8.2	Bilinear forms and inner products on real vector spaces	96
8.3	Hermitian forms and inner products on complex vector spaces	106
8.4	Orthonormal sets and orthonormal bases	108
8.5	Adjoints and Isometries	111
9	What's next?	120

1 The Story So Far...

The main objects of study in linear algebra are vector spaces and linear maps between them. In this section we will recap their definition and the main statements about them we have learned so far. As mentioned in the beginning, we will only deal with vector spaces over the real numbers \mathbb{R} , the complex numbers \mathbb{C} or the rational numbers \mathbb{Q} in this lecture. The thing that \mathbb{R} , \mathbb{Q} and \mathbb{C} have in common is that they all fit the definition of a **field**. This is a set together with an addition and a multiplication map, such that, in particular, the set contains elements that behave like zero and one and we can divide by any non-zero element. We will spare the reader the full definition of this structure, which is even more involved than that of a vector space. Instead, whenever a field F appears in the following, the reader should replace F by \mathbb{Q} , \mathbb{R} or \mathbb{C} in her or his mind. That said, it should be mentioned that there are more general fields as well, for example, ones that have only finitely many elements. Without further ado, here is now the definition of a vector space over a field F .

Definition 1.0.1. A **vector space** V **over a field** F is a set on which two operations, called **addition** and **scalar multiplication**, are defined, so that for each $x, y \in V$ there is a unique element $x + y \in V$ and for each $a \in F$ and each $x \in V$ there is a unique element $ax \in V$, such that the following conditions hold:

VS 1) For all $x, y \in V$ we have $x + y = y + x$.

VS 2) For all $x, y, z \in V$ we have $(x + y) + z = x + (y + z)$.

VS 3) There exists an element $0 \in V$ such that $x + 0 = x$ for all $x \in V$.

VS 4) For all $x \in V$ there exists an element $y \in V$ such that $x + y = 0$.

VS 5) For each $x \in V$, $1x = x$ (with $1 \in F$).

VS 6) For all $a, b \in F$ and each $x \in V$ we have $(ab)x = a(bx)$.

VS 7) For all $a \in F$ and each pair $x, y \in V$ we have $a(x + y) = ax + ay$.

VS 8) For all $a, b \in F$ and each $x \in V$ we have $(a + b)x = ax + bx$.

Example 1.0.2. Let us fix $F = \mathbb{R}$ and let $n \in \mathbb{N}$. The set $V = \mathbb{R}^n$ is a vector space over \mathbb{R} . If $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$, then the element $x + y$ is given by

$$(x_1, \dots, x_n) + (y_1, \dots, y_n) = (x_1 + y_1, \dots, x_n + y_n) .$$

Moreover, for $a \in \mathbb{R}$ the element ax is given by $ax = (ax_1, \dots, ax_n)$. Likewise, \mathbb{C}^n is a vector space over \mathbb{C} and \mathbb{Q}^n is a vector space over \mathbb{Q} with similar definitions of the addition and the scalar multiplication. In general F^n is a vector space over F .

Whenever we have a vector space V over a field F , we can look for subsets $W \subseteq V$ that are vector spaces with respect to the addition and scalar multiplication restricted from V to W . Such subsets are called subspaces (or linear subspaces).

Definition 1.0.3. Let V be a vector space over a field F . A subset W of V is called a **(linear) subspace** of V if W is a vector space over F with respect to the operations of addition and scalar multiplication defined on V .

Example 1.0.4. The linear subspaces of \mathbb{R}^n have a geometric interpretation and are easy to draw for $n \leq 2$. The one-dimensional subspaces of \mathbb{R}^2 for example correspond to the lines in \mathbb{R}^2 passing through the origin. (For the definition of dimension see Def. 1.0.8.)




Figure 1: The blue line is a linear subspace of \mathbb{R}^2 .

Not every subset S of a vector space V is automatically a linear subspace. However, each subset S can be extended to the linear subspace generated by S , which is called the span of S :

Definition 1.0.5. Given a non-empty subset $S \subset V$ the **(linear) span of S** , denoted $\text{span}(S)$ is the linear subspace of V consisting of all linear combinations of the vectors in S . We define $\text{span}(\emptyset) = \{0\}$.

A subset $S \subset V$ is said to **generate the vector space V** if $V = \text{span}(S)$.

Consider the vector space \mathbb{R}^3 over the field \mathbb{R} and let $v_1, v_2, v_3 \in \mathbb{R}^3$ be the following vectors:

$$v_1 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad v_2 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \quad \text{and} \quad v_3 = \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix}.$$

Observe that for any vector $w \in \mathbb{R}^3$ we can find unique scalars $\alpha_1, \alpha_2, \alpha_3 \in \mathbb{R}$ such that $w = \alpha_1 v_1 + \alpha_2 v_2 + \alpha_3 v_3$. For example,

$$\begin{pmatrix} 5 \\ 3 \\ 2 \end{pmatrix} = 2 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + 3 \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} + 0 \begin{pmatrix} 0 \\ 0 \\ -1 \end{pmatrix} = 2v_1 + 3v_2 + 0v_3$$

and $\alpha_1 = 2, \alpha_2 = 3, \alpha_3 = 0$ are the only scalars that produce the vector $w = (5, 3, 2)$. This is no longer true for the following vectors v'_1, v'_2 and v'_3 :

$$v'_1 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad v'_2 = \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \quad \text{and} \quad v'_3 = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}.$$

To see an example we can use the same vector w as above. We have

$$\begin{pmatrix} 5 \\ 3 \\ 2 \end{pmatrix} = 2 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + 3 \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} + 0 \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix} = 3 \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} + 2 \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} + 1 \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix} .$$

Therefore the decomposition of this particular vector w is no longer unique. Depending on w the equation $w = \alpha_1 v'_1 + \alpha_2 v'_2 + \alpha_3 v'_3$ might have more than one solution $\alpha_1, \alpha_2, \alpha_3$ or it could even have no solution at all. Observe that $-v'_1 + v'_2 - v'_3 = 0$.

The difference between the two sets $\beta = \{v_1, v_2, v_3\}$ and $\beta' = \{v'_1, v'_2, v'_3\}$ is that β is linearly independent whereas β' is linearly dependent. These are two of the **most important concepts** of linear algebra. They are defined as follows:

Definition 1.0.6. Let V be a vector space over a field F . A subset $S \subset V$ is called **linearly dependent** if there exists a finite number of distinct vectors $u_1, \dots, u_n \in S$ and scalars $\alpha_1, \dots, \alpha_n \in F$, not all zero, such that

$$\alpha_1 u_1 + \dots + \alpha_n u_n = 0 .$$

We say that a subset $S \subset V$ is **linearly independent** if it is not linearly dependent. Hence, S is linearly independent if and only if for any finite number of distinct vectors $u_1, \dots, u_n \in S$, the only solution $\alpha_1, \dots, \alpha_n \in F$ to the equation $\alpha_1 u_1 + \dots + \alpha_n u_n = 0$ is given by $\alpha_1 = \dots = \alpha_n = 0$.

The subset $\beta = \{v_1, v_2, v_3\} \subset \mathbb{R}^3$ we discussed above is not only linearly independent, but it also spans \mathbb{R}^3 . Subsets with these two properties are of fundamental importance in Linear Algebra:

Definition 1.0.7. A subset $\beta \subset V$ is a **basis for V** if β is linearly independent and spans V .

Let V be a vector space over a field F . If β and β' are both bases for V , then there is a bijection between them. In particular, the number of elements of any basis for V depends only on V and not on the choice of basis.

Definition 1.0.8. The **dimension** $\dim(V)$ of V is the cardinality (i.e. the number of elements) of any basis $\beta \subset V$.

Given two vector spaces V and W , we can consider the maps $T: V \rightarrow W$ from V to W that transform the addition and the scalar multiplication of V into that of W . Such maps are crucial for everything that follows. Hence, we make the following definition:

Definition 1.0.9. Let V and W be vector spaces over the same field F . We call a map $T: V \rightarrow W$ a **linear transformation (or linear map) from V to W** if for all $x, y \in V$ and $a \in F$ we have

(a) $T(x + y) = T(x) + T(y)$ and

(b) $T(ax) = aT(x)$.

The following theorem is often helpful if we want to check the linearity of a map.

Lemma 1.0.10. *Let V, W be a vector spaces over the field F and let $T: V \rightarrow W$ be a map. Then T is a linear transformation if and only if*

$$T(ax + y) = aT(x) + T(y) \tag{1}$$

for all $x, y \in V$ and $a \in F$.

Proof. If we set $a = 1$ in (1) but choose arbitrary x and y , we see that condition a) in Def. 1.0.9 follows. Now note that for $a = 1$ and $x = y = 0$ we have $T(0) = T(0) + T(0)$, which implies $T(0) = 0$ after subtracting $T(0)$ from both sides. If we now choose a and x to be arbitrary and set $y = 0$ we therefore obtain b). \square

Whenever we have a linear transformation $T: V \rightarrow W$ between two vector spaces V and W over a field F , there are two important subspaces of V , respectively W , associated to T , which we define now.

Definition 1.0.11. The **null space (or kernel) of T** is denoted by $\ker(T)$ (sometimes also by $N(T)$) and is defined as follows

$$\ker(T) = \{v \in V \mid T(v) = 0\} ,$$

i.e. it consists of all vectors in V that are mapped to 0 in W by T . This is a subspace of V . The **range (or image) of T** is denoted by $\text{Im}(T)$ and defined as

$$\text{Im}(T) = \{T(v) \in W \mid v \in V\} .$$

The range of T consists of all vectors in the image of T . It is a subspace of W .

If $\ker(T)$ is finite-dimensional we define the **nullity** of T , denoted by $\text{nullity}(T)$, to be the dimension of $\ker(T)$. Likewise, if $\text{Im}(T)$ is finite-dimensional we define the **rank** of T , denoted by $\text{rank}(T)$, to be the dimension of $\text{Im}(T)$.

A linear transformation $T: V \rightarrow W$ is injective if and only if $\ker(T) = \{0\}$. It is surjective if and only if $\text{Im}(T) = W$. More generally, the following dimension formula holds, which is also known as the **rank-nullity theorem**:

$$\text{nullity}(T) + \text{rank}(T) = \dim(\ker(T)) + \dim(\text{Im}(T)) = \dim(V) .$$

Given a basis for V the image of T can be obtained from the images of the basis vectors in the following way:

Theorem 1.0.12. *Let V and W be vector spaces and let $T: V \rightarrow W$ be a linear transformation. If $\beta = \{v_1, \dots, v_n\}$ is a basis for V , then*

$$\text{Im}(T) = \text{span}(T(\beta)) = \text{span}\{T(v_1), \dots, T(v_n)\} .$$

Proof. First remember that if a vector space contains a set S of vectors, then it also contains $\text{span}(S)$. Since $T(v_i) \in \text{Im}(T)$ for all $i \in \{1, \dots, n\}$, we obtain

$$\text{span}(\beta) = \text{span}\{T(v_1), \dots, T(v_n)\} \subset \text{Im}(T) .$$

Now suppose that $w \in \text{Im}(T)$. By definition there exists a $v \in V$ with $w = T(v)$. As β is a basis for V , we have that there exist $a_1, \dots, a_n \in F$ with the property that

$$v = a_1v_1 + \dots + a_nv_n .$$

As T is linear, it follows that

$$w = T(v) = a_1T(v_1) + \dots + a_nT(v_n) \in \text{span}(\beta)$$

Since w was arbitrary, this implies that $\text{Im}(T) \subset \text{span}(\beta)$. □

A very important property about a linear transformation is that its action on an arbitrary vector $v \in V$ is completely determined by its action on some chosen basis of V . The following theorem illustrates this property:

Theorem 1.0.13. *Let V be a finite-dimensional vector space and let W be a vector space (both over the field F). Let $\beta = \{v_1, \dots, v_n\}$ be a basis for V . Consider some arbitrary vectors $w_1, \dots, w_n \in W$. Then there exists exactly one linear transformation $T: V \rightarrow W$ such that*

$$T(v_i) = w_i$$

for $i \in \{1, \dots, n\}$.

Proof. For $x \in V$ we have a unique representation with respect to the basis β , i.e. there exist unique scalars $a_1, \dots, a_n \in F$ with the property

$$x = \sum_{i=1}^n a_i v_i .$$

We define $T: V \rightarrow W$ by

$$T(x) = \sum_{i=1}^n a_i w_i .$$

If $x = v_i$, then $a_j = 0$ for all $j \neq i$ and $a_i = 1$. Therefore we obtain $T(v_i) = w_i$ from our definition. We claim that the map T is linear. To see this, let $x \in V$ have a decomposition as above and let $y = \sum_{i=1}^n b_i v_i$. Let $c \in F$. Then $cx + y = \sum_{i=1}^n (ca_i + b_i)v_i$ and

$$T(cx + y) = \sum_{i=1}^n (ca_i + b_i)w_i = c \sum_{i=1}^n a_i w_i + \sum_{i=1}^n b_i w_i = cT(x) + T(y) .$$

To see that T is unique, let U be another linear transformation with the property $U(v_i) = w_i$ for all $i \in \{1, \dots, n\}$. Then for any $x = \sum_{i=1}^n a_i v_i \in V$ we have that

$$U(x) = U\left(\sum_{i=1}^n a_i v_i\right) = \sum_{i=1}^n a_i U(v_i) = \sum_{i=1}^n a_i w_i = T(x) ,$$

where the second equality follows by linearity of U . This proves $U = T$. □

If $T: V \rightarrow W$ is a linear transformation between two finite-dimensional vector spaces **of the same dimension**, then it turns out that injectivity, surjectivity and bijectivity of T actually coincide as the next theorem shows. We leave the proof as an exercise.

Theorem 1.0.14. *Let V and W be vector spaces of finite dimension with*

$$\dim(V) = \dim(W)$$

Suppose $T: V \rightarrow W$ is a linear transformation. Then the following are equivalent:

- a) T is injective (equivalently $\ker(T) = \{0\}$),
- b) T is surjective,
- c) $\dim(\text{Im}(T)) = \dim(V)$.

Exercise 1.0.15. Prove Thm. 1.0.14.

2 Linear Transformations and Matrices

The theory of linear transformations is closely linked to matrices. In fact, every linear transformation $T: V \rightarrow W$ gives rise to a matrix $[T]_{\beta}^{\gamma}$, which describes T completely. It depends on the choice of ordered bases β for V and γ for W . These are defined as follows:

Definition 2.0.1. Let V be a finite-dimensional vector space. An **ordered basis** for V is a basis β for V endowed with a specific order.

Let V and W be finite-dimensional vector spaces over a field F and let $T: V \rightarrow W$ be a linear transformation between them. Let $n = \dim(V)$ and let $m = \dim(W)$. Let $\beta = \{v_1, \dots, v_n\}$ be an ordered basis for V and let $\gamma = \{w_1, \dots, w_m\}$ be an ordered basis for W . Since β and γ are bases, there are unique scalars $a_{ij} \in F$ with $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, n\}$ such that

$$T(v_j) = \sum_{i=1}^m a_{ij} w_i .$$

Definition 2.0.2. We call the $m \times n$ -matrix $A = (a_{ij})_{i,j}$ the **matrix representation of T in the ordered bases β and γ** and write $A = [T]_{\beta}^{\gamma}$. If the bases involved are clear from the context we will call $[T]_{\beta}^{\gamma}$ just the **matrix representation of T** .

To determine the matrix representation of $T: V \rightarrow W$ with respect to the ordered bases $\beta = \{v_1, \dots, v_n\}$ and $\gamma = \{w_1, \dots, w_m\}$ we therefore perform the following steps:

- a) Determine the images of the basis elements in β , i.e. determine $T(v_j)$ for all j .
- b) Write each $T(v_j)$ as a linear combination of elements in γ , i.e. find $a_{ij} \in F$ such that

$$T(v_j) = \sum_{i=1}^m a_{ij} w_i .$$

- c) The coefficient $a_{ij} \in F$ gives the entry in the i th row and j th column of the matrix. In particular, for fixed j the coefficients a_{ij} with $i = 1, \dots, m$ form the j th column of the matrix.

Example 2.0.3. Consider the linear transformation

$$T: P_3(\mathbb{R}) \rightarrow P_2(\mathbb{R}) \quad , \quad T(p) = p' \quad (2)$$

(i.e. taking the first derivative of the polynomial). Let

$$\beta = \{1, x, x^2, x^3\} \quad , \quad \gamma = \{1, 1+x, 1+x^2\} \quad (3)$$

Now we follow the instructions outlined above: First we determine $T(1)$, $T(x)$, $T(x^2)$ and $T(x^3)$ and then write each resulting vector as a linear combination of the basis elements from γ . Finally, the coefficients in front of the elements from γ form the columns of the matrix and we obtain:

$$\left. \begin{aligned} T(1) = 0 &= 0 \cdot 1 + 0 \cdot (1+x) + 0 \cdot (1+x^2) \\ T(x) = 1 &= \boxed{1} \cdot 1 + \boxed{0} \cdot (1+x) + \boxed{0} \cdot (1+x^2) \\ T(x^2) = 2x &= -2 \cdot 1 + 2 \cdot (1+x) + 0 \cdot (1+x^2) \\ T(x^3) = 3x^2 &= -3 \cdot 1 + 0 \cdot (1+x) + 3 \cdot (1+x^2) \end{aligned} \right\} [T]_{\beta}^{\gamma} = \begin{pmatrix} 0 & \boxed{1} & -2 & -3 \\ 0 & \boxed{0} & 2 & 0 \\ 0 & \boxed{0} & 0 & 3 \end{pmatrix}$$

We will see throughout the whole course that these matrices provide very versatile tools to study linear transformations. In particular, the composition of linear transformations turns into a multiplication operation for matrices.

But first we will take a closer look at the set of all linear transformations $\mathcal{L}(V, W)$ between two vector spaces V and W over the same field F . In particular, we will prove that $\mathcal{L}(V, W)$ is in fact a vector space itself. To do so we must first define the addition of two linear maps and the multiplication of a linear map by a scalar.

Definition 2.0.4. Let V and W be vector spaces over the field F . Let $T, U: V \rightarrow W$ be linear transformations and let $a \in F$. We define the map $T + U$ to be

$$(T + U)(v) = T(v) + U(v)$$

for all $v \in V$ and we define aT to be

$$(aT)(v) = aT(v)$$

for all $v \in V$.

Remark 2.0.5. In fact, the last definition works in much greater generality: Whenever we have arbitrary maps $f, g: X \rightarrow W$ from an arbitrary set X to a vector space W , we can define the sum and scalar multiplication in the way described above, i.e. $(f+g)(x) = f(x) + g(x)$ and $(af)(x) = af(x)$.

Theorem 2.0.6. *Let V and W be vector spaces over the field F and let $T, U: V \rightarrow W$ be linear transformations.*

- a) *For all $a \in F$, $aT + U$ is again a linear transformation.*
 b) *Using the operations from Def. 2.0.4 and the zero map as a zero element, the collection of all linear transformations from V to W is a vector space over F .*

Proof. To prove a) we need to check that $(aT + U)(cx + y) = c(aT + U)(x) + (aT + U)(y)$ for all $c \in F$ and $x, y \in V$ (see Lem. 1.0.10).

$$\begin{aligned}
 (aT + U)(cx + y) &= (aT)(cx + y) + U(cx + y) \\
 &= a(T(cx + y)) + U(cx + y) \\
 &= a(cT(x) + T(y)) + cU(x) + U(y) \\
 &= acT(x) + cU(x) + aT(y) + U(y) \\
 &= c(aT(x) + U(x)) + (aT + U)(y) \\
 &= c(aT + U)(x) + (aT + U)(y)
 \end{aligned}$$

In the first two lines we have used the definition of $aT + U$. In the third we used the linearity of T and U . In the fourth line we just reshuffled the terms and worked our way back using the definition of $aT + U$ again in lines 5 and 6.

We leave the proof of b) as an exercise (see below) and just verify two of the axioms here: Let T_0 be the zero map. To prove that VS 3) is satisfied we need to check that $T + T_0 = T$ for all linear maps $T: V \rightarrow W$. This means we need to check that $(T + T_0)(v) = T(v)$ for all $v \in V$, but this is true, since $T_0(v) = 0$ and $T(v) + 0 = T(v)$. To show that VS 4) holds, we need to find for a given linear map $T: V \rightarrow W$ another linear map $U: V \rightarrow W$ such that $T + U = T_0$. Let $U = -T = (-1)T$, then $(T + U)(v) = T(v) + U(v) = T(v) - T(v) = 0$ and therefore $T + U = T_0$. \square

Exercise 2.0.7. Fill in the remaining details in the proof of Thm. 2.0.6 b).

Definition 2.0.8. Let V and W be vector spaces over the field F . We denote the vector space of all linear transformations from V to W by $\mathcal{L}(V, W)$. In the case that $W = V$ we also write $\mathcal{L}(V)$ instead of $\mathcal{L}(V, V)$.

If we are given three vector spaces V_1, V_2 and V_3 together with linear maps $T: V_1 \rightarrow V_2$, $U: V_2 \rightarrow V_3$ we can form their composition $U \circ T: V_1 \rightarrow V_3$, which is defined by first applying T to a vector in V_1 and then U to the resulting vector in V_2 . We will also denote this composition by UT . The following theorem shows that UT is again a linear transformation.

Theorem 2.0.9. *Let V_1, V_2 and V_3 be vector spaces over the field F . Let $T: V_1 \rightarrow V_2$ and $U: V_2 \rightarrow V_3$ be linear transformations. Then their composition $UT = U \circ T$ is also a linear transformation.*

Proof. Let $x, y \in V_1$ and $a \in F$. We have to show that $UT(ax + y) = aUT(x) + UT(y)$ (see Lem. 1.0.10). This follows from the computation:

$$\begin{aligned} UT(ax + y) &= U(T(ax + y)) = U(aT(x) + T(y)) \\ &= aU(T(x)) + U(T(y)) = aUT(x) + UT(y) . \end{aligned} \quad \square$$

The following theorem summarises the most important properties of the composition of linear transformations:

Theorem 2.0.10. *Let V_i for $i \in \{1, \dots, 4\}$ be vector spaces over the field F .*

a) *If $T: V_1 \rightarrow V_2$ and $U_1, U_2: V_2 \rightarrow V_3$ are linear transformations, then*

$$(U_1 + U_2)T = U_1T + U_2T .$$

Likewise, if $T_1, T_2: V_1 \rightarrow V_2$ and $U: V_2 \rightarrow V_3$ are linear transformations, then

$$U(T_1 + T_2) = UT_1 + UT_2 .$$

b) *If $T_3: V_1 \rightarrow V_2$, $T_2: V_2 \rightarrow V_3$ and $T_1: V_3 \rightarrow V_4$ are linear transformations, then*

$$(T_1T_2)T_3 = T_1(T_2T_3) .$$

c) *If $T: V \rightarrow W$ is a linear transformation and $I_V: V \rightarrow V$, $I_W: W \rightarrow W$ are the identities on V , respectively W , then $TI_V = I_WT = T$.*

d) *If $T_2: V_1 \rightarrow V_2$, $T_1: V_2 \rightarrow V_3$ are linear transformations and $a \in F$, then*

$$a(T_1T_2) = (aT_1)T_2 = T_1(aT_2) .$$

Proof. We will only prove the first result and leave the rest as an exercise. Let T, U_1 and U_2 be linear transformations as in the first part of a) and let $x \in V_1$. Then

$$((U_1 + U_2)T)(x) = (U_1 + U_2)(T(x)) = U_1(T(x)) + U_2(T(x)) = U_1T(x) + U_2T(x) .$$

Likewise, let T_1, T_2 and U be as in the second part of a). Let $x \in V_1$. Then

$$\begin{aligned} (U(T_1 + T_2))(x) &= U((T_1 + T_2)(x)) = U(T_1(x) + T_2(x)) \\ &= U(T_1(x)) + U(T_2(x)) = UT_1(x) + UT_2(x) . \end{aligned}$$

All other statements follow in a similar fashion by evaluating both sides on a vector $x \in V_1$ and checking that they agree using the linearity of the maps involved. \square

Exercise 2.0.11. Prove the statements b), c) and d) in Theorem 2.0.10.

Remark 2.0.12. The vector space $\mathcal{L}(V)$ is quite special: We can compose any two elements $T, S \in \mathcal{L}(V)$ and obtain another element $TS \in \mathcal{L}(V)$. In particular, the composition yields a map

$$\circ: \mathcal{L}(V) \times \mathcal{L}(V) \rightarrow \mathcal{L}(V) \quad , \quad (T, S) \mapsto TS . \quad (4)$$

The statements in Thm. 2.0.10 a), b) and d) imply that this map is linear if we keep either T or S fixed. Such a map is called bilinear. Moreover, Thm. 2.0.10 b) states that the composition operation on $\mathcal{L}(V)$ is associative as well. To summarise: $\mathcal{L}(V)$ is a vector space with a bilinear, associative multiplication (4) and a unit element $I_V \in \mathcal{L}(V)$ (see Thm. 2.0.10 c)). In this case we also say that $\mathcal{L}(V)$ has the structure of a *unital algebra*.

As we have seen in Def. 2.0.2 after choosing ordered bases of V and W we can associate a matrix to any linear transformation $T: V \rightarrow W$. The following theorem shows that this operation is well-behaved with respect to taking sums and scalar multiples of linear transformations:

Theorem 2.0.13. *Let V and W be finite-dimensional vector spaces over a field F with ordered bases β and γ , respectively. Let $T, U: V \rightarrow W$ be linear transformations and let $a \in F$. Then*

$$a) \quad [T + U]_{\beta}^{\gamma} = [T]_{\beta}^{\gamma} + [U]_{\beta}^{\gamma},$$

$$b) \quad [aT]_{\beta}^{\gamma} = a [T]_{\beta}^{\gamma}.$$

Proof. Let $\beta = \{v_1, \dots, v_n\}$ and let $\gamma = \{w_1, \dots, w_m\}$. By definition of the matrix representation

$$T(v_j) = \sum_{i=1}^m a_{ij} w_i \quad \text{and} \quad U(v_j) = \sum_{i=1}^m b_{ij} w_i$$

for all $1 \leq j \leq n$. It follows that

$$(T + U)(v_j) = \sum_{i=1}^m (a_{ij} + b_{ij}) w_i$$

for all $1 \leq j \leq n$. From this we see that

$$\left([T + U]_{\beta}^{\gamma}\right)_{ij} = a_{ij} + b_{ij} = \left([T]_{\beta}^{\gamma} + [U]_{\beta}^{\gamma}\right)_{ij}$$

Part b) of the theorem follows in a similar way. □

2.1 Composition and matrix multiplication

As we discussed above, if two linear transformations are composed, the result is again linear (see Thm. 2.0.9). Likewise, matrices can be multiplied to give new matrices. This works as follows:

Let F be a field (e.g. $F = \mathbb{R}$ or $F = \mathbb{C}$). Consider two matrices $A \in M_{l \times m}(F)$ and $B \in M_{m \times n}(F)$. Note that the number of columns of A is equal to the number of rows in B . If this is the case, the two matrices can be **multiplied** (which is usually denoted by a dot just like ordinary multiplication) to give a matrix $A \cdot B \in M_{l \times n}(F)$. If A has entries $a_{ij} \in F$ and B has entries $b_{ij} \in F$, then the product $A \cdot B$ will have entries c_{ij} defined by

$$c_{ij} = \sum_{k=1}^m a_{ik}b_{kj} . \quad (5)$$

For the multiplication of a general 2×3 -matrix with a 3×2 -matrix this is illustrated below:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} + a_{12}b_{21} + a_{13}b_{31} & a_{11}b_{12} + a_{12}b_{22} + a_{13}b_{32} \\ a_{21}b_{11} + a_{22}b_{21} + a_{23}b_{31} & a_{21}b_{12} + a_{22}b_{22} + a_{23}b_{32} \end{pmatrix}$$

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} + a_{12}b_{21} + a_{13}b_{31} & a_{11}b_{12} + a_{12}b_{22} + a_{13}b_{32} \\ a_{21}b_{11} + a_{22}b_{21} + a_{23}b_{31} & a_{21}b_{12} + a_{22}b_{22} + a_{23}b_{32} \end{pmatrix}$$

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} + a_{12}b_{21} + a_{13}b_{31} & a_{11}b_{12} + a_{12}b_{22} + a_{13}b_{32} \\ a_{21}b_{11} + a_{22}b_{21} + a_{23}b_{31} & a_{21}b_{12} + a_{22}b_{22} + a_{23}b_{32} \end{pmatrix}$$

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \\ b_{31} & b_{32} \end{pmatrix} = \begin{pmatrix} a_{11}b_{11} + a_{12}b_{21} + a_{13}b_{31} & a_{11}b_{12} + a_{12}b_{22} + a_{13}b_{32} \\ a_{21}b_{11} + a_{22}b_{21} + a_{23}b_{31} & a_{21}b_{12} + a_{22}b_{22} + a_{23}b_{32} \end{pmatrix}$$

The entries of $A \cdot B$ defined in (5) can be obtained from the three rules listed below:

- The rows of the first matrix are multiplied with the columns of the second.
- This works as follows: Multiply the k th entry of the row with the k th entry of the column and add up the results.
- When the i th row of the first matrix is multiplied with the j th column of the second in this way, the result is the entry in the i th row and j th column of the result.

Remark 2.1.1. Even though the multiplication of matrices is associative in the sense that $(A \cdot B) \cdot C = A \cdot (B \cdot C)$, whenever the product is defined, the multiplication is not commutative, i.e. in general $A \cdot B \neq B \cdot A$. If it is clear from the context that the matrices involved should be multiplied, then the dot is sometimes omitted, i.e. $A \cdot B$ would be written as AB .

As we will see in the next theorem, the matrix of the composition of two linear maps corresponds to the product of the matrices associated to the linear maps themselves if the bases are chosen in a compatible way.

Theorem 2.1.2. *Let V, W and Z be finite-dimensional vector spaces over the field F with ordered bases α, β and γ , respectively. Let $T: V \rightarrow W$ and $U: W \rightarrow Z$ be linear transformations. Then*

$$[UT]_{\alpha}^{\gamma} = [U]_{\beta}^{\gamma} \cdot [T]_{\alpha}^{\beta} .$$

Proof. Let $n = \dim(V)$, $m = \dim(W)$ and $p = \dim(Z)$. Let $\alpha = \{v_1, \dots, v_n\}$, $\beta = \{w_1, \dots, w_m\}$ and $\gamma = \{z_1, \dots, z_p\}$ be the ordered bases. Let $A = [U]_{\beta}^{\gamma}$ and $B = [T]_{\alpha}^{\beta}$. By the definition of the matrix representation (see Def. 2.0.2)

$$T(v_j) = \sum_{k=1}^m B_{kj} w_k \quad \text{and} \quad U(w_j) = \sum_{i=1}^p A_{ij} z_i .$$

For $1 \leq j \leq n$, we have

$$\begin{aligned} (UT)(v_j) &= U(T(v_j)) = U\left(\sum_{k=1}^m B_{kj} w_k\right) = \sum_{k=1}^m B_{kj} U(w_k) \\ &= \sum_{k=1}^m B_{kj} \left(\sum_{i=1}^p A_{ik} z_i\right) = \sum_{i=1}^p \left(\sum_{k=1}^m A_{ik} B_{kj}\right) z_i = \sum_{i=1}^p C_{ij} z_i \end{aligned}$$

where $C_{ij} = \sum_{k=1}^m A_{ik} B_{kj}$. Observe that C agrees with the result of the matrix multiplication of A and B , i.e. $C = AB$ and therefore $[UT]_{\alpha}^{\gamma} = [U]_{\beta}^{\gamma} [T]_{\alpha}^{\beta}$. \square

Example 2.1.3. Let $T: P_3(\mathbb{R}) \rightarrow P_2(\mathbb{R})$ be the linear transformation given by differentiation, i.e. $T(f)(x) = f'(x)$. Let $\beta = \{1, x, x^2, x^3\}$ and $\gamma = \{1, x, x^2\}$. The matrix representation of T with respect to the ordered bases β and γ is given by

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

Let $U: P_2(\mathbb{R}) \rightarrow P_3(\mathbb{R})$ be the linear transformation given by integration, i.e.

$$U(f)(x) = \int_0^x f(y) dy .$$

Its matrix representation with respect to γ and β is given by

$$[U]_{\gamma}^{\beta} = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix} .$$

Now consider the linear transformation $TU: P_2(\mathbb{R}) \rightarrow P_2(\mathbb{R})$ obtained as the composition of the two maps. By Theorem 2.1.2 its matrix representation with respect to γ satisfies

$$\begin{aligned} [TU]_\gamma^\gamma &= [T]_\beta^\gamma [U]_\gamma^\beta \\ &= \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = [I_{P_2(\mathbb{R})}]_\gamma^\gamma \end{aligned}$$

where $I_{P_2(\mathbb{R})}$ is the identity transformation. Indeed, we can compute TU directly and obtain

$$TU(a_2x^2 + a_1x + a_0) = T\left(\frac{1}{3}a_2x^3 + \frac{1}{2}a_1x^2 + a_0x\right) = a_2x^2 + a_1x + a_0 .$$

Hence, TU is indeed the identity on the vector space $P_2(\mathbb{R})$. What about the other composition, i.e. $UT: P_3(\mathbb{R}) \rightarrow P_3(\mathbb{R})$? Using Thm. 2.1.2 again we obtain

$$\begin{aligned} [UT]_\beta^\beta &= [U]_\gamma^\beta [T]_\beta^\gamma \\ &= \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix} \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{aligned}$$

which is **not** the identity matrix! In fact, if we perform the operations of differentiation and integration in this order, the constant polynomials will be in the null space of UT , because $T(a_0) = a_0T(1) = 0$ for any $a_0 \in F$.

This example also highlights the importance of Thm. 2.1.2 for applications: Suppose we wanted to teach a computer how to differentiate and integrate polynomials of fixed finite degree $d \in \mathbb{N}$. All we have to do is find the matrix of this linear transformation in a suitable basis, like the monomials x^k for $k \in \{0, \dots, d\}$. Performing multiple differentiation and integration operations on a polynomial will then boil down to computing the product of the corresponding matrices, which is something that is easily implemented in any programming language.

Given a vector space V of dimension $n = \dim(V) < \infty$ over a field F with ordered basis β we can associate to any $x \in V$ its coordinate vector $[x]_\beta \in F^n$.

Definition 2.1.4. Let $\beta = \{u_1, \dots, u_n\}$ be an ordered basis for a finite-dimensional vector space V over the field F . For $x \in V$, let $a_i \in F$ for $i \in \{1, \dots, n\}$ be the unique scalars given by

$$x = \sum_{i=1}^n a_i u_i .$$

We define the **coordinate vector of x with respect to β** , denoted by $[x]_\beta$, by

$$[x]_\beta = \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \in F^n .$$

The following theorem shows how these coordinate vectors behave under linear transformations.

Theorem 2.1.5. *Let V and W be finite-dimensional vector spaces over the field F with ordered bases β and γ , respectively, and let $T: V \rightarrow W$ be a linear transformation. Then for each $u \in V$ we have*

$$[T(u)]_\gamma = [T]_\beta^\gamma [u]_\beta .$$

Proof. Fix $u \in V$ and define the linear transformation $f: F \rightarrow V$ by $f(a) = au$ for all $a \in F$. Let $g: F \rightarrow W$ be the linear transformation given by $g(a) = aT(u)$ for all $a \in F$. Observe that $Tf = g$. Let $\alpha = \{1\} \subset F$ and consider α as an ordered basis for the vector space F over F . Let $n = \dim(V)$. We can think of the coordinate vector $[u]_\beta$ as an $n \times 1$ -matrix. With this identification we have $[g(1)]_\gamma = [g]_\alpha^\gamma$ and $[f(1)]_\beta = [f]_\alpha^\beta$. Using Thm. 2.1.2 we obtain

$$[T(u)]_\gamma = [g(1)]_\gamma = [g]_\alpha^\gamma = [Tf]_\alpha^\gamma = [T]_\beta^\gamma [f]_\alpha^\beta = [T]_\beta^\gamma [f(1)]_\beta = [T]_\beta^\gamma [u]_\beta . \quad \square$$

Example 2.1.6. Let us illustrate Thm. 2.1.5 by looking at the linear transformation $T: P_3(\mathbb{R}) \rightarrow P_2(\mathbb{R})$ given by differentiation, i.e. $Tf(x) = f'(x)$. Let $\beta = \{1, x, x^2, x^3\}$ and $\gamma = \{1, x, x^2\}$ be the standard ordered bases for $P_3(\mathbb{R})$ and $P_2(\mathbb{R})$, respectively. Let $A = [T]_\beta^\gamma$. We have that

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix} .$$

Let $f(x) = 3x^3 - x^2 + 5x + 7$. Then $f \in P_3(\mathbb{R})$ and

$$[f]_\beta = \begin{pmatrix} 7 \\ 5 \\ -1 \\ 3 \end{pmatrix} .$$

Calculating Tf directly we obtain $Tf(x) = 9x^2 - 2x + 5$ with coordinate vector

$$[Tf]_\gamma = \begin{pmatrix} 5 \\ -2 \\ 9 \end{pmatrix} .$$

According to Thm. 2.1.5 we should obtain the same result for $A[f]_\beta$. Indeed we have

$$[T]_\beta^\gamma [f]_\beta = A [f]_\beta = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} 7 \\ 5 \\ -1 \\ 3 \end{pmatrix} = \begin{pmatrix} 5 \\ -2 \\ 9 \end{pmatrix} = [Tf]_\gamma .$$

The two Theorems 2.1.2 and 2.1.5 are crucial in Linear Algebra. They tell us that any calculation with linear maps can be reduced to a calculation that only involves matrices

and coordinate vectors after choosing appropriate bases for the vector spaces involved. In particular, any Linear Algebra problem that only involves finite-dimensional vector spaces can in principle be solved by a computer program.

As we have seen we can associate a matrix to any linear transformation between finite-dimensional vector spaces after choosing ordered bases. We will end this chapter by discussing in which way any matrix gives rise to a linear transformation itself.

Definition 2.1.7. Let A be an $m \times n$ -matrix with entries from a field F . We denote by L_A the linear transformation given by

$$L_A: F^n \rightarrow F^m \quad , \quad x \mapsto Ax .$$

We call L_A the **left-multiplication transformation**.

Exercise 2.1.8. Check that L_A is indeed a linear transformation.

Example 2.1.9. Let

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 1 & 0 & 1 \end{pmatrix} .$$

Then $A \in M_{2 \times 3}(\mathbb{R})$. Therefore the left-multiplication transformation corresponding to A is a map $L_A: \mathbb{R}^3 \rightarrow \mathbb{R}^2$. If

$$x = \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix}$$

then

$$L_A(x) = Ax = \begin{pmatrix} 1 & 2 & 3 \\ 1 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 \\ -1 \\ 2 \end{pmatrix} = \begin{pmatrix} 5 \\ 3 \end{pmatrix} .$$

3 Invertibility of Linear Transformations

If two sets X and Y have the same number of elements, then there are maps $f: X \rightarrow Y$ and $g: Y \rightarrow X$ with the property that $f \circ g = I_Y$ and $g \circ f = I_X$. In this case we say that f is a bijection between the two sets (so is g). If we do not want to mention the maps, we might also say that the sets are in bijection. The maps f and g allow us to compare different sets and it turns out that the number of elements, the cardinality, distinguishes between sets that are not in bijection.

In this section we will study the analogous question not for sets, but for vector spaces. Of course each vector space V over any field F is in particular also a set and we could simply look for bijections of these underlying sets. But this neglects all of the additional structure (like the addition and the scalar multiplication) and is therefore not a good way to compare vector spaces.

On the other hand we have already met a natural notion of map between vector spaces, namely that of linear transformations. By definition these transformations are well-behaved with respect to addition and scalar multiplication. Therefore we make the following definition:

Definition 3.0.1. Let V and W be vector spaces (over a field F) and let $T: V \rightarrow W$ be a linear transformation. A map $U: W \rightarrow V$ is said to be an **inverse** of T if $TU = I_W$ and $UT = I_V$. If T has an inverse, then T is said to be **invertible**. In this case T is also called an **isomorphism**. Two vector spaces V, W are called **isomorphic** if there exists an isomorphism $T: V \rightarrow W$.

Note that the definition only requires T to be a linear map, whereas it seems that the inverse U can be any map. As we will see in the next theorem, the inverse U will automatically be a linear transformation as well. In particular, any bijective linear map $T: V \rightarrow W$ is an isomorphism.

Theorem 3.0.2. *Let V and W be vector spaces over the field F and let $T: V \rightarrow W$ be a linear transformation. Suppose that $U: W \rightarrow V$ is an inverse of T . Then U is a linear transformation.*

Proof. Let $y_1, y_2 \in W$ and let $c \in F$. Let $x_1 = U(y_1)$ and $x_2 = U(y_2)$. Since U is an inverse of T we have $T(x_1) = T(U(y_1)) = y_1$ and likewise $T(x_2) = y_2$. We have

$$U(cy_1 + y_2) = U(cT(x_1) + T(x_2)) = U(T(cx_1 + x_2)) = cx_1 + x_2 = cU(y_1) + U(y_2) ,$$

where we used that T is a linear transformation. This proves that U is a linear map (see Lem. 1.0.10). \square

The inverse of a linear transformation is unique. From now on we will use the notation T^{-1} for the inverse of T . We have seen that there is a close connection between linear transformations and matrices. Therefore it makes sense to define invertibility for matrices as well.

Definition 3.0.3. Let F be a field and let $n \in \mathbb{N}$. An $n \times n$ -matrix $A \in M_{n \times n}(F)$ is called **invertible** if there exists an $n \times n$ -matrix $B \in M_{n \times n}(F)$ with the property that

$$AB = I \quad \text{and} \quad BA = I .$$

Example 3.0.4. Consider the linear transformation T defined as follows

$$T: \mathbb{R}^3 \rightarrow P_2(\mathbb{R}) \quad , \quad (a_0, a_1, a_2) \mapsto a_2x^2 + a_1x + a_0 .$$

Note that a polynomial is uniquely fixed by the coefficients $a_0, a_1, a_2 \in \mathbb{R}$. Therefore T is invertible and the inverse $U: P_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ is given by

$$U: P_2(\mathbb{R}) \rightarrow \mathbb{R}^3 \quad , \quad a_2x^2 + a_1x + a_0 \mapsto (a_0, a_1, a_2) .$$

Note that

$$\begin{aligned} T(U(a_2x^2 + a_1x + a_0)) &= T(a_0, a_1, a_2) = a_2x^2 + a_1x + a_0 , \\ U(T(a_0, a_1, a_2)) &= U(a_2x^2 + a_1x + a_0) = (a_0, a_1, a_2) . \end{aligned}$$

Therefore T is invertible with $T^{-1} = U$ and $P_2(\mathbb{R})$ is isomorphic to \mathbb{R}^3 .

Example 3.0.5. Consider the two matrices $A, B \in M_{2 \times 2}(\mathbb{R})$ given by

$$A = \begin{pmatrix} 1 & 2 \\ 0 & 2 \end{pmatrix} \quad , \quad B = \begin{pmatrix} 1 & -1 \\ 0 & \frac{1}{2} \end{pmatrix} .$$

They satisfy $AB = I_2$ and $BA = I_2$. In particular, A is invertible with inverse B .

Example 3.0.6. Let F be a field. As we will see later a matrix $A \in M_{2 \times 2}(F)$ given by

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

is invertible if and only if $ad - bc \neq 0$. In this case the inverse of A is given by

$$A^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} = \begin{pmatrix} \frac{d}{ad-bc} & \frac{-b}{ad-bc} \\ \frac{-c}{ad-bc} & \frac{a}{ad-bc} \end{pmatrix}$$

It is a good exercise to check that this matrix indeed satisfies the defining condition

$$A \cdot A^{-1} = I_2 = A^{-1} \cdot A .$$

Lemma 3.0.7. *Let V and W be vector spaces over the field F and let $T: V \rightarrow W$ be an isomorphism between them. Then V is finite-dimensional if and only if W is finite-dimensional. In this case, $\dim(V) = \dim(W)$.*

Proof. Suppose V is finite-dimensional, let $n = \dim(V)$ and $\beta = \{x_1, \dots, x_n\}$ be a basis for V . By Thm. 1.0.12 we have $\text{span}(T(\beta)) = \text{Im}(T)$. In particular, we have found a finite spanning set for $\text{Im}(T)$. Since any basis for $\text{Im}(T)$ has at most as many elements as any of its spanning sets, $\text{Im}(T)$ is finite-dimensional. But since T is an isomorphism, it is bijective. Hence, we have $\text{Im}(T) = W$. Therefore W is finite-dimensional. Conversely, if W is finite-dimensional then so is V by a similar argument using T^{-1} .

Now suppose that V and W are finite-dimensional. Since T is bijective, we have $\ker(T) = \{0\}$ and $\text{Im}(T) = W$. By the rank-nullity theorem we have

$$\dim(V) = \dim(\text{Im}(T)) + \dim(\ker(T)) = \dim(W) + 0 = \dim(W) . \quad \square$$

Remark 3.0.8. From the above theorem we see that if V and W are isomorphic finite-dimensional vector spaces, then we must have $\dim(V) = \dim(W)$. What about the converse statement? Is having $\dim(V) = \dim(W)$ enough to deduce that V and W are isomorphic? Remember that to prove such a statement we need to construct an isomorphism $T: V \rightarrow W$. If you paid close attention in the chapter about bases you might already have a clue how to construct such a linear transformation, but we will revisit this later.

Example 3.0.9. Let $V = P_2(\mathbb{R})$ and let $W = \mathbb{R}^2$. From the isomorphism we constructed in Example 3.0.4 we know that $\dim(V) = 3$. Moreover, $\dim(W) = 2$. Hence, both spaces are finite-dimensional. If they were isomorphic, we would have $\dim(V) = \dim(W)$ according to Lemma 3.0.7, which is not the case. Therefore V and W are **not isomorphic**.

Choosing bases for finite-dimensional vector spaces V and W we can associate a matrix to each linear map $T: V \rightarrow W$. Invertibility of T should be closely related to the invertibility of the corresponding matrix. The following theorem makes this precise.

Theorem 3.0.10. *Let V and W be finite-dimensional vector spaces over the field F . Let β be an ordered basis for V and let γ be an ordered basis for W . Let $T: V \rightarrow W$ be a linear transformation. Then T is an isomorphism if and only if $[T]_{\beta}^{\gamma}$ is invertible. In this case,*

$$[T^{-1}]_{\gamma}^{\beta} = \left([T]_{\beta}^{\gamma}\right)^{-1} .$$

Proof. First suppose that T is invertible. By Lemma 3.0.7 we have $\dim(V) = \dim(W)$. Let $n = \dim(V)$. So $[T]_{\beta}^{\gamma}$ is an $n \times n$ -matrix. Note that $T^{-1}: W \rightarrow V$ satisfies $TT^{-1} = I_W$ and $T^{-1}T = I_V$. Thus

$$I_n = [I_V]_{\beta}^{\beta} = [T^{-1}T]_{\beta}^{\beta} = [T^{-1}]_{\gamma}^{\beta} [T]_{\beta}^{\gamma} .$$

Likewise, $[T]_{\beta}^{\gamma} [T^{-1}]_{\gamma}^{\beta} = I_n$. Hence $[T]_{\beta}^{\gamma}$ is invertible and

$$[T^{-1}]_{\gamma}^{\beta} = \left([T]_{\beta}^{\gamma}\right)^{-1} .$$

Now suppose that the matrix $A = [T]_{\beta}^{\gamma}$ is invertible. By Definition 3.0.3 this means that there exists an $n \times n$ -matrix B such that $AB = I_n$ and $BA = I_n$. In particular, A has to be an $n \times n$ -matrix itself. Therefore $\dim(V) = \dim(W) = n$, since the number of columns of A is the dimension of V and the number of rows is the dimension of W . Let $w_i \in W$ for $i \in \{1, \dots, n\}$ be the vectors in the basis γ , i.e. $\gamma = \{w_1, w_2, \dots, w_n\}$ and let $v_i \in V$ for $i \in \{1, \dots, n\}$ be such that $\beta = \{v_1, v_2, \dots, v_n\}$. Using B and the elements of γ and β we can now construct a linear transformation $U: W \rightarrow V$, which satisfies

$$U(w_j) = \sum_{i=1}^n B_{ij} v_i$$

for all $j \in \{1, \dots, n\}$. By construction $[U]_{\gamma}^{\beta} = B$. We have to check that U is the inverse of T , i.e. that $UT = I_V$ and $TU = I_W$. Observe that

$$[UT]_{\beta}^{\beta} = [U]_{\gamma}^{\beta} [T]_{\beta}^{\gamma} = BA = I_n .$$

This means that $UT: V \rightarrow V$ is a linear transformation, which is the identity on the basis elements v_i for all $i \in \{1, \dots, n\}$. Using the linearity of UT we see that it has to be the identity on all vectors in V . A similar argument shows that $TU = I_W$. Therefore T is an isomorphism with inverse $T^{-1} = U$. \square

Exercise 3.0.11. Let $T: V \rightarrow W$ and $U: W \rightarrow V$ be the linear transformations from the proof of Theorem 3.0.10. Fill in the details of the proof that $TU = I_W$.

Exercise 3.0.12. Let A be an $n \times n$ -matrix with entries in the field F . Let $L_A: F^n \rightarrow F^n$ be the linear transformation given by $L_A(v) = Av$. Show that L_A is invertible if and only if A is invertible and that in this case $(L_A)^{-1} = L_{A^{-1}}$.

Example 3.0.13. Consider the linear transformation $T: P_2(\mathbb{R}) \rightarrow \mathbb{R}^3$ given by

$$T: P_2(\mathbb{R}) \rightarrow \mathbb{R}^3 \quad , \quad f \mapsto (f(0), f(1), f(2)) .$$

The ordered set $\beta = \{1, x, x^2\}$ is an ordered basis for $P_2(\mathbb{R})$. Likewise, the ordered set $\gamma = \{e_1, e_2, e_3\}$ is an ordered basis for \mathbb{R}^3 . Evaluating T on the basis vectors in β gives the following results

$$\begin{aligned} T(1) &= \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = e_1 + e_2 + e_3 , \\ T(x) &= \begin{pmatrix} 0 \\ 1 \\ 2 \end{pmatrix} = e_2 + 2e_3 , \\ T(x^2) &= \begin{pmatrix} 0 \\ 1 \\ 4 \end{pmatrix} = e_2 + 4e_3 . \end{aligned}$$

Therefore the matrix representation of T in the ordered bases β and γ is

$$[T]_{\beta}^{\gamma} = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 2 & 4 \end{pmatrix} .$$

We leave it as an exercise to check that the matrix

$$B = \begin{pmatrix} 1 & 0 & 0 \\ -\frac{3}{2} & 2 & -\frac{1}{2} \\ \frac{1}{2} & -1 & \frac{1}{2} \end{pmatrix}$$

is an inverse of $[T]_{\beta}^{\gamma}$. We obtain from Theorem 3.0.10 that T is invertible and that B is the matrix representation of T^{-1} in the ordered bases γ and β . In particular, we see that $P_2(\mathbb{R})$ is isomorphic to \mathbb{R}^3 .

We are now in the position to answer the questions posed in Remark 3.0.8.

Theorem 3.0.14. *Let V and W be finite-dimensional vector spaces over the field F . Then V is isomorphic to W if and only if $\dim(V) = \dim(W)$.*

Proof. First suppose that V and W are isomorphic. This means that there exists an isomorphism $T: V \rightarrow W$. By Lemma 3.0.7 we have that $\dim(V) = \dim(W)$.

Now suppose that $\dim(V) = \dim(W)$ and let $n = \dim(V)$. Let $\beta = \{v_1, \dots, v_n\}$ be an ordered basis for V and let $\gamma = \{w_1, \dots, w_n\}$ be an ordered basis for W . By Thm. 1.0.13 there exists a linear transformation $T: V \rightarrow W$ with the property that

$$T(v_i) = w_i$$

for all $i \in \{1, \dots, n\}$. Using Thm. 1.0.12, we have

$$\text{Im}(T) = \text{span}(T(\beta)) = \text{span}(\gamma) = W ,$$

which implies that T is surjective. Using the rank-nullity theorem, $\dim(\text{Im}(T)) = \dim(W)$ and $\dim(V) = \dim(W)$ we obtain

$$\begin{aligned} \dim(V) &= \dim(\text{Im}(T)) + \dim(\ker(T)) = \dim(W) + \dim(\ker(T)) = \dim(V) + \dim(\ker(T)) \\ &\Rightarrow \dim(\ker(T)) = 0 . \end{aligned}$$

Therefore $\ker(T) = \{0\}$ and T is injective as well. Thus, T is a bijective linear transformation and therefore an isomorphism. \square

Remark 3.0.15. A word of caution is in order here: If V and W are isomorphic vector spaces, then this means that **there exists** a bijective linear transformation $T: V \rightarrow W$. It is not true, that all linear transformations $V \rightarrow W$ are bijective. For example, the zero map $Z: V \rightarrow W$ with $Z(v) = 0$ is a linear transformation. However, Z is only an isomorphism if V and W are both the zero vector space, i.e. if we have $\dim(V) = \dim(W) = 0$. Otherwise Z is never bijective.

Corollary 3.0.16. *Let V be a vector space over a field F . Then V is isomorphic to F^n if and only if $\dim(V) = n$.*

As we know from Theorem 2.0.6 the linear transformations between two vector spaces over F form a new vector space over F . Furthermore, we can associate a matrix to each such linear transformation after choosing ordered bases for the vector spaces involved. In the next theorem we extend this observation to the statement that this association yields an isomorphism between the vector space of linear transformations and the corresponding vector space of matrices.

Theorem 3.0.17. *Let V and W be finite-dimensional vector spaces over a field F , let $n = \dim(V)$, $m = \dim(W)$ and let β and γ be ordered bases for V and W , respectively. Then the map $\Phi: \mathcal{L}(V, W) \rightarrow M_{m \times n}(F)$, defined by $\Phi(T) = [T]_{\beta}^{\gamma}$ for $T \in \mathcal{L}(V, W)$ is an isomorphism.*

Proof. By Thm. 2.0.13 the map Φ is a linear transformation. Therefore it suffices to show that Φ is bijective. This is accomplished if we show that for every $m \times n$ -matrix A there is a unique linear transformation $T: V \rightarrow W$ such that $\Phi(T) = A$. Let $\beta = \{v_1, \dots, v_n\}$, $\gamma = \{w_1, \dots, w_m\}$ and let A be a given $m \times n$ -matrix. By Thm. 1.0.13 there exists a unique linear transformation $T: V \rightarrow W$ with the property that

$$T(v_j) = \sum_{i=1}^m A_{ij} w_i$$

for $1 \leq j \leq n$. But then we have $[T]_\beta^\gamma = A$ by construction and therefore $\Phi(T) = A$. Thus Φ is an isomorphism. \square

Corollary 3.0.18. *Let V and W be finite-dimensional vector spaces over F of dimension n and m , respectively. Then $\mathcal{L}(V, W)$ is finite-dimensional and*

$$\dim(\mathcal{L}(V, W)) = mn .$$

Proof. Note that $\dim(M_{m \times n}(F)) = mn$. By Theorem 3.0.17 the vector space $\mathcal{L}(V, W)$ is isomorphic to $M_{m \times n}(F)$ and by Theorem 3.0.14

$$\dim(\mathcal{L}(V, W)) = \dim(M_{m \times n}(F)) = mn . \quad \square$$

As we have seen in Corollary 3.0.16 any n -dimensional vector space V is isomorphic to F^n . However, this corollary does not give an explicit isomorphism, but just states the existence of one. In the next theorem we will construct an explicit isomorphism $\varphi_\beta: V \rightarrow F^n$ that depends on the choice of a basis β of V . It is given by mapping a vector $x \in V$ to the coordinate vector $[x]_\beta$ of x relative to β (see Def. 2.1.4).

Theorem 3.0.19. *For any finite-dimensional vector space V over the field F with ordered basis β and $n = \dim(V)$, the map $\varphi_\beta: V \rightarrow F^n$ given by $\varphi_\beta(x) = [x]_\beta$ is an isomorphism.*

Proof. Let $\beta = \{v_1, \dots, v_n\}$. We have to prove that φ_β is a bijective linear transformation. We will first show that it is linear: Let $x, y \in V$ and $c \in F$. Let $[x]_\beta = (a_1, \dots, a_n) \in F^n$ and $[y]_\beta = (b_1, \dots, b_n) \in F^n$. By definition $x = \sum_{j=1}^n a_j v_j$, $y = \sum_{j=1}^n b_j v_j$ and

$$\sum_{j=1}^n (ca_j + b_j)v_j = c \left(\sum_{j=1}^n a_j v_j \right) + \sum_{j=1}^n b_j v_j = cx + y .$$

Therefore we have

$$\begin{aligned} \varphi_\beta(cx + y) &= [cx + y]_\beta = (ca_1 + b_1, \dots, ca_n + b_n) \\ &= c(a_1, \dots, a_n) + (b_1, \dots, b_n) \\ &= c[x]_\beta + [y]_\beta = c\varphi_\beta(x) + \varphi_\beta(y) \end{aligned}$$

and φ_β is indeed a linear transformation. Suppose that $x \in V$ is such that $(a_1, \dots, a_n) = \varphi_\beta(x) = [x]_\beta = (0, \dots, 0)$. This means $x = \sum_{j=1}^n a_j v_j = \sum_{j=1}^n 0 v_j = 0$. Thus, $\ker(\varphi_\beta) = \{0\}$. Therefore φ_β is injective. Now let $(c_1, \dots, c_n) \in F^n$ be an arbitrary vector and let $x = \sum_{j=1}^n c_j v_j \in V$. By construction $\varphi_\beta(x) = [x]_\beta = (c_1, \dots, c_n)$. Hence, φ_β is surjective as well. \square

We will use isomorphisms induced by ordered bases to prove a structural result about linear transformations in the next theorem.

Theorem 3.0.20. Let V and W be finite-dimensional vector spaces over the field F of dimension n and m , respectively. Let β and γ be ordered bases for V , respectively W . Let $T: V \rightarrow W$ be a linear transformation and let $A = [T]_{\beta}^{\gamma}$ be the matrix representation of T with respect to the ordered bases β and γ . Then

$$\varphi_{\gamma} \circ T = L_A \circ \varphi_{\beta} \tag{6}$$

Proof. Let $x \in V$. If the left hand side of (6) is applied to x we obtain

$$(\varphi_{\gamma} \circ T)(x) = \varphi_{\gamma}(T(x)) = [T(x)]_{\gamma} .$$

If the right hand side of (6) is evaluated on x we get

$$(L_A \circ \varphi_{\beta})(x) = L_A(\varphi_{\beta}(x)) = A[x]_{\beta} = [T]_{\beta}^{\gamma}[x]_{\beta} .$$

However, $[T(x)]_{\gamma} = [T]_{\beta}^{\gamma}[x]_{\beta}$ by Thm. 2.1.5. Therefore $(\varphi_{\gamma} \circ T)(x) = (L_A \circ \varphi_{\beta})(x)$ for all $x \in V$. This proves the statement. \square

Remark 3.0.21. The linear transformations from Thm. 3.0.20 appear as arrows in the diagram shown below.

$$\begin{array}{ccc} V & \xrightarrow{T} & W \\ \varphi_{\beta} \downarrow & & \downarrow \varphi_{\gamma} \\ F^n & \xrightarrow{L_A} & F^m \end{array}$$

There are two ways of composing the arrows in this diagram from V to F^m : The first option first runs through T and then through φ_{γ} giving the composition $\varphi_{\gamma} \circ T$ (Mind the order of composition here!), whereas the second option runs through φ_{β} and then through L_A giving $L_A \circ \varphi_{\beta}$. The statement $\varphi_{\gamma} \circ T = L_A \circ \varphi_{\beta}$ of Thm. 3.0.20 tells us that both compositions agree and in terms of the diagram we could phrase this by saying that it does not matter which option of following the arrows we pick, we always end up with the same map. Diagrams with this property appear quite often in mathematics. Since the condition $\varphi_{\gamma} \circ T = L_A \circ \varphi_{\beta}$ is a bit similar to commutativity, in the sense that it does not matter in which way we compose, such diagrams are called *commutative diagrams*.

Remark 3.0.22. Since $\varphi_{\gamma}: W \rightarrow F^m$ in Theorem 3.0.20 is an isomorphism, we can rewrite (6) as

$$T = \varphi_{\gamma}^{-1} \circ L_A \circ \varphi_{\beta} . \tag{7}$$

Or in words: Any linear transformation $T: V \rightarrow W$ between finite-dimensional vector spaces can be obtained as the composition of the inverse of a coordinate isomorphism with a left multiplication transformation and another coordinate isomorphism. So, up to choosing a basis any linear transformation can be described by left multiplication with a matrix.

Exercise 3.0.23. Reprove the statement of Theorem 2.1.2 using (7) and the following fact about left multiplication transformations: $L_A \circ L_B = L_{AB}$.

Remark 3.0.24. Consider the isomorphisms $T: V \rightarrow V$, i.e. from the vector space V to itself. These elements form a subset of $\mathcal{L}(V)$. We have already seen in Remark 2.0.12 that the vector space $\mathcal{L}(V)$ is quite special in the sense that we can compose any two elements $T, S \in \mathcal{L}(V)$. This composition is associative. Let $GL(V) \subset \mathcal{L}(V)$ be the subset of all isomorphisms. Since the composition of two isomorphisms is again an isomorphism, $GL(V)$ is closed under composition. Moreover, I_V is a neutral element for this composition and by definition any $T \in GL(V)$ has an inverse $T^{-1} \in GL(V)$, which satisfies $TT^{-1} = T^{-1}T = I_V$. Altogether we see that $GL(V)$ is a group! This group is called the *general linear group* of the vector space V .

Example 3.0.25. Let $T: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the linear transformation given by $T(x, y) = (2x - y, 2y - x)$. With respect to the standard basis $\alpha = \{e_1, e_2\}$ given by $e_1 = (1, 0)$ and $e_2 = (0, 1)$ it has

$$[T]_{\alpha}^{\alpha} = \begin{pmatrix} 2 & -1 \\ -1 & 2 \end{pmatrix}$$

as its matrix representation. We would like to compute T^{2017} by hand. This seems impossible at first, since it requires us to perform over 2000 matrix multiplications. However, we can make our life a lot easier by a clever choice of basis. Let $\beta = \{f_1, f_2\}$ be the ordered basis with $f_1 = (1, 1)$ and $f_2 = (1, -1)$. Observe that

$$\begin{aligned} T(1, 1) &= (1, 1) = f_1 \\ T(1, -1) &= (3, -3) = 3f_2 \end{aligned}$$

Thus, if we compute the matrix representation of T with respect to the ordered basis β we obtain

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

By Thm. 3.0.20, Remark 3.0.22 and the fact that $(L_A)^n = L_{A^n}$ we can now compute powers of T :

$$\begin{aligned} T &= \varphi_{\beta}^{-1} \circ L_A \circ \varphi_{\beta} \\ T^2 &= \varphi_{\beta}^{-1} \circ (L_A)^2 \circ \varphi_{\beta} = \varphi_{\beta}^{-1} \circ L_{A^2} \circ \varphi_{\beta} \\ T^3 &= \varphi_{\beta}^{-1} \circ L_{A^3} \circ \varphi_{\beta} \\ &\vdots \\ T^n &= \varphi_{\beta}^{-1} \circ L_{A^n} \circ \varphi_{\beta} \end{aligned}$$

However, since A is a diagonal matrix, it is easy to compute its n th power:

$$A^n = \begin{pmatrix} 1 & 0 \\ 0 & 3^n \end{pmatrix}.$$

Therefore we just need to compute the isomorphism $\varphi_\beta: \mathbb{R}^2 \rightarrow \mathbb{R}^2$. With $e_1 = \frac{1}{2}(f_1 + f_2)$ and $e_2 = \frac{1}{2}(f_1 - f_2)$ we obtain for $(x, y) \in \mathbb{R}^2$

$$\begin{aligned}(x, y) &= xe_1 + ye_2 = \frac{1}{2}x(f_1 + f_2) + \frac{1}{2}y(f_1 - f_2) \\ &= \frac{1}{2}(x + y)f_1 + \frac{1}{2}(x - y)f_2\end{aligned}$$

and hence $\varphi_\beta((x, y)) = (\frac{1}{2}(x + y), \frac{1}{2}(x - y))$. Likewise, we compute $\varphi_\beta^{-1}((x, y)) = (x + y, x - y)$. (It is a coincidence that this turned out to be a similar formula in both cases!) Now we just need to compose and obtain

$$\begin{aligned}T^{2017}((x, y)) &= \varphi_\beta^{-1}(L_{A^{2017}}(\varphi_\beta((x, y)))) = \varphi_\beta^{-1}\left(L_{A^{2017}}\left(\left(\frac{1}{2}(x + y), \frac{1}{2}(x - y)\right)\right)\right) \\ &= \varphi_\beta^{-1}\left(\left(\frac{1}{2}(x + y), \frac{3^{2017}}{2}(x - y)\right)\right) \\ &= \left(\frac{1}{2}(x + y) + \frac{3^{2017}}{2}(x - y), \frac{1}{2}(x + y) - \frac{3^{2017}}{2}(x - y)\right)\end{aligned}$$

This is a complete description of the linear transformation T^{2017} in terms of coordinates with respect to the standard basis. We see that the basis $\beta = \{f_1, f_2\}$ is very helpful in this computation, since the matrix representation of T with respect to β is diagonal and we can easily compute powers of diagonal matrices. This motivates the question: For which matrices can we find a basis with the property that its matrix representation with respect to that basis becomes diagonal?

4 Systems of Linear Equations

In this section we are going to revisit the procedure for solving systems of linear equations. This is a problem that is ubiquitous in mathematics as well as in physics. As we will see in the two examples below, many problems can be rephrased in terms of linear equations. We will then establish the connection between systems of linear equations and the theory we have learned so far.

Example 4.0.1. This example highlights a common problem in engineering. Consider the circuit shown in Figure 2. It consists of a battery as a power source and three parallel resistors. We would like to compute the currents I_1, I_2 and I_3 flowing through each of the three branches. To achieve this, we have guessed a direction of the current, which can be arbitrary and is indicated by the arrow over I_k . The system of currents obeys Kirchhoff's laws

- i) At each junction in the circuit the sum over the ingoing currents equals the sum of outgoing currents.
- ii) In each closed loop in the circuit, the sum over the electrical potential differences adds up to zero. The potential difference at a resistor R_k with current I_l flowing

through is given by $\pm R_k I_l$, where the sign depends on whether the loop runs along the arrow or against it. Likewise, the potential difference at a battery is plus or minus its voltage value.




Figure 2: A circuit with a battery and resistors.

In the circuit we have two junctions at which three wires meet. Since all currents flow **into** the junction on the left and **out of** the junction on the right, the first of the two laws stated above leads to the following two linear equations

$$\begin{aligned} I_1 + I_2 + I_3 &= 0 , \\ -I_1 - I_2 - I_3 &= 0 . \end{aligned}$$

Observe that these two equations are actually equivalent in the sense that a triple (I_1, I_2, I_3) is a solution of the first if and only if it is a solution of the second. The loop around the upper half of the circuit yields the equation

$$24 + 2I_1 - 3I_2 = 0 \quad \Leftrightarrow \quad -2I_1 + 3I_2 = 24$$

and the loop in the lower half gives

$$3I_2 - 6I_3 = 0 .$$

Altogether we arrive at the following system of linear equations

$$\begin{aligned} I_1 + I_2 + I_3 &= 0 \\ -2I_1 + 3I_2 &= 24 \\ 3I_2 - 6I_3 &= 0 \end{aligned}$$

Note that we can rewrite the above system of linear equations in terms of a matrix multiplication as follows

$$\begin{pmatrix} 1 & 1 & 1 \\ -2 & 3 & 0 \\ 0 & 3 & -6 \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \\ I_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 24 \\ 0 \end{pmatrix}$$

Exercise 4.0.2. Solve the above system of linear equations by reducing the matrix to its echelon form. How many solutions does it have? The intermediate steps of one possible solution are shown below.

$$\left(\begin{array}{ccc|c} 1 & 1 & 1 & 0 \\ -2 & 3 & 0 & 24 \\ 0 & 3 & -6 & 0 \end{array} \right) \rightsquigarrow \left(\begin{array}{ccc|c} 1 & 1 & 1 & 0 \\ 0 & 5 & 2 & 24 \\ 0 & 0 & 36 & 72 \end{array} \right) \rightsquigarrow \left(\begin{array}{ccc|c} 1 & 0 & 0 & -6 \\ 0 & 1 & 0 & 4 \\ 0 & 0 & 1 & 2 \end{array} \right)$$

As can be seen from Exercise 4.0.2, the value obtained for I_1 is negative. This means that the direction we guessed for the current I_1 in the beginning was actually wrong and it flows in the direction opposite to the arrow shown in Figure 2.

Exercise 4.0.3. Find the system of linear equations for the circuit, in which the direction of I_1 is reversed and solve it. If you have done this correctly, the absolute values of I_1, I_2, I_3 should be the same as in Exercise 4.0.2, but the solution for I_1 should be positive.

Example 4.0.4. Another example is based on approximation problems: Suppose we want to find a polynomial p of degree d such that the value of p at the k points $x_1, \dots, x_k \in \mathbb{R}$ agrees with k given values $y_1, \dots, y_k \in \mathbb{R}$. A general polynomial of degree d is of the form

$$p(t) = \sum_{j=0}^d a_j t^j = a_d t^d + a_{d-1} t^{d-1} + \dots + a_1 t + a_0 .$$

To find a polynomial that satisfies $p(x_i) = y_i$, we have to find the right values of the $(d+1)$ coefficients $a_0, \dots, a_d \in \mathbb{R}$. The condition $p(x_i) = y_i$ leads to the equations

$$\begin{aligned} a_d x_1^d + a_{d-1} x_1^{d-1} + \dots + a_1 x_1 + a_0 &= y_1 \\ a_d x_2^d + a_{d-1} x_2^{d-1} + \dots + a_1 x_2 + a_0 &= y_2 \\ &\vdots \\ a_d x_k^d + a_{d-1} x_k^{d-1} + \dots + a_1 x_k + a_0 &= y_k \end{aligned}$$

Note that we fixed the values of $x_i \in \mathbb{R}$ and $y_i \in \mathbb{R}$ beforehand and we have to look at the above system of equations as equations in the coefficients a_0, \dots, a_d . As such, we can rewrite it as an equation for the vector $(a_0, \dots, a_d) \in \mathbb{R}^{d+1}$.

$$\begin{pmatrix} 1 & x_1 & x_1^2 & \dots & x_1^{d-1} & x_1^d \\ 1 & x_2 & x_2^2 & \dots & x_2^{d-1} & x_2^d \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & x_{k-1} & x_{k-1}^2 & \dots & x_{k-1}^{d-1} & x_{k-1}^d \\ 1 & x_k & x_k^2 & \dots & x_k^{d-1} & x_k^d \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ a_2 \\ \vdots \\ a_{d-1} \\ a_d \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_{k-1} \\ y_k \end{pmatrix} \quad (8)$$

Exercise 4.0.5. Check that the vector equation (8) indeed reproduces the system of linear equations from above.

To answer the question if and how many polynomials p with $p(x_i) = y_i$ exist for given values $x_i \in \mathbb{R}$ and $y_i \in \mathbb{R}$ we now have to solve the above system of linear equations. Let us make this a little more precise and find all the polynomials p of degree 2 that satisfy $p(0) = 0$ and $p(1) = 2$. In this case we have $d = 2$, $k = 2$, $x_1 = 0, x_2 = 1$ and $y_1 = 0, y_2 = 2$. The system of linear equations is then given by

$$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 2 \end{pmatrix} . \quad (9)$$

We obtain $a_0 = 0$ and $a_1 + a_2 = 2$. Hence, there are infinitely many solutions that can be described as follows: Let $a = a_1$. By the above we have $a_2 = 2 - a$. Then for each $a \in \mathbb{R}$ the polynomial $p(t) = at + (2 - a)t^2$ is a solution. It is easy to check that $p(0) = 0$ and $p(1) = 2$. The above computation shows us that these are all possible solutions.

To talk about systems of linear equations like the ones in the two examples above, we first need some vocabulary: In general a system of equations of the form

$$\begin{aligned} a_{11}x_1 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + \cdots + a_{mn}x_n &= b_m \end{aligned} \quad (10)$$

where $a_{ij} \in F$ (for $1 \leq i \leq m$ and $1 \leq j \leq n$), $b_k \in F$ (for $1 \leq k \leq m$) are scalars in a field F and x_1, \dots, x_n are n variables taking values in F , is called a **system of m linear equations in n unknowns over the field F** .

The $m \times n$ -matrix A given by

$$A = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mn} \end{pmatrix}$$

is called the **coefficient matrix** of the system. Let

$$x = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} \in F^n \quad \text{and} \quad b = \begin{pmatrix} b_1 \\ \vdots \\ b_m \end{pmatrix} \in F^m .$$

Then we can rewrite the system of linear equations as a single matrix equation as follows:

$$Ax = b .$$

An n -tuple $(s_1, \dots, s_n) \in F^n$ is called a **solution** to the system if $As = b$. The set

$$S = \{s \in F^n \mid As = b\}$$

is called the **solution set** of the system. The system of linear equations is called **consistent** if the solution set is non-empty, otherwise it is called **inconsistent**.

A system of m linear equations in N unknowns $Ax = b$ is said to be **homogeneous** if $b = 0$. Otherwise it is said to be **nonhomogeneous**.

4.1 Solving Systems of Linear Equations

In this section we will briefly recapitulate how to solve systems of linear equations like the one in (10). The following three operations will be crucial for us:

Definition 4.1.1. Let A be an $m \times n$ matrix. The following three operations are called **elementary row operations**:

- (1) interchanging any two rows of A ;
- (2) multiplying any row of A by a nonzero scalar;
- (3) adding any scalar multiple of a row of A to another row.

Depending on whether the operation is obtained from (1), (2) or (3) we will say that we performed an operation of **type 1**, **type 2** or **type 3**.

Given a system of linear equations $Ax = b$ it will come in handy to use the following compact notation to combine A and b into one matrix.

Definition 4.1.2. Let A be an $m \times n$ matrix and let B be an $m \times p$ matrix. The **augmented matrix** $(A|B)$ is defined to be the $m \times (n + p)$ matrix whose first n columns are the columns of A and whose last p columns are the ones of B .

We will illustrate the recipe for solving linear equations of the form $Ax = b$ using the following example (we will look for solutions $(x_1, x_2, x_3, x_4) \in \mathbb{R}^4$):

$$\begin{aligned} 3x_1 + 3x_2 + 2x_3 - 2x_4 &= 1 \\ x_1 + x_2 + x_3 &= 3 \\ x_1 + x_2 + 2x_3 - x_4 &= 2 \end{aligned} \tag{11}$$

(Step 1) Identify the coefficient matrix A of the system and write down the augmented matrix $(A|b)$. For the system in (11) this looks like this:

$$\left(\begin{array}{cccc|c} 3 & 3 & 2 & -2 & 1 \\ 1 & 1 & 1 & 0 & 3 \\ 1 & 1 & 2 & -1 & 2 \end{array} \right)$$

(Step 2) Create a 1 in the leftmost nonzero column using row operations. This can be accomplished in (11) by interchanging the first and third rows, which yields:

$$\left(\begin{array}{cccc|c} 1 & 1 & 2 & -1 & 2 \\ 1 & 1 & 1 & 0 & 3 \\ 3 & 3 & 2 & -2 & 1 \end{array} \right)$$

(Step 3) By means of type 3 row operations, use the first row to obtain zeros in the remaining positions of the leftmost nonzero column. In the example, we add (-1) times the first row to the second row and add (-3) times the first row to the third row to obtain

$$\left(\begin{array}{cccc|c} 1 & 1 & 2 & -1 & 2 \\ 0 & 0 & -1 & 1 & 1 \\ 0 & 0 & -4 & 1 & -5 \end{array} \right)$$

(Step 4) As in (Step 2) create a 1 in the next row in the leftmost possible column without using the previous rows. In the example the leftmost possible position is the third column and we can create a 1 there by multiplying the second row by -1 . This produces the following matrix:

$$\left(\begin{array}{cccc|c} 1 & 1 & 2 & -1 & 2 \\ 0 & 0 & 1 & -1 & -1 \\ 0 & 0 & -4 & 1 & -5 \end{array} \right)$$

(Step 5) Use type 3 row operations to obtain zeroes below the 1 created in the preceding step. In the example we have to add 4 times the the second row to the third to obtain

$$\left(\begin{array}{cccc|c} 1 & 1 & 2 & -1 & 2 \\ 0 & 0 & 1 & -1 & -1 \\ 0 & 0 & 0 & -3 & -9 \end{array} \right)$$

(Step 6) Repeat the previous two steps on each succeeding row until no nonzero rows remain. In the example this is achieved by dividing the last row by (-3) (in other words multiplying by $-\frac{1}{3}$). The result of this step is an **upper triangular matrix**, i.e. all of its nonzero entries are on or above the diagonal. In the example this looks as follows:

$$\left(\begin{array}{cccc|c} 1 & 1 & 2 & -1 & 2 \\ 0 & 0 & 1 & -1 & -1 \\ 0 & 0 & 0 & 1 & 3 \end{array} \right)$$

(Step 7) Beginning with the last row add multiples of each row to the rows above it to create zeroes above the leading 1 in each row. In the example above this can

be done in two steps: Add the third row to the second and to the first. Then add the (-2) times the second row to the first. The result is

$$\left(\begin{array}{cccc|c} 1 & 1 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 2 \\ 0 & 0 & 0 & 1 & 3 \end{array} \right)$$

This reduction of the augmented matrix corresponds to a system of linear equations that is much easier to solve. It is given by

$$\begin{aligned} x_1 + x_2 &= 1 \\ x_3 &= 2 \\ x_4 &= 3 \end{aligned}$$

If we set $t = x_1$, then $x_2 = 1 - x_1 = 1 - t$. Therefore the solution set S is

$$S = \left\{ \left(\begin{array}{c} t \\ 1-t \\ 2 \\ 3 \end{array} \right) \mid t \in \mathbb{R} \right\} = \left\{ \left(\begin{array}{c} 0 \\ 1 \\ 2 \\ 3 \end{array} \right) + t \left(\begin{array}{c} 1 \\ -1 \\ 0 \\ 0 \end{array} \right) \mid t \in \mathbb{R} \right\}$$

Note that the vector $(1, -1, 0, 0) \in \mathbb{R}^4$ is a solution of the corresponding homogeneous system of linear equations, which we obtain by setting the right hand side equal to the zero vector. The augmented matrix of the homogeneous system is given by:

$$\left(\begin{array}{cccc|c} 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{array} \right)$$

and the solution set S_h of this system is

$$S_h = \left\{ t \left(\begin{array}{c} 1 \\ -1 \\ 0 \\ 0 \end{array} \right) \mid t \in \mathbb{R} \right\} \subset \mathbb{R}^4$$

Exercise 4.1.3. Show that S_h is a linear subspace of \mathbb{R}^4 . What is the dimension of S_h ? What about the solution set $S \subset \mathbb{R}^4$ of the nonhomogeneous system? Is it also a linear subspace? Why not?

In fact, we will see later that the solution set of a *homogeneous* system of linear equations is always a linear subspace.

The augmented matrix we obtained for the example above in (Step 7) allows us to directly read off the solution set. It seems sensible to give this convenient form of the augmented matrix a name.

Definition 4.1.4. A matrix is said to be in **reduced row echelon form** if the following three conditions are satisfied.

- (a) Any row containing a nonzero entry precedes any row in which all the entries are zero (if any).
- (b) The first nonzero entry in each row is the only nonzero entry in its column.
- (c) The first nonzero entry in each row is 1 and it occurs in a column to the right of the first nonzero entry in the preceding row.

4.2 Elementary Matrices

In this section we will discuss why the three elementary row operations listed in Def. 4.1.1 do not change the solution set of a system of linear equations. We will see that the row operations can be obtained by multiplying the augmented matrix with certain invertible matrices.

Definition 4.2.1. An $n \times n$ -matrix E is called **elementary** if it can be obtained by performing a single elementary row operation on the identity matrix I_n . An elementary matrix E is said to be **of type 1, 2 or 3** according to whether the elementary operation performed on I_n is of type 1, 2 or 3, respectively.

Example 4.2.2. The following 4×4 -matrices E_1 and E_2 are examples of elementary matrices of type 1. The matrix E_1 is obtained from I_4 by interchanging the first two rows and E_2 is obtained from I_4 by interchanging the first row with the the fourth one.

$$E_1 = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad E_2 = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

The 3×3 -matrix E_3 shown below is obtained from I_3 by multiplying the second row by (-4) , hence it is an elementary matrix of type 2.

$$E_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -4 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

The following 5×5 -matrix E_4 is obtained from I_5 by adding (-2) times the last row to the first. Therefore it is an elementary matrix of type 3.

$$E_4 = \begin{pmatrix} 1 & 0 & 0 & 0 & -2 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

We will now show that performing an elementary row operation really is the same as multiplying the matrix by an elementary matrix.

Theorem 4.2.3. *Let F be a field and let $A \in M_{m \times n}(F)$. Suppose that $B \in M_{m \times n}(F)$ is obtained from A by performing an elementary row operation. Let $E \in M_{m \times m}(F)$ be the elementary matrix obtained from I_m by performing the same elementary row operation. Then we have*

$$EA = B .$$

Proof. We will compare the matrices EA and B column by column. For this we need the following observation: Let $e_i \in F^n$ for $i \in \{1, \dots, n\}$ be the vector that has a 1 in the i th entry and 0 in every other entry. The vectors e_1, \dots, e_n form a basis for F^n . Observe that for $k \in \{1, \dots, n\}$ the vector Ae_k coincides with the k th column of the matrix A .

Suppose that E (respectively B) is obtained by interchanging the i th and j th row of I_m (respectively A). Without loss of generality we will assume $i < j$. Let $v = (v_1, \dots, v_m) \in F^m$, let $k \in \{1, \dots, m\}$ and note that

$$(0 \quad \dots \quad 0 \quad \underbrace{1}_k \quad 0 \quad \dots \quad 0) \begin{pmatrix} v_1 \\ \vdots \\ v_m \end{pmatrix} = v_k \quad (12)$$

In particular, Ev is the vector with the i th and j th entry interchanged, i.e.

$$Ev = \begin{pmatrix} v_1 \\ \vdots \\ v_j \\ \vdots \\ v_i \\ \vdots \\ v_m \end{pmatrix} .$$

Therefore the vector $E Ae_k = E(Ae_k) \in F^m$ is obtained from the k th column vector of A by interchanging the i th and j th entry, but this is also the k th column vector of B . This shows the statement for elementary row operations of type 1.

Next suppose that E (respectively B) is obtained by multiplying the i th row of I_m (respectively A) by a scalar $\lambda \in F$ with $\lambda \neq 0$. If we multiply the i th row of E with a vector $v = (v_1, \dots, v_m)$, we get

$$(0 \quad \dots \quad 0 \quad \underbrace{\lambda}_i \quad 0 \quad \dots \quad 0) \begin{pmatrix} v_1 \\ \vdots \\ v_m \end{pmatrix} = \lambda v_i$$

Using the above observation together with (12), we see that $E Ae_k = E(Ae_k)$ is obtained from the k th column vector of A by multiplying the i th entry by λ . But this agrees with the k th column vector of B . This proves the statement for row operations of type 2.

Now suppose that E (respectively B) is obtained by adding λ times the j th row of I_m (respectively A) to the i th row of I_m (respectively A). We will assume that $i < j$. The proof for the case $j < i$ is very similar. The i th row of E looks like this

$$(0 \quad \dots \quad 0 \quad \underbrace{1}_i \quad 0 \quad \dots \quad 0 \quad \underbrace{\lambda}_j \quad 0) .$$

If we multiply this row with a vector $\overset{i}{v} = (v_1, \dots, v_m) \in F^m$ we get

$$(0 \quad \dots \quad 0 \quad \underbrace{1}_i \quad 0 \quad \dots \quad 0 \quad \underbrace{\lambda}_j \quad 0) \begin{pmatrix} v_1 \\ \vdots \\ v_m \end{pmatrix} = v_i + \lambda v_j .$$

In particular, $E A e_k = E(A e_k)$ is obtained from the k th column vector of A by adding λ times the j th entry to the i th one. This agrees with the k th column vector of B . In all cases we therefore have $EA = B$. \square

Theorem 4.2.4. *Elementary matrices are invertible, and the inverse of an elementary matrix of type k is again an elementary matrix of type k .*

Proof. Each of the three types of elementary row operations can be reversed by an elementary row operation of the same type: Performing a type 1 operation on a matrix A twice will return A to its original state. The inverse of multiplying a row by $\lambda \in F$ with $\lambda \neq 0$, i.e. a type 2 operation, is given by multiplying the row by $\frac{1}{\lambda}$, which is again a type 2 operation. Lastly, adding λ times the i th row to the j th one is reversed by adding $-\lambda$ times the i th row to the j th and both of these operations are of type 3.

Now let $E \in M_{n \times n}(F)$ be an elementary $n \times n$ -matrix. Let $E' \in M_{n \times n}(F)$ be obtained by performing the elementary row operation on I_n that is inverse to the one used to obtain E . In particular, E' is an elementary matrix of the same type as E . Now it follows from Theorem 4.2.3 that

$$E E' = E' E = I_n . \quad \square$$

We used elementary row operations to change a system of linear equations to one which has a particularly simple form.

Definition 4.2.5. Two systems of linear equations are called **equivalent** if they have the same solution set.

Theorem 4.2.6. *Let F be a field, let $A \in M_{m \times n}(F)$ and let $b \in F^m$. Consider the system $Ax = b$ of m linear equations in n unknowns with augmented matrix $(A|b)$. Let $C \in M_{m \times m}(F)$ be an invertible matrix. Then the systems corresponding to $(A|b)$ and $(CA|Cb)$ are equivalent.*

Proof. Let K be the solution set for $Ax = b$. Let K' be the one of $CAx = Cb$. If $w \in K$, then $Aw = b$. Therefore, $(CA)w = Cb$ and hence $w \in K'$. Thus, $K \subseteq K'$. Conversely, if $w \in K'$, then $CAw = Cb$. Hence,

$$Aw = C^{-1}(CAw) = C^{-1}(Cb) = b .$$

So $w \in K$. Therefore $K' \subseteq K$ and $K' = K$. \square

Corollary 4.2.7. *Let $Ax = b$ be a system of m linear equations in n unknowns with augmented matrix $(A|b)$ as in Theorem 4.2.6 . If $(A'|b')$ is obtained from $(A|b)$ by a finite number of elementary row operations, then $A'x = b'$ is equivalent to $Ax = b$.*

Proof. If $(A'|b')$ is obtained from $(A|b)$ by a finite number of row operations, then by Thm. 4.2.3 the matrix $(A'|b')$ is obtained from $(A|b)$ by multiplication by a matrix C that is a product of elementary matrices E_1, \dots, E_p , i.e. $C = E_p E_{p-1} \cdots E_1$. Then we have

$$(A'|b') = C(A|b) = (CA|Cb)$$

Since each E_i is invertible, so is C . Now, $CA = A'$ and $Cb = b'$. Thus, by Thm. 4.2.6 $A'x = b'$ and $Ax = b$ are equivalent. \square

4.3 The Rank of a Matrix

Given a linear transformation $T: V \rightarrow W$ between two vector spaces over a field F , the rank of T is defined as the dimension of the range $\text{Im}(T)$. We can now use the left multiplication transformation L_A associated to a matrix $A \in M_{m \times n}(F)$ to extend the definition of rank to matrices.

Definition 4.3.1. Let $m, n \in \mathbb{N}$ and let $A \in M_{m \times n}(F)$ be an $m \times n$ -matrix with entries in the field F . We define the **rank** of A , denoted by $\text{rank}(A)$, to be the rank of the left multiplication transformation $L_A: F^n \rightarrow F^m$.

Now we have two options of associating a rank to a linear transformation $T: V \rightarrow W$. We can use the usual definition $\text{rank}(T) = \dim(\text{Im}(T))$ or we could form the matrix $A = [T]_{\beta}^{\gamma}$ of T with respect to ordered bases β and γ and look at $\text{rank}(A) = \text{rank}(L_A)$. Luckily, these two numbers turn out to be equal. But before we can prove this, we need the following Lemma, which is interesting on its own and states that compositions with isomorphisms do not change the rank.

Lemma 4.3.2. *Let V, W, Z be vector spaces over the field F .*

a) *If $T: V \rightarrow W$ is a linear transformation and $S: W \rightarrow Z$ an isomorphism, then*

$$\text{rank}(ST) = \text{rank}(T) .$$

b) *If $T: W \rightarrow Z$ is a linear transformation and $S: V \rightarrow W$ is an isomorphism , then*

$$\text{rank}(TS) = \text{rank}(T) .$$

Proof. We will first prove a). Let $n = \text{rank}(T)$ and let $\beta = \{w_1, \dots, w_n\}$ be a basis for $\text{Im}(T)$. We have to show that $S(\beta) = \{S(w_1), \dots, S(w_n)\}$ is a basis of $\text{Im}(ST)$. Note that $\beta = S^{-1}(S(\beta))$ is a linearly independent subset of $S^{-1}(Z)$ and $S(w_j) \in S(\beta)$ has the property that $S^{-1}(S(w_j)) = w_j$. By Exercise 5 on Sheet 3, $S(\beta)$ is linearly independent. Now let $z \in \text{Im}(ST)$. By definition there is a vector $v \in V$ with $ST(v) = z$. Since

$T(v) \in \text{Im}(T)$ and β is a basis for $\text{Im}(T)$ we can find coefficients $\alpha_1, \dots, \alpha_n \in F$ with $T(v) = \sum_{i=1}^n \alpha_i w_i$. But then

$$z = ST(v) = S(T(v)) = S\left(\sum_{i=1}^n \alpha_i w_i\right) = \sum_{i=1}^n \alpha_i S(w_i).$$

Therefore $\text{Im}(ST) = \text{span}(S(\beta))$, the set $S(\beta)$ is a basis for $\text{Im}(ST)$ and $\text{rank}(ST) = n = \text{rank}(T)$.

The proof of b) is much easier. Since S is an isomorphism, it is in particular surjective, which implies

$$\text{Im}(TS) = \{TS(v) \in Z \mid v \in V\} = \{T(w) \in Z \mid w \in W\} = \text{Im}(T).$$

Thus, $\text{rank}(TS) = \dim(\text{Im}(TS)) = \dim(\text{Im}(T)) = \text{rank}(T)$. \square

Now we are ready to prove the claim we made above:

Theorem 4.3.3. *Let V and W be finite-dimensional vector spaces over the field F with ordered bases β and γ , respectively. Let $T: V \rightarrow W$ be a linear transformation. Then*

$$\text{rank}(T) = \text{rank}\left([T]_{\beta}^{\gamma}\right).$$

Proof. Let $A = [T]_{\beta}^{\gamma}$ and let $\varphi_{\beta}, \varphi_{\gamma}$ be the coordinate isomorphisms with respect to β and γ , respectively (see Thm. 3.0.19). By Thm. 3.0.20 and Remark 3.0.22, the linear transformation T satisfies $T = \varphi_{\gamma}^{-1} \circ L_A \circ \varphi_{\beta}$. Thus, Lemma 4.3.2 implies that

$$\text{rank}(T) = \text{rank}(\varphi_{\gamma}^{-1} \circ L_A \circ \varphi_{\beta}) = \text{rank}(L_A \circ \varphi_{\beta}) = \text{rank}(L_A) = \text{rank}(A). \quad \square$$

The next corollary is just a reformulation of Lemma 4.3.2 in terms of matrices. Note that if A is an invertible $n \times n$ -matrix with inverse B , then L_A is an isomorphism with inverse L_B .

Corollary 4.3.4. *Let A be an $m \times n$ -matrix with entries in a field F . Let P and Q be invertible $m \times m$ and $n \times n$ -matrices with entries in F , respectively. Then*

$$\text{rank}(PAQ) = \text{rank}(A)$$

Proof. Note that L_P and L_Q are isomorphisms and that $L_{PAQ} = L_P L_A L_Q$. By Lemma 4.3.2

$$\text{rank}(PAQ) = \text{rank}(L_{PAQ}) = \text{rank}(L_P L_A L_Q) = \text{rank}(L_A) = \text{rank}(A). \quad \square$$

The next theorem provides a very helpful interpretation of the rank of a matrix in terms of its column vectors. We will later use this description to see how we can easily compute the rank of a given matrix.

Theorem 4.3.5. *The rank of any matrix $A \in M_{m \times n}(F)$ with entries in a field F equals the maximum number of its linearly independent columns; that is, the rank of a matrix is the dimension of the subspace of F^m generated by its columns.*

Proof. By definition we have $\text{rank}(A) = \text{rank}(L_A) = \dim(\text{Im}(L_A))$. Let $\beta = \{e_1, \dots, e_n\}$ be the standard basis for F^n . By Thm. 1.0.12 the rank of $\text{Im}(L_A)$ can be expressed as follows

$$\text{Im}(L_A) = \text{span}(L_A(\beta)) = \text{span}(\{L_A(e_1), \dots, L_A(e_n)\}) .$$

Now note that $L_A(e_j) = Ae_j =: c_j$ is the j th column c_j of the matrix A for all $j \in \{1, \dots, n\}$. Therefore

$$\text{Im}(L_A) = \text{span}(\{c_1, \dots, c_n\})$$

and $\text{rank}(A) = \dim(\text{Im}(L_A)) = \dim(\text{span}(\{c_1, \dots, c_n\}))$. □

Example 4.3.6. Let us try to find the rank of the following 3×3 -matrix A with entries in the field \mathbb{R} :

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

Let $c_1, c_2, c_3 \in \mathbb{R}^3$ be the column vectors, i.e.

$$c_1 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix} , \quad c_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} , \quad c_3 = \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

We immediately see that $c_1 + c_2 = c_3$ and that c_1 and c_2 are linearly independent, since they are not scalar multiples of one another. Therefore the maximum number of linearly independent columns of A is 2 and Theorem 4.3.5 implies $\text{rank}(A) = 2$.

Combining Theorem 4.2.3, Theorem 4.2.4 and Corollary 4.3.4 we see that the rank of a matrix does not change under elementary row operations. In fact, it is also true that the rank does not change under elementary column operations. Before we prove this statement we first need to discuss the definition of these operations. It is very similar to Def. 4.1.1:

Definition 4.3.7. Let A be an $m \times n$ matrix. The following three operations are called **elementary column operations**:

- (1) interchanging any two columns of A ;
- (2) multiplying any column of A by a nonzero scalar;
- (3) adding any scalar multiple of a column of A to another column.

Depending on whether the operations is obtained from (1), (2) or (3) we will say that we performed an operation of **type 1**, **type 2** or **type 3**.

Since elementary row operations correspond to left multiplication by elementary matrices by Thm. 4.2.3, it seems plausible that something similar should be true for elementary column operations as well. We will use the transpose of a matrix to prove the corresponding results. Given a matrix $A \in M_{m \times n}(F)$ with entries A_{ij} in a field F , the transpose of A is the matrix $B \in M_{n \times m}(F)$ with entries $B_{ij} = A_{ij}$, i.e. the i th column of B is the i th row of A . We also write A^t for B .

Example 4.3.8. Let A be the following matrix with entries in \mathbb{R} :

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix}$$

The transpose of A is given by

$$A^t = \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}$$

Example 4.3.9. In this example we will study how the product of two matrices behaves with respect to transposition. Let A be the matrix from Example 4.3.8 and let $B \in M_{3 \times 2}(\mathbb{R})$ be given by

$$B = \begin{pmatrix} -1 & 0 \\ 1 & 1 \\ 0 & 0 \end{pmatrix}.$$

Then we have

$$AB = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \begin{pmatrix} -1 & 0 \\ 1 & 1 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 2 \\ 1 & 5 \end{pmatrix}.$$

The transpose of this matrix is

$$(AB)^t = \begin{pmatrix} 1 & 1 \\ 2 & 5 \end{pmatrix}.$$

Compare this result with the following

$$B^t A^t = \begin{pmatrix} -1 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 2 & 5 \end{pmatrix}.$$

Thus, we get $(AB)^t = B^t A^t$. In fact, this is not a coincidence as the next Lemma shows.

Lemma 4.3.10. *Let F be a field and let $A \in M_{m \times n}(F)$, $B \in M_{n \times r}(F)$ be matrices with entries in F . Then*

$$(AB)^t = B^t A^t.$$

Proof. For any matrix $C \in M_{s \times t}(F)$ denote by C_{ij} for $1 \leq i \leq s$ and $1 \leq j \leq t$ the entry in the position (i, j) of the matrix C . Then

$$\begin{aligned} (B^t A^t)_{ik} &= \sum_{j=1}^n (B^t)_{ij} (A^t)_{jk} = \sum_{j=1}^n B_{ji} A_{kj} \\ &= \sum_{j=1}^n A_{kj} B_{ji} = (AB)_{ki} = ((AB)^t)_{ik} . \quad \square \end{aligned}$$

Lemma 4.3.11. *Let F be a field, $n \in \mathbb{N}$ and let $E \in M_{n \times n}(F)$ be an elementary matrix (see Def. 4.2.1). Then E^t is an elementary matrix of the same type.*

Proof. Let $I_n \in M_{n \times n}(F)$ be the identity matrix. Observe that if we view the i th row of I_n as a column vector, it agrees with the i th column of I_n . Let E be obtained by interchanging the rows i and j of I_n . This means that the i th column of E has a 1 in position j and only zeroes otherwise. Likewise, the j th column of E has a 1 in position i and zeroes otherwise. In particular, the effect of interchanging the i th and j th row is the same as interchanging the i th and j th column. Thus, for elementary matrices of type 1 we have $E^t = E$.

Since the only nonzero entries of I_n are on the diagonal, multiplying the i th row of I_n by a nonzero scalar yields the same result as multiplying the i th column by the same scalar. Therefore $E^t = E$ for elementary matrices of type 2 as well.

Now let E be a matrix obtained from I_n by adding λ times the i th row to the j th row for some scalar $\lambda \in F$ with $\lambda \neq 0$. The matrix E^t is then obtained from I_n by adding λ times the i th column to the j th column of I_n . Hence, the only row with non-zero off-diagonal entries is the i th row, which has a λ in position j . Therefore this is the same matrix that is obtained by adding λ times the j th row of I_n to the i th row. In particular, E^t is again an elementary matrix of type 3 (This time it is a different one!). \square

Theorem 4.3.12. *Let F be a field and let $A \in M_{m \times n}(F)$. Suppose that $B \in M_{m \times n}(F)$ is obtained from A by performing an elementary column operation. Let $E \in M_{n \times n}(F)$ be the elementary matrix obtained from I_n by performing the same elementary column operation. Then we have*

$$AE = B .$$

Proof. Observe that $B^t \in M_{n \times m}(F)$ is obtained from $A^t \in M_{n \times m}(F)$ by performing an elementary row operation. Now Thm. 4.2.3 implies that

$$B^t = \tilde{E}A^t$$

where the matrix $\tilde{E} \in M_{n \times n}(F)$ is the matrix obtained from I_n by performing the same elementary row operation. Therefore $E = \tilde{E}^t$ is the matrix obtained from I_n by performing the same elementary column operation as the one used to obtain B from A . Note that this is again an elementary matrix by Lemma 4.3.11. We obtain using Lemma 4.3.10

$$B = (\tilde{E}A^t)^t = (A^t)^t \tilde{E}^t = AE . \quad \square$$

Corollary 4.3.13. *Elementary row and column operations preserve the rank of a matrix.*

Proof. Let F be a field and let $A \in M_{m \times n}(F)$. Let $E \in M_{m \times m}(F)$, $G \in M_{n \times n}(F)$ be elementary matrices. By Corollary 4.3.4 we have

$$\begin{aligned}\text{rank}(EA) &= \text{rank}(A) , \\ \text{rank}(AG) &= \text{rank}(A) .\end{aligned}$$

The statement follows from Thm. 4.2.3 and Thm. 4.3.12. \square

Example 4.3.14. Let $A \in M_{3 \times 3}(\mathbb{R})$ be the following matrix

$$A = \begin{pmatrix} 1 & 2 & 1 \\ 1 & 0 & 3 \\ 1 & 1 & 2 \end{pmatrix}$$

We would like to determine the rank of the matrix A . By Corollary 4.3.13 we are allowed to perform row and column operations without changing the rank. For example, we could subtract the first row from the second and third to get

$$\begin{pmatrix} 1 & 2 & 1 \\ 0 & -2 & 2 \\ 0 & -1 & 1 \end{pmatrix} .$$

We can now add multiples of the first column to the second and third to obtain

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & -2 & 2 \\ 0 & -1 & 1 \end{pmatrix} .$$

If we now multiply the second row by $\frac{1}{2}$ and subtract it from the third we get

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & 0 \end{pmatrix} .$$

Finally, we can add the second column to the third and afterwards multiply the second column by (-1) . This yields the matrix

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} .$$

By the corollary this matrix has the same rank as A . But we can easily read off the rank of this matrix. It is 2.

In the last example we could use row and column operations to reduce the matrix to one, whose non-zero entries are on the diagonal and are equal to 1 or 0. We will see in the next theorem that this can always be achieved.

Theorem 4.3.15. *Let F be a field and let $A \in M_{m \times n}(F)$. Then, by means of a finite number of row and column operations, A can be transformed into a matrix that has one of the following forms*

$$D = \begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix} \quad , \quad D = \begin{pmatrix} I_r \\ 0 \end{pmatrix} \quad , \quad D = (I_r \ 0) \quad , \quad D = I_r \quad , \quad D = 0 \quad (13)$$

for some $r \in \{1, \dots, n\}$ where I_r is the $r \times r$ -matrix with ones along the diagonal and zeroes everywhere else. The rank of A is then equal to r , or it is zero in case $D = 0$.

Proof. If A is the zero matrix, we are done. Suppose $A \neq 0$. We can use elementary row operations (see Def. 4.1.1) to bring A into reduced row echelon form (see Def. 4.1.4). Interchanging columns (i.e. using elementary column operations of type 1) if necessary we can now turn the columns containing 1 as their only nonzero entry into the leftmost columns of the matrix. The result will be one of the following matrices

$$\begin{pmatrix} I_r & B \\ 0 & 0 \end{pmatrix} \quad , \quad \begin{pmatrix} I_r \\ 0 \end{pmatrix} \quad , \quad (I_r \ B) \quad , \quad I_r$$

for some $r \in \{1, \dots, n\}$ and some matrix $B \in M_{r \times (n-r)}(F)$. The r columns of I_r form the standard basis for F^r . Therefore any column of B can be represented as a linear combination of the r columns of I_r . Hence, we can use type 3 column operations to turn B into another zero matrix. The resulting matrix D will then be of the form mentioned in the statement. Since we only used elementary row and column operations to obtain this form and the rank of D is clearly equal to r , we have $\text{rank}(A) = r$ by Cor. 4.3.13. \square

Since elementary row and column operations correspond to multiplications by invertible matrices on the left or right, respectively, we can rephrase the statement of the last theorem in the a way that will become very useful in the next chapter about the determinant.

Corollary 4.3.16. *Let F be a field and let $A \in M_{m \times n}(F)$ with $\text{rank}(A) = r$. Then there exist invertible matrices $B \in M_{m \times m}(F)$ and $C \in M_{n \times n}(F)$ such that $D = BAC$, where D is one of the matrices in the list (13).*

Proof. By Theorem 4.3.15, A can be transformed by means of a finite number of row and column operations into one of the matrices in the list (13). Combining Thm. 4.2.3 and Thm. 4.3.12 we see that each of these row and column operations corresponds to a left multiplication by an elementary matrix $E_i \in M_{m \times m}(F)$ for $i \in \{1, \dots, p\}$ or a right multiplication by an elementary matrix $G_j \in M_{n \times n}(F)$ for $j \in \{1, \dots, q\}$. By Thm. 4.2.4 all these matrices are invertible. Let $B = E_p E_{p-1} \dots E_1$ and $C = G_1 G_2 \dots G_q$. Then B and C are invertible and

$$D = E_p E_{p-1} \dots E_1 A G_1 G_2 \dots G_q = BAC \quad . \quad \square$$

Looking back at Example 4.3.6 we make the following curious observation: We already noted that the maximal number of linearly independent columns of the matrix A is 2. The row vectors of that matrix (written as column vectors) are

$$r_1 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad r_2 = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}, \quad r_3 = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}.$$

Hence, $r_1 = r_3$ and the set $\{r_1, r_2\}$ is linearly independent. Therefore the maximal number of linearly independent rows is also 2. This is in fact not a coincidence and we are now finally in the position to explain this phenomenon in the next corollary.

Corollary 4.3.17. *Let F be a field and let $A \in M_{m \times n}(F)$. Then $\text{rank}(A) = \text{rank}(A^t)$.*

Proof. If $X \in M_{r \times r}(F)$ is an invertible matrix, then so is X^t . In fact, $(X^{-1})^t$ is an inverse of X^t by the following computation that uses Lemma 4.3.10

$$\begin{aligned} X^t(X^{-1})^t &= (X^{-1}X)^t = I_r^t = I_r, \\ (X^{-1})^t X^t &= (XX^{-1})^t = I_r^t = I_r. \end{aligned}$$

By Cor. 4.3.16, there exist invertible matrices $B \in M_{m \times m}(F)$ and $C \in M_{n \times n}(F)$ such that $D = BAC$, where D is one of the matrices in the list (13). Taking transposes, we have

$$D^t = (BAC)^t = C^t A^t B^t.$$

By the first observation B^t, C^t are both invertible. By Cor. 4.3.4, $\text{rank}(D^t) = \text{rank}(A^t)$ and $\text{rank}(D) = \text{rank}(BAC) = \text{rank}(A)$. But for all matrices D in the list (13) we have $\text{rank}(D) = \text{rank}(D^t)$. Thus, $\text{rank}(A) = \text{rank}(A^t)$. \square

The next theorem is a structural result about invertible matrices that is quite surprising. But this is more than just an interesting fact: It will help us in the next chapter.

Corollary 4.3.18. *Every invertible matrix is a product of elementary matrices.*

Proof. Suppose that $A \in M_{n \times n}(F)$ is an invertible matrix. By Corollary 4.3.16, there are invertible matrices $B \in M_{n \times n}(F)$ and $C \in M_{n \times n}(F)$ such that $D = BAC$. By the proof of Corollary 4.3.16 B and C can be chosen to be products of elementary matrices, i.e. $B = E_p E_{p-1} \dots E_1$ and $C = G_1 G_2 \dots G_q$.

If A is invertible, D is invertible as well. But the only invertible matrix in the list (13) is $D = I_n$. Moreover, the inverse of an elementary matrix is an elementary matrix as well by Thm. 4.2.4. Therefore,

$$\begin{aligned} A &= B^{-1} D C^{-1} = (E_p E_{p-1} \dots E_1)^{-1} I_n (G_1 G_2 \dots G_q)^{-1} \\ &= E_1^{-1} \dots E_{p-1}^{-1} E_p^{-1} G_q^{-1} \dots G_2^{-1} G_1^{-1} \end{aligned}$$

is a product of elementary matrices. \square

5 Determinants

Historically, the determinant was first defined as a number associated to a system of n linear equations in n unknowns that determines whether the system has a unique solution. This is based on the following observation:

Theorem 5.0.1. *Let F be a field and let $A \in M_{m \times n}(F)$ and $b \in F^m$. Let $s \in F^n$ be a solution of the system of linear equations $Ax = b$. Then the solution set S of the system is given by*

$$S = \{s + v \in F^n \mid v \in \ker(L_A)\} .$$

In particular, $Ax = b$ has a unique solution if and only there is at least one solution $s \in F^n$ and $\ker(L_A) = \{0\}$. If $m = n$, then $Ax = b$ has a unique solution if and only if A is invertible.

Proof. First observe that $\ker(L_A) \subset F^n$ is the solution set of the homogeneous system of linear equations $Ax = 0$. If $s' \in F^n$ is a solution of $Ax = b$, then

$$A(s' - s) = As' - As = b - b = 0 .$$

Thus, $v = s' - s \in \ker(L_A)$ and $s' = s + v$. This proves $S \subseteq \{s + v \in F^n \mid v \in \ker(L_A)\}$. On the other hand if $v \in \ker(L_A)$, then

$$A(s + v) = As + Av = b + 0 = b .$$

Therefore, $\{s + v \in F^n \mid v \in \ker(L_A)\} \subseteq S$ as well, which directly implies the statement about the uniqueness of solutions.

Now suppose $m = n$. If A is invertible, then $s = A^{-1}b$ satisfies $As = AA^{-1}b = b$ and is a solution of the system. Moreover, the matrix representation of L_A with respect to the standard basis of F^n agrees with A (This is a very good exercise!). Thus, the invertibility of A is equivalent to the one of L_A by Thm. 3.0.10, which in turn is equivalent to $\ker(L_A) = \{0\}$ by Thm. 1.0.14. This implies both directions of the second statement. \square

The above theorem implies that a system of n linear equations in n unknowns has a unique solution if and only if the $n \times n$ -matrix A is invertible. Let us explore this question further for 2×2 -matrices: Let $A \in M_{2 \times 2}(F)$ be given by

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

This matrix is invertible if and only if the scalar $\det(A) = ad - bc \in F$ is not zero. In fact, if $\det(A) \neq 0$, then

$$B = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} = \begin{pmatrix} \frac{d}{ad-bc} & \frac{-b}{ad-bc} \\ \frac{-c}{ad-bc} & \frac{a}{ad-bc} \end{pmatrix}$$

is an inverse of A , i.e. $AB = I_2$ and $BA = I_2$. On the other hand if $\det(A) = 0$, then the vector

$$v = \begin{pmatrix} d \\ -c \end{pmatrix}$$

satisfies $Av = 0$. Thus, $\ker(L_A) \neq \{0\}$ and A can not be invertible. The number $\det(A)$ is called the determinant of the 2×2 -matrix A .

Remark 5.0.2. It is important to note here that $A \mapsto \det(A)$ is **not** a linear transformation. To see this let $A = I_2$ and $B = -I_2$. Then the sum $A + B$ is the zero matrix. In particular, it is not invertible and $\det(A + B) = 0$. However, $\det(A) = 1 = \det(B)$. Therefore $\det(A) + \det(B) = 2$.

In the following definition we will generalise the determinant to arbitrary dimensions. To characterise it we have to think of \det as a map that takes the row vectors of a matrix A and gives a scalar in F . If $r_i \in F^n$ for $i \in \{1, \dots, n\}$ are given row vectors, then we will write

$$\det(r_1, \dots, r_n)$$

for $\det(A)$, where A is the matrix with i th row given by the row vector r_i .

Definition 5.0.3. Let $n \in \mathbb{N}$ and let F be a field. The **determinant** is a map $\det: M_{n \times n}(F) \rightarrow F$ that is uniquely determined by the following properties:

- 1) Suppose $r_1, \dots, r_n \in F^n$ and $y \in F^n$ are row vectors and $c \in F$. Then for any $i \in \{1, \dots, n\}$, \det satisfies

$$\begin{aligned} & \det(r_1, \dots, r_{i-1}, cr_i + y, r_{i+1}, \dots, r_n) \\ &= c \det(r_1, \dots, r_{i-1}, r_i, r_{i+1}, \dots, r_n) + \det(r_1, \dots, r_{i-1}, y, r_{i+1}, \dots, r_n), \end{aligned}$$

i.e. \det is a linear transformation when considered as a map on a single row of a matrix while keeping the other rows fixed.

- 2) If two rows of a matrix A agree, then $\det(A) = 0$, i.e. if $r_1, \dots, r_n \in F^n$ are the rows of A and $r_i = r_j = y$ for some $i, j \in \{1, \dots, n\}$, then

$$\det(A) = \det(r_1, \dots, y, \dots, y, \dots, r_n) = 0.$$

- 3) The determinant of I_n is 1.

Remark 5.0.4. There is a theorem hiding behind the word “uniquely” in the above definition. We have to check that the determinant is indeed the only map with these three properties. It is also not clear from this definition that such a map should even exist in every dimension. We will see in the next example, that our definition in dimension 2 works and will solve the general case later.

Example 5.0.5. Let us check that the determinant of a 2×2 -matrix has the properties stated in the definition. Let $A \in M_{2 \times 2}(F)$ be given by

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

Let $y = (a' \ b')$ be a row vector and $k \in F$. Then

$$\det \begin{pmatrix} ka + a' & kb + b' \\ c & d \end{pmatrix} = (ka + a')d - c(kb + b') = k(ad - cb) + (a'd - cb')$$

$$k \det \begin{pmatrix} a & b \\ c & d \end{pmatrix} + \det \begin{pmatrix} a' & b' \\ c & d \end{pmatrix} = k(ad - cb) + (a'd - cb')$$

The proof for the second row is very similar. Therefore the first property holds. It follows from the formula that $\det(I_2) = 1$. Therefore the third property is also satisfied. Now note that

$$\det \begin{pmatrix} a & b \\ a & b \end{pmatrix} = ab - ab = 0 .$$

Hence, the second property is also true for \det .

In the next few theorems we will start proving that the properties 1) to 3) in Definition 5.0.3 fix the map \det uniquely. To achieve this we will look at arbitrary maps $\delta: M_{n \times n}(F) \rightarrow F$ with properties 1) to 3) and deduce step by step that they all have to agree. We start by making the following helpful observations:

Theorem 5.0.6. *Let F be a field. Let $\delta: M_{n \times n}(F) \rightarrow F$ be a map with the properties 1) to 3) stated in Def. 5.0.3.*

- a) *If $A \in M_{n \times n}(F)$ is a matrix containing a zero row, then $\delta(A) = 0$.*
- b) *If $A \in M_{n \times n}(F)$ is any matrix and $B \in M_{n \times n}(F)$ is obtained from A by interchanging two rows of A , then $\delta(B) = -\delta(A)$.*
- c) *If $A \in M_{n \times n}(F)$ is any matrix and B is obtained from A by multiplying the i th row by a scalar $\lambda \in F$, $\lambda \neq 0$, then $\delta(B) = \lambda \delta(A)$.*
- d) *If $A \in M_{n \times n}(F)$ is any matrix and B is obtained from A by adding λ times the i th row to the j th, then $\delta(A) = \delta(B)$.*
- e) *If $A \in M_{n \times n}(F)$ is a matrix that is not invertible, then $\delta(A) = 0$.*

Proof. To prove a) let $r_1, \dots, r_n \in F^n$ be the rows of A and let $i \in \{1, \dots, n\}$ be such that $r_i = (0 \ \dots \ 0)$. Note that $r_i = 0r_i$. Therefore by property 1) of δ (with $c = 0$ and $y = 0$):

$$\delta(r_1, \dots, r_i, \dots, r_n) = \delta(r_1, \dots, 0r_i, \dots, r_n) = 0\delta(r_1, \dots, r_i, \dots, r_n) = 0 .$$

To prove b) let $A \in M_{n \times n}(F)$ and let $r_1, \dots, r_n \in F^n$ be the rows of A . Let $i, j \in \{1, \dots, n\}$ be the indices of the two rows that are interchanged to obtain B and suppose

$i < j$. Using properties 1) and 2) in Def. 5.0.3 we obtain

$$\begin{aligned}
0 &= \delta(r_1, \dots, \underbrace{\mathbf{r}_i + \mathbf{r}_j}_i, \dots, \underbrace{\mathbf{r}_i + \mathbf{r}_j}_j, \dots, r_n) \\
&= \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_i + \mathbf{r}_j, \dots, r_n) + \delta(r_1, \dots, \mathbf{r}_j, \dots, \mathbf{r}_i + \mathbf{r}_j, \dots, r_n) \\
&= \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_i, \dots, r_n) + \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots, r_n) \\
&\quad + \delta(r_1, \dots, \mathbf{r}_j, \dots, \mathbf{r}_i, \dots, r_n) + \delta(r_1, \dots, \mathbf{r}_j, \dots, \mathbf{r}_j, \dots, r_n) \\
&= \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots, r_n) + \delta(r_1, \dots, \mathbf{r}_j, \dots, \mathbf{r}_i, \dots, r_n) \\
&= \delta(A) + \delta(B) .
\end{aligned}$$

The statement in c) follows directly from property 1) in Def. 5.0.3 (with $c = \lambda$ and $y = 0$). Now let B be obtained from A as in d) and let $r_1, \dots, r_n \in F^n$ be the rows of A . Then

$$\begin{aligned}
\delta(B) &= \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_j + \lambda \mathbf{r}_i, \dots, r_n) \\
&= \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots, r_n) + \lambda \delta(r_1, \dots, \mathbf{r}_i, \dots, \mathbf{r}_i, \dots, r_n) \\
&= \delta(A) + 0 ,
\end{aligned}$$

where we used property 1) of Def. 5.0.3 for the second equality and property 2) for the last one. This proves d).

Let $A \in M_{n \times n}(F)$ be a matrix that is not invertible. By a sequence of elementary row operations we can then bring A into reduced row echelon form. Call the resulting matrix B . By part b) to d) of this theorem $\delta(B)$ is equal to $\lambda \delta(A)$ for a scalar $\lambda \in F$ with $\lambda \neq 0$. But since A is not invertible, the matrix B will contain a row of zeroes. Thus, we have $\delta(B) = 0$ by part a) and hence $\delta(A) = 0$ as well. \square

Exercise 5.0.7. Let $A, B \in M_{n \times n}(F)$. Suppose that A is not invertible. Prove that AB and BA are also not invertible.

Theorem 5.0.8. Let F be a field. Let $\delta: M_{n \times n}(F) \rightarrow F$ be a map with the properties 1) to 3) stated in Def. 5.0.3. Let $A, B \in M_{n \times n}(F)$. Then

$$\delta(AB) = \delta(A) \delta(B) .$$

Proof. We will first prove this theorem in the special case that $A = E$ is an elementary matrix (see Def. 4.2.1). In this case EB is the matrix obtained from B by the same elementary row operation that was used to obtain E from I_n (see Thm. 4.2.3).

If E is a type 1 elementary matrix, then Thm. 5.0.6 b) and property 3) in Def. 5.0.3 imply $\delta(E) = -1$. If E is of type 2, i.e. obtained from I_n by multiplying a row by a nonzero scalar $\lambda \in F$, then Thm. 5.0.6 c) and property 3) combine to give $\delta(E) = \lambda$. In case E is a type 3 elementary matrix, we obtain $\delta(E) = 1$ in the same way using Thm. 5.0.6 d). Thus, in each case, we obtain by combining Thm. 4.2.3 and Thm. 5.0.6 b) to d) that

$$\delta(EB) = \delta(E) \delta(B) .$$

Now suppose that $A \in M_{n \times n}(F)$ is any matrix. If A or B is not invertible, then so is AB by Exercise 5.0.7. Thus, in both cases $\delta(A)\delta(B) = \delta(AB) = 0$ by Thm. 5.0.6 e). Hence, the theorem only remains to be proven for invertible matrices A and B . If A is invertible, then it is a product of elementary matrices, say $A = E_1 \dots E_p$, by Cor. 4.3.18. But by our previous observation

$$\begin{aligned} \delta(AB) &= \delta(E_1 \dots E_p B) = \delta(E_1)\delta(E_2 \dots E_p B) \\ &= \delta(E_1) \dots \delta(E_p)\delta(B) = \delta(E_1) \dots \delta(E_{p-1}E_p)\delta(B) \\ &= \delta(E_1 \dots E_p)\delta(B) = \delta(A)\delta(B) . \end{aligned}$$

This proves the statement for all A and B . □

Theorem 5.0.9. *Let F be a field. Let $\delta_1, \delta_2: M_{n \times n}(F) \rightarrow F$ be two maps with the properties 1) to 3) stated in Def. 5.0.3. Then $\delta_1 = \delta_2$. In particular, $\det = \delta_1 = \delta_2$ is uniquely determined by properties 1) to 3).*

Proof. We need to show that $\delta_1(A) = \delta_2(A)$ for all $A \in M_{n \times n}(F)$. If A is not invertible, then we have $\delta_1(A) = \delta_2(A) = 0$ by Thm. 5.0.6 e). If $A = E \in M_{n \times n}(F)$ is an elementary matrix, then

- $\delta_1(E) = \delta_2(E) = -1$ if E is of type 1 by Thm. 5.0.6 b) and property 3) in Def. 5.0.3,
- $\delta_1(E) = \delta_2(E) = \lambda$ if E is of type 2 with multiplication by $\lambda \in F$, $\lambda \neq 0$ by Thm. 5.0.6 c) and property 3) in Def. 5.0.3,
- $\delta_1(E) = \delta_2(E) = 1$ if E is of type 3 by Thm. 5.0.6 d) and property 3) in Def. 5.0.3.

If $A \in M_{n \times n}(F)$ is an arbitrary invertible matrix, then $A = E_1 \dots E_p$ for elementary matrices $E_i \in M_{n \times n}(F)$ by Cor. 4.3.18. Thus, by Thm. 5.0.8

$$\delta_1(A) = \delta_1(E_1 \dots E_p) = \delta_1(E_1) \dots \delta_1(E_p) = \delta_2(E_1) \dots \delta_2(E_p) = \delta_2(E_1 \dots E_p) = \delta_2(A) .$$

This proves the statement for all $A \in M_{n \times n}(F)$. □

Exercise 5.0.10. Let $A \in M_{n \times n}(F)$. Prove that $\det(A) = \det(A^t)$ using Theorem 5.0.6 and Theorem 5.0.8.

5.1 That horrible formula

We have seen in Thm. 5.0.9 that the determinant is uniquely determined by properties 1) to 3) in Def. 5.0.3. What we still do not know up to this point is whether there **exists** a map $\det: M_{n \times n}(F) \rightarrow F$ with these properties. In your *Vectors and Matrices* lecture you have already met a definition of the determinant. We will review it below and show that it indeed has the properties list in Def. 5.0.3.

Remember that a permutation is a bijective map $\sigma: \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ from the set of the first n natural numbers to itself. They form a group, denoted by S_n . A transposition is a permutation switching just two numbers and leaving the rest fixed.

For example, $\tau \in S_3$ with $\tau(1) = 2$, $\tau(2) = 1$ and $\tau(3) = 3$ is a transposition. Any permutation σ can be obtained as a composition of finitely many transpositions. However, this decomposition into transpositions is not unique. The sign $\text{sign}(\sigma) \in \{\pm 1\}$ of a permutation is given by $(-1)^{k(\sigma)}$, where $k(\sigma)$ is the number of transpositions in a decomposition of σ , i.e. $\sigma = \tau_1 \dots \tau_{k(\sigma)}$. Even though the decomposition is not unique, $(-1)^{k(\sigma)}$ is well-defined.

Example 5.1.1. The permutation $\sigma \in S_3$ given by $\sigma(1) = 2$, $\sigma(2) = 3$ and $\sigma(3) = 1$ has sign 1, since it can be obtained as the composition of two transpositions: The first transposition switches 1 and 2 and the second transposition interchanges 2 and 3.

We are now ready to show that the determinant exists by giving an explicit formula:

Theorem 5.1.2. *Let F be a field. For a matrix $A \in M_{n \times n}(F)$ denote the entries of A by $a_{i,j} \in F$. Let $\delta: M_{n \times n}(F) \rightarrow F$ be given by*

$$\delta(A) = \sum_{\sigma \in S_n} \text{sign}(\sigma) a_{1,\sigma(1)} \dots a_{n,\sigma(n)} .$$

Then δ has the three properties listed in Def. 5.0.3. In particular, there is a map with the properties listed in Def. 5.0.3 and $\delta = \det$.

Proof. Let r_i be the i th row of A . Then $r_i = (a_{i,1}, \dots, a_{i,n})$. Let $c \in F$ and let $y = (y_1, \dots, y_n) \in F^n$ be another row vector. Then

$$\begin{aligned} & \delta(r_1, \dots, \mathbf{c}\mathbf{r}_i + \mathbf{y}, \dots, r_n) \\ &= \sum_{\sigma \in S_n} \text{sign}(\sigma) a_{1,\sigma(1)} \dots (\mathbf{c}\mathbf{a}_{i,\sigma(i)} + \mathbf{y}_{\sigma(i)}) \dots a_{n,\sigma(n)} \\ &= \mathbf{c} \left(\sum_{\sigma \in S_n} \text{sign}(\sigma) a_{1,\sigma(1)} \dots \mathbf{a}_{i,\sigma(i)} \dots a_{n,\sigma(n)} \right) + \sum_{\sigma \in S_n} \text{sign}(\sigma) a_{1,\sigma(1)} \dots \mathbf{y}_{\sigma(i)} \dots a_{n,\sigma(n)} \\ &= \mathbf{c} \delta(r_1, \dots, \mathbf{r}_i, \dots, r_n) + \delta(r_1, \dots, \mathbf{y}, \dots, r_n) \end{aligned}$$

and therefore δ satisfies property 1) of Def. 5.0.3.

To prove that property 2) also holds, let r_i be the rows of the matrix A and assume that for two distinct $i, j \in \{1, \dots, n\}$ we have $r_i = r_j = y$. In terms of the entries of A , this means that $a_{i,s} = a_{j,s} = y_s$ for all $s \in \{1, \dots, n\}$. The summand in the definition of δ corresponding to $\sigma \in S_n$ is of the form

$$\text{sign}(\sigma) a_{1,\sigma(1)} \dots a_{i-1,\sigma(i-1)} a_{i+1,\sigma(i+1)} \dots a_{j-1,\sigma(j-1)} a_{j+1,\sigma(j+1)} \dots a_{n,\sigma(n)} y_{\sigma(i)} y_{\sigma(j)} , \quad (14)$$

where we moved the factors arising from y_s to the end. Let $\tau \in S_n$ be the transposition that interchanges i and j . The elements $\sigma \in S_n$ and $\sigma \circ \tau \in S_n$ lead to summands that have the same absolute value as the one in (14), but different signs, since $\sigma \circ \tau$ is σ composed with an additional transposition. Therefore all summands cancel and the sum evaluates to zero. This proves that δ has property 2) from Def. 5.0.3.

The entry $a_{i,j} \in F$ of the matrix $A = I_n$ is equal to 1 if $i = j$ and equal to 0 otherwise. Therefore the only summand in the definition of δ that is nonzero is the one corresponding to the trivial permutation $\iota \in S_n$ (with $\iota(k) = k$ for all $k \in \{1, \dots, n\}$), which evaluates to 1, since the sign of ι is also 1. Thus, δ has property 1) as well. \square

Example 5.1.3. Let us write an element in S_3 by juxtaposing the images of 1, 2 and 3. So, for example, (213) is the transposition that interchanges 1 and 2. Then the signs of all elements in S_3 are shown in the following list:

σ	sign(σ)
(123)	1
(231)	1
(312)	1
(213)	-1
(321)	-1
(132)	-1

Therefore the determinant of a 3×3 -matrix is given by

$$\det(A) = a_{1,1}a_{2,2}a_{3,3} + a_{1,2}a_{2,3}a_{3,1} + a_{1,3}a_{2,1}a_{3,2} - a_{1,2}a_{2,1}a_{3,3} - a_{1,3}a_{2,2}a_{3,1} - a_{1,1}a_{2,3}a_{3,2} .$$

Example 5.1.4. The group S_2 has only two elements: the trivial permutation and the transposition of 1 and 2. The neutral element has sign 1 and the transposition has sign -1 . Therefore the determinant of a 2×2 -matrix according to the above formula is

$$\det(A) = a_{1,1}a_{2,2} - a_{1,2}a_{2,1} ,$$

which is precisely the formula we met at the beginning of this chapter.

5.2 Summary of the properties of the determinant

It might be worthwhile to summarise the properties of \det , which we found along the way while proving that it is uniquely characterised by the ones listed in Def. 5.0.3. Apart from the properties listed in Def. 5.0.3, the determinant also has the following properties:

- The determinant is multiplicative, i.e. if $A, B \in M_{n \times n}(F)$, then

$$\det(AB) = \det(A) \det(B) .$$

This is Thm. 5.0.8.

- A matrix $A \in M_{n \times n}(F)$ is invertible if and only if $\det(A) \neq 0$. One direction is Thm. 5.0.6 e). For the other direction we have to show that if A is invertible, then $\det(A) \neq 0$. But this follows from the fact that $\det(E) \neq 0$ for any elementary matrix $E \in M_{n \times n}(F)$ (see proof of Thm. 5.0.9) together with Thm. 5.0.8 and Cor. 4.3.18.

- The determinant of the transpose A^t of a matrix $A \in M_{n \times n}(F)$ satisfies

$$\det(A) = \det(A^t) .$$

This is Exercise 5.0.10.

- Let $A \in M_{n \times n}(F)$ be an invertible matrix. Then $\det(A) \neq 0$ and

$$\det(A) \det(A^{-1}) = \det(AA^{-1}) = \det(I_n) = 1$$

(by Thm. 5.0.8) and therefore $\det(A^{-1}) = \det(A)^{-1}$.

- If $A, B \in M_{n \times n}(F)$ and B is obtained from A by interchanging two rows, then $\det(B) = -\det(A)$ by Thm. 5.0.6 b).
- If $A, B \in M_{n \times n}(F)$ and B is obtained from A by multiplying a row of A by $\lambda \in F$ with $\lambda \neq 0$, then $\det(B) = \lambda \det(A)$ by Thm. 5.0.6 c).
- If $A, B \in M_{n \times n}(F)$ and B is obtained from A by adding a multiple of one row of A to another row of A , then $\det(B) = \det(A)$ by Thm. 5.0.6 d).

5.3 The determinant of a linear transformation $T: V \rightarrow V$

Let V be a finite-dimensional vector space over the field F and let β be an ordered basis for V . In this section we will discuss whether the definition of the determinant can be extended from matrices to linear transformations of the form $T: V \rightarrow V$. We would like to define $\det(T)$ as follows:

$$\det(T) := \det \left([T]_{\beta}^{\beta} \right) .$$

But since the matrix $[T]_{\beta}^{\beta}$ depends on the choice of basis β for V , it is not clear whether this expression depends on our choice of β . To compare how two choices differ let γ be another ordered basis for V and let $A_{\beta} = [T]_{\beta}^{\beta}$ and $A_{\gamma} = [T]_{\gamma}^{\gamma}$. By Thm. 3.0.20 we can write T in two ways

$$T = \varphi_{\gamma}^{-1} \circ L_{A_{\gamma}} \circ \varphi_{\gamma} = \varphi_{\beta}^{-1} \circ L_{A_{\beta}} \circ \varphi_{\beta} .$$

Thus, to compare A_{γ} and A_{β} we obtain from the last computation

$$L_{A_{\gamma}} = \left(\varphi_{\gamma} \circ \varphi_{\beta}^{-1} \right) \circ L_{A_{\beta}} \circ \left(\varphi_{\gamma} \circ \varphi_{\beta}^{-1} \right)^{-1} .$$

Let $n = \dim(V)$ and let $\alpha = \{e_1, \dots, e_n\}$ be the standard basis of F^n . It is not difficult to check that the linear transformation $\varphi_{\gamma} \circ \varphi_{\beta}^{-1}: F^n \rightarrow F^n$ is given by left multiplication by the matrix

$$C_{\gamma, \beta} = \left[\varphi_{\gamma} \circ \varphi_{\beta}^{-1} \right]_{\alpha}^{\alpha} .$$

Altogether we obtain the following equation for A_{γ} and A_{β}

$$A_{\gamma} = C_{\gamma, \beta} A_{\beta} C_{\gamma, \beta}^{-1} .$$

Taking determinants gives

$$\begin{aligned}\det\left([T]_\gamma^\gamma\right) &= \det(A_\gamma) = \det(C_{\gamma,\beta} A_\beta C_{\gamma,\beta}^{-1}) = \det(C_{\gamma,\beta}) \det(A_\beta) \det(C_{\gamma,\beta}^{-1}) \\ &= \det(C_{\gamma,\beta}) \det(A_\beta) \det(C_{\gamma,\beta})^{-1} = \det(A_\beta) \\ &= \det\left([T]_\beta^\beta\right) .\end{aligned}$$

Therefore the expression $\det\left([T]_\beta^\beta\right)$ does not depend on the choice of basis and our definition of the determinant of T makes sense.

5.4 Minors and cofactor expansion – a nice formula for a change

There are several ways to compute the determinant of a square matrix $A \in M_{n \times n}(F)$ that are computationally more efficient than the one given in Thm. 5.1.2. In this section we will learn another formula for \det that is very efficient in computing the determinant of a matrix which contains a row or a column with many zeroes.

Example 5.4.1. Consider the following quite randomly chosen matrix $A \in M_{3 \times 3}(\mathbb{R})$:

$$A = \begin{pmatrix} 4 & 1 & 9 \\ 2 & 3 & 3 \\ 7 & 1 & 0 \end{pmatrix}$$

For each $i, j \in \{1, 2, 3\}$ let $M_{ij} \in M_{2 \times 2}(\mathbb{R})$ the matrix obtained from A by removing the i th row and the j th column. For example, to get M_{21} we need to remove the second row and the first column, i.e.

$$\begin{pmatrix} \color{red}{4} & \color{red}{1} & \color{red}{9} \\ \color{red}{2} & \color{red}{3} & \color{red}{3} \\ 7 & 1 & 0 \end{pmatrix} \rightsquigarrow M_{21} = \begin{pmatrix} 1 & 9 \\ 1 & 0 \end{pmatrix}$$

Define the scalar c_{ij} as follows:

$$c_{ij} = (-1)^{i+j} \det(M_{ij}) .$$

For c_{21} in the above example we obtain

$$c_{21} = (-1)^{2+1} \det \begin{pmatrix} 1 & 9 \\ 1 & 0 \end{pmatrix} = -(-9) = 9 .$$

The nine values c_{ij} computed from A this way are listed below.

$$\begin{array}{lll} c_{11} = -3 & c_{12} = 21 & c_{13} = -19 \\ c_{21} = 9 & c_{22} = -63 & c_{23} = 3 \\ c_{31} = -24 & c_{32} = 6 & c_{33} = 10 \end{array}$$

Now we look at the first column of A and compute the following sum:

$$a_{11}c_{11} + a_{21}c_{21} + a_{31}c_{31} = 4 \cdot (-3) + 2 \cdot 9 + 7 \cdot (-24) = -162 .$$

If we perform the same computation with the second column we get:

$$a_{12}c_{12} + a_{22}c_{22} + a_{32}c_{32} = 1 \cdot 21 + 3 \cdot (-63) + 1 \cdot 6 = -162 .$$

Let us try the analogous computation for the first row instead of the first column, i.e.

$$a_{11}c_{11} + a_{12}c_{12} + a_{13}c_{13} = 4 \cdot (-3) + 1 \cdot 21 + 9 \cdot (-19) = -162 .$$

We see that we always end up with the same result. This might just be a coincidence or a special feature of the matrix A we chose at the beginning. However, we will see later that it is not! In fact, in this example we have $\det(A) = -162$ and all the numbers we computed above therefore agree with the determinant.

Example 5.4.2. Consider an arbitrary matrix $A \in M_{2 \times 2}(\mathbb{R})$, i.e.

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

Let M_{ij} again be the matrix obtained from A by removing the i th row and the j th column. Note that these are 1×1 -matrices over \mathbb{R} , i.e. real numbers. We have

$$M_{11} = d \quad , \quad M_{21} = b \quad , \quad M_{12} = c \quad , \quad M_{22} = a .$$

As in Example 5.4.1 we define the numbers c_{ij} to be

$$c_{ij} = (-1)^{i+j} \det(M_{ij}) .$$

Since the determinant of a 1×1 -matrix is just the number itself, we obtain

$$c_{11} = d \quad , \quad c_{21} = -b \quad , \quad c_{12} = -c \quad , \quad c_{22} = a .$$

If we now perform the computation from the above example for the first column, we end up with

$$a_{11}c_{11} + a_{21}c_{21} = a \cdot d + c \cdot (-b) = ad - bc = \det(A) .$$

The same expansion over the second row of A yields

$$a_{21}c_{21} + a_{22}c_{22} = c \cdot (-b) + d \cdot a = ad - bc = \det(A) .$$

It is straightforward to check that this computation **always** yields the determinant no matter which row or column we pick.

Definition 5.4.3. Let F be a field and let $A \in M_{n \times n}(F)$ be a square matrix with entries in F . Let M_{ij} for $i, j \in \{1, \dots, n\}$ be the matrix obtained from A by removing the i th row and the j th column. The scalars

$$m_{ij} = \det(M_{ij})$$

for $i, j \in \{1, \dots, n\}$ are called the **minors** of A . The scalars c_{ij} given by

$$c_{ij} = (-1)^{i+j} m_{ij} = (-1)^{i+j} \det(M_{ij})$$

are called the **cofactors** of the matrix A .

Definition 5.4.4. Let F be a field and let $A \in M_{n \times n}(F)$ be a square matrix with entries $a_{ij} \in F$. Denote the cofactors of A by c_{ij} . For a fixed $j \in \{1, \dots, n\}$ the sum

$$\sum_{i=1}^n a_{ij} c_{ij} = a_{1j} c_{1j} + \dots + a_{nj} c_{nj}$$

is called the **cofactor expansion along the j th column**. Likewise for a fixed $i \in \{1, \dots, n\}$ the sum

$$\sum_{j=1}^n a_{ij} c_{ij} = a_{i1} c_{i1} + \dots + a_{in} c_{in}$$

is called the **cofactor expansion along the i th row**.

Examples 5.4.1 and 5.4.2 give rise to the conjecture that the cofactor expansion along any row or any column provides a recursive way of computing the determinant. This is in fact true as the next theorem shows.

Theorem 5.4.5. Let F be a field and let $A \in M_{n \times n}(F)$ be a square matrix with entries $a_{ij} \in F$. Denote the cofactors of A by c_{ij} . For a fixed $j \in \{1, \dots, n\}$ we have

$$\det(A) = \sum_{i=1}^n a_{ij} c_{ij} ,$$

i.e. the cofactor expansion along the j th column agrees with the determinant. Likewise for any fixed $i \in \{1, \dots, n\}$ we obtain

$$\det(A) = \sum_{j=1}^n a_{ij} c_{ij} ,$$

i.e. the cofactor expansion along the i th row also agrees with the determinant.

Proof. First note that the second statement is a consequence of the first one. Indeed, we have $\det(A) = \det(A^t)$ by Exercise 5.0.10 and the cofactors c'_{ij} of A^t satisfy $c'_{ij} = c_{ji}$. Therefore it suffices prove the statement for expansions along columns.

Fix $j \in \{1, \dots, n\}$ and let $\delta: M_{n \times n}(F) \rightarrow F$ be the map defined by $\delta(A) = \sum_{k=1}^n a_{kj} c_{kj}$, where c_{kj} are the cofactors of A . We have to show that δ has the properties listed in Def. 5.0.3. Then the statement follows from Thm. 5.0.9.

Let $i \in \{1, \dots, n\}$ and let $x, y \in F^n$. Let $A^x \in M_{n \times n}(F)$ be a matrix where the i th row is given by x . Denote by A^y the matrix having the same rows as A^x except for the i th one, which is replaced by y . Let $\lambda \in F$ and define $A^{\lambda x + y}$ in the same way with the i th row given by $\lambda x + y$. To see that Def. 5.0.3 1) holds, we have to show that

$$\delta(A^{\lambda x + y}) = \lambda \delta(A^x) + \delta(A^y) .$$

Let M_{kl}^x be the matrix obtained from A^x by removing the k th row and l th column. Define M_{kl}^y and $M_{kl}^{\lambda x + y}$ similarly by removing rows and column from A^y and $A^{\lambda x + y}$, respectively. Similarly, denote the cofactors of A^x , A^y and $A^{\lambda x + y}$ by c_{kl}^x , c_{kl}^y and $c_{kl}^{\lambda x + y}$, respectively. Note that for $k \in \{1, \dots, n\}$ with $k \neq i$ we have

$$\det(M_{kj}^{\lambda x + y}) = \lambda \det(M_{kj}^x) + \det(M_{kj}^y) ,$$

which yields

$$c_{kj}^{\lambda x + y} = \lambda c_{kj}^x + c_{kj}^y .$$

Moreover, we have $M_{ij}^{\lambda x + y} = M_{ij}^x = M_{ij}^y$, because to obtain this matrix we have to remove the i th row, which is the only one that is different in A^x , A^y and $A^{\lambda x + y}$. Therefore, $c_{ij}^x = c_{ij}^y = c_{ij}^{\lambda x + y}$. If we denote the matrix entries of A^x by a_{kl}^x and similarly for A^y and $A^{\lambda x + y}$, then we have $a_{kj}^{\lambda x + y} = a_{kj}^x = a_{kj}^y$ for $k \in \{1, \dots, n\}$ with $k \neq i$, since only the i th row is different. For the i th row we get $a_{ij}^{\lambda x + y} = \lambda a_{ij}^x + a_{ij}^y$. Altogether we obtain

$$\begin{aligned} \delta(A^{\lambda x + y}) &= \sum_{k=1}^n a_{kj}^{\lambda x + y} c_{kj}^{\lambda x + y} = a_{1j}^{\lambda x + y} c_{1j}^{\lambda x + y} + \dots + a_{ij}^{\lambda x + y} c_{ij}^{\lambda x + y} + \dots + a_{nj}^{\lambda x + y} c_{nj}^{\lambda x + y} \\ &= a_{1j}^x (\lambda c_{1j}^x + c_{1j}^y) + \dots + (\lambda a_{ij}^x + a_{ij}^y) c_{ij}^x + \dots + a_{nj}^x (\lambda c_{nj}^x + c_{nj}^y) \\ &= \lambda a_{1j}^x c_{1j}^x + a_{1j}^y c_{1j}^y + \dots + \lambda a_{ij}^x c_{ij}^x + a_{ij}^y c_{ij}^y + \dots + \lambda a_{nj}^x c_{nj}^x + a_{nj}^y c_{nj}^y \\ &= \lambda a_{1j}^x c_{1j}^x + a_{1j}^y c_{1j}^y + \dots + \lambda a_{ij}^x c_{ij}^x + a_{ij}^y c_{ij}^y + \dots + \lambda a_{nj}^x c_{nj}^x + a_{nj}^y c_{nj}^y \\ &= \lambda (a_{1j}^x c_{1j}^x + \dots + a_{ij}^x c_{ij}^x + \dots + a_{nj}^x c_{nj}^x) + a_{1j}^y c_{1j}^y + \dots + a_{ij}^y c_{ij}^y + \dots + a_{nj}^y c_{nj}^y \\ &= \lambda \delta(A^x) + \delta(A^y) . \end{aligned}$$

This finishes the proof that δ satisfies 5.0.3 1).

Now assume that $A \in M_{n \times n}(F)$ is a matrix, in which the k th and l th row agree. Without loss of generality we may assume $k < l$. Let M_{aj} be the matrix obtained from A by removing the a th row and the j th column and denote the cofactors of A by c_{aj} . If $a \neq k$ and $a \neq l$, then $\det(M_{aj}) = 0$, since the resulting matrix will still have two rows that agree. Therefore $c_{aj} = 0$ for $a \neq k$ and $a \neq l$.

The two matrices M_{kj} and M_{lj} agree up to a permutation of the rows: More precisely, to obtain the matrix M_{lj} from M_{kj} we have to move the $(l-1)$ th row of M_{kj} up to the k th one. This involves $(l-k-1)$ transpositions of rows. Therefore

$$\det(M_{kj}) = (-1)^{l-k-1} \det(M_{lj}) .$$

This yields $c_{kj} = (-1)^{k+j} \det(M_{kj}) = (-1)^{l+j-1} \det(M_{lj}) = -c_{lj}$. Since $a_{kj} = a_{lj}$ we obtain in this case

$$\delta(A) = c_{kj}a_{kj} + c_{lj}a_{lj} = c_{kj}a_{kj} - c_{kj}a_{kj} = 0 .$$

This shows that δ satisfies Def. 5.0.3 2) as well.

For the unit matrix $A = I_n$ we have $a_{kj} = 0$ for all $k \in \{1, \dots, n\}$ with $k \neq j$ and $a_{jj} = 1$. Moreover, $c_{jj} = (-1)^{2j} \det(I_{n-1}) = 1$. Therefore

$$\delta(I_n) = a_{jj}c_{jj} = 1 \cdot 1 = 1 .$$

Thus, δ also has property 3) in Def. 5.0.3. This finishes the proof. \square

Remark 5.4.6. Note that the definition of the cofactors of a matrix $A \in M_{n \times n}(F)$ involve determinants of smaller matrices. Therefore Thm. 5.4.5 indeed provides a recursive way of computing $\det(A)$.

Example 5.4.7. In this example we will see how the cofactor expansion can be used to compute the determinant of a matrix that contains a lot of zeroes. Consider the following matrix

$$A = \begin{pmatrix} 0 & 1 & 2 & 3 \\ 0 & 0 & 1 & 2 \\ 1 & 10 & 13 & 15 \\ 0 & 0 & 3 & 4 \end{pmatrix} .$$

We can compute the determinant of A using the cofactor expansion along the first column. If we denote the cofactors of A by c_{ij} again, we obtain

$$\det(A) = 0 \cdot c_{11} + 0 \cdot c_{21} + 1 \cdot c_{31} + 0 \cdot c_{41} = c_{31} .$$

This means we just have to compute the cofactor c_{31} . Using its definition we obtain

$$c_{31} = (-1)^{3+1} \det \begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 3 & 4 \end{pmatrix} .$$

We can compute this determinant also by cofactor expansion along the first column. This yields

$$\det \begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 2 \\ 0 & 3 & 4 \end{pmatrix} = 1 \cdot \det \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} = 4 - 6 = -2 .$$

Therefore $\det(A) = -2$ as well.

Definition 5.4.8. Let $A \in M_{n \times n}(F)$. The matrix $C = (c_{ij})_{i,j=1}^n$ containing the cofactors c_{ij} of A is called the **cofactor matrix** of A . Its transpose matrix, i.e.

$$\text{adj}(A) = C^t$$

is called the **adjugate matrix** of A .

The matrix $\text{adj}(A)$ can be used to compute the inverse of a matrix A . This is based on the following theorem:

Theorem 5.4.9. *Let F be a field and let $A \in M_{n \times n}(F)$. Then the adjugate matrix $\text{adj}(A)$ satisfies*

$$\text{adj}(A) \cdot A = A \cdot \text{adj}(A) = \det(A) \cdot I_n .$$

Proof. Denote the entries of A by a_{ik} and its cofactors by c_{ik} . Let $B = A \cdot \text{adj}(A)$ with entries b_{ik} . Then by definition of the adjugate matrix and matrix multiplication we get

$$b_{ik} = (A \cdot C^t)_{ik} = \sum_{j=1}^n a_{ij} c_{kj} .$$

For $i = k$ this is exactly the cofactor expansion of the determinant along the i th row from Thm. 5.4.5. Thus, $b_{ii} = \det(A)$. For $i \neq k$ the entry b_{ik} yields the cofactor expansion of the matrix A' , where A' is obtained from A by replacing its k th row by its i th row. But then the resulting matrix has two rows that agree, which implies $\det(A') = 0 = b_{ik}$. Altogether we obtain that B is the matrix with $\det(A)$ along the diagonal and zeroes everywhere else. The proof for $\text{adj}(A) \cdot A$ is very similar and we omit it. \square

Corollary 5.4.10. *Let F be a field and let $A \in M_{n \times n}(F)$. Suppose that $\det(A) \neq 0$. Then the inverse of A is given by*

$$A^{-1} = \frac{1}{\det(A)} \text{adj}(A) .$$

Proof. There is not much left to check here. It is a direct consequence of Thm. 5.4.9 that

$$A \cdot \left(\frac{1}{\det(A)} \text{adj}(A) \right) = \left(\frac{1}{\det(A)} \text{adj}(A) \right) \cdot A = I_n .$$

\square

Example 5.4.11. Let $A \in M_{2 \times 2}(\mathbb{R})$ be an arbitrary 2×2 -matrix over the reals as in Example 5.4.2, i.e.

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

We can read off the cofactor matrix C of A from Example 5.4.2:

$$C = \begin{pmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{pmatrix} = \begin{pmatrix} d & -c \\ -b & a \end{pmatrix}$$

The adjugate matrix is then given by

$$\text{adj}(A) = \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}$$

and if $\det(A) = ad - bc \neq 0$ the inverse of A is given by

$$A^{-1} = \frac{1}{\det(A)} \text{adj}(A) = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix} .$$

Exercise 5.4.12. Compute the adjugate and the inverse of the matrix $A \in M_{3 \times 3}(\mathbb{R})$ given by

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 0 & 1 & 0 \\ 1 & 1 & 2 \end{pmatrix} .$$

As the following theorem, which is known as **Cramer's rule**, shows we can use determinants to solve systems of linear equations.

Theorem 5.4.13. Let F be a field, let $A \in M_{n \times n}(F)$ be a square matrix with $\det(A) \neq 0$ and let $b \in F^n$. For $i \in \{1, \dots, n\}$ define $A_i \in M_{n \times n}(F)$ to be the matrix obtained from A by replacing the i th column with b . Then the unique solution $x \in F^n$ of the system of linear equations given in matrix form by $Ax = b$ is given by $x = (x_1, \dots, x_n)$ with

$$x_i = \frac{\det(A_i)}{\det(A)} .$$

Proof. Since A is invertible, the solution $x \in F^n$ satisfies $x = A^{-1}b$. By Thm. 5.4.9

$$x = A^{-1}b = \frac{1}{\det(A)} \operatorname{adj}(A)b .$$

But the entries of the vector $\operatorname{adj}(A)b$ are given by

$$(\operatorname{adj}(A)b)_i = \sum_{j=1}^n b_j c_{ji} .$$

The right hand side of this equation is the determinant of the matrix obtained from A by replacing the i th column by b , i.e. the right hand side agrees with $\det(A_i)$. \square

6 Diagonalisation, Eigenvalues and Eigenvectors

Let V be a finite-dimensional vector space and let $T: V \rightarrow V$ be a linear transformation. It is a common problem in Linear Algebra to compute powers of T , i.e. the n -fold composition of T with itself, denoted by

$$T^n = \underbrace{T \circ \dots \circ T}_{n \text{ times}} .$$

If $v \in V$ is a vector with the property that $T(v) = \lambda v$ for some scalar $\lambda \in F$, then $T^n(v)$ is very easy to compute:

$$T^n(v) = T^{n-1}(T(v)) = T^{n-1}(\lambda v) = \lambda T^{n-1}(v) = \dots = \lambda^n v .$$

These vectors will play a crucial role in this chapter. Therefore we make the following definition.

Definition 6.0.1. Let V be a vector space over the field F and let $T: V \rightarrow V$ be a linear transformation. A nonzero vector $v \in V$ is called an **eigenvector** of T if there exists a scalar $\lambda \in F$ such that $T(v) = \lambda v$. The scalar λ is called the **eigenvalue** corresponding to the eigenvector v .

Let $A \in M_{n \times n}(F)$. A nonzero vector $v \in F^n$ is called an **eigenvector** of A if v is an eigenvector of L_A , that is, if $Av = \lambda v$ for some scalar $\lambda \in F$. The scalar λ is called the **eigenvalue** of the matrix A corresponding to the eigenvector v .

We have already seen in Example 3.0.25 that a basis consisting of such vectors, i.e. a basis $\beta = \{v_1, \dots, v_n\}$ with the property that $T(v_i) = \lambda_i v_i$ for scalars $\lambda_i \in F$ can be used to compute $T^n(v)$ for arbitrary vectors $v \in V$. In particular, the matrix representation of T with respect to β is a diagonal matrix. This motivates the following definition:

Definition 6.0.2. A linear transformation $T: V \rightarrow V$ on a finite-dimensional vector space V is called **diagonalizable** if there is an ordered basis β for V such that $[T]_\beta^\beta$ is a diagonal matrix. A matrix $A \in M_{n \times n}(F)$ is called **diagonalizable** if L_A is diagonalizable.

Let $T: V \rightarrow V$ be a diagonalizable linear transformation and let $\beta = \{v_1, \dots, v_n\}$ be a basis of V such that $[T]_\beta^\beta$ is a diagonal matrix. Then each vector v_i is an eigenvector of T . To see why this is true let $D = [T]_\beta^\beta$ with entries $D_{ii} = \lambda_i$ and $D_{ij} = 0$ for $i \neq j$. Using the definition of the matrix representation we have

$$T(v_j) = \sum_j D_{ij} v_i = D_{jj} v_j = \lambda_j v_j .$$

Therefore the vector $v_i \in \beta$ is an eigenvector of T corresponding to the eigenvalue λ_i . In particular, T is diagonalizable if and only if there exists a basis for V that consists entirely of eigenvectors of T .

We will illustrate the importance of these concepts by means of the following examples that will accompany us all throughout this section.

Example 6.0.3. The Fibonacci sequence $(f_n)_{n \in \mathbb{N}_0}$ is defined as follows: We have $f_0 = 1$, $f_1 = 1$ and $f_n = f_{n-2} + f_{n-1}$, i.e. each element is the sum of the two previous elements. The first numbers obtained in this way are 1, 1, 2, 3, 5, ... **Is there a way to compute f_n directly without having to compute f_0, \dots, f_{n-1} ?**

Let us rephrase this question in terms of a Linear Algebra problem: Let $A \in M_{2 \times 2}(\mathbb{R})$ be the following matrix

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

Let $e_1 \in \mathbb{R}^2$ be the first vector of the standard basis for \mathbb{R}^2 . Then

$$Ae_1 = \begin{pmatrix} 1 \\ 1 \end{pmatrix} , \quad A^2e_1 = \begin{pmatrix} 2 \\ 1 \end{pmatrix} , \quad A^3e_1 = \begin{pmatrix} 3 \\ 2 \end{pmatrix} , \quad A^4e_1 = \begin{pmatrix} 5 \\ 3 \end{pmatrix} .$$

The first coordinate of the vector $A^n e_1$ seems to be the n th Fibonacci number f_n , while the second appears to be f_{n-1} . Indeed, this is true for $n = 0$ (with $f_{-1} = 0$) and $n = 1$ and

$$A \begin{pmatrix} f_{n-1} \\ f_{n-2} \end{pmatrix} = \begin{pmatrix} f_{n-1} + f_{n-2} \\ f_{n-1} \end{pmatrix} = \begin{pmatrix} f_n \\ f_{n-1} \end{pmatrix}$$

implies the general statement by induction. Therefore the problem of finding a formula for f_n boils down to computing A^n for arbitrary $n \in \mathbb{N}_0$. Thus, if we could find a basis $\beta = \{v_1, v_2\}$ for \mathbb{R}^2 consisting of eigenvectors of A corresponding to the eigenvalues λ_1 and λ_2 , respectively, then we could write e_1 as a linear combination of v_1 and v_2 , say $e_1 = \alpha_1 v_1 + \alpha_2 v_2$, and we would get

$$A^n e_1 = \alpha_1 A^n v_1 + \alpha_2 A^n v_2 = \alpha_1 \lambda_1^n v_1 + \alpha_2 \lambda_2^n v_2 = \begin{pmatrix} f_n \\ f_{n-1} \end{pmatrix}. \quad (15)$$

Example 6.0.4. In the early days of the internet search engines used the PageRank algorithm to sort the results of a search according to their importance. Since this algorithm is based on finding an eigenvector of a matrix associated to a network, we will discuss it a little in this example.

From the perspective of a search engine the internet consists of a collection of pages and links between them. We can visualise it by drawing the pages as circles and a link from page A to page B by a little arrow pointing from the circle of A to the one of B . An example of such a network with five pages and ten links is shown in Figure 3. (In mathematical terms, such a figure is called a *directed graph*, but this is not important here.)




Figure 3: Our example search results.

Suppose our internet search returned the pages A, B, C, D, E as results and we would like to sort these pages according to how important they are. Finding the right measure for this is a tricky question. The PageRank algorithm approaches this problem as follows: If we start on some page of our search results and randomly click from page to page, then the pages that we will visit more often, must be more important than the rest.

The probability on which page we will end up can be encoded in a matrix in the following way: Observe that the only link on page A will take us to page B . Therefore if we start on page A , the probability that we end up on page B after one click is 1. There are two links on page B : One will take us back to page A , the other one to page C .

Thus, if we start on page B , the probability to end up on page A after one click is $\frac{1}{2}$ and the probability to end up on C is also $\frac{1}{2}$. We write these probabilities into a matrix P as follows: Identify the numbers $1, \dots, 5$ with the pages A, \dots, E . The entry in column i and row j corresponds to the probability to end up on page j after one click if we start from page i . In our example, the matrix P looks as follows:

$$P = \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{3} & 1 & 0 \\ 1 & 0 & \frac{1}{3} & 0 & \frac{1}{3} \\ 0 & \frac{1}{2} & 0 & 0 & \frac{1}{3} \\ 0 & 0 & 0 & 0 & \frac{1}{3} \\ 0 & 0 & \frac{1}{3} & 0 & 0 \end{pmatrix}$$

Let $v \in \mathbb{R}^5$ be a vector with nonnegative entries that sum up to 1. We can think of such a vector as a probability distribution for our starting page. The probability distribution after one click is then given by Pv . For example, if we take the vector $v = e_2$ corresponding to the probability distribution that is concentrated on page B , then Pe_2 has probability $\frac{1}{2}$ to be on page A and probability $\frac{1}{2}$ to be on page C . Likewise, the probability distribution after n clicks with starting distribution v is given by $P^n v$. The matrix P is also called a **Markov transition matrix** and the underlying mechanism is called a **Markov chain**.

The PageRank algorithm looks for *stationary* probability distributions, i.e. those distributions that do not change after one click. It will sort the results according to their values in the vector v from high to low. By our above observations this means we are looking for a vector $v \in \mathbb{R}^5$ with nonnegative entries summing up to 1 that satisfies

$$Pv = v .$$

In particular, we are interested in eigenvectors of P with respect to the eigenvalue 1.

Example 6.0.5. Let $C^\infty(\mathbb{R})$ denote the set of all functions $f: \mathbb{R} \rightarrow \mathbb{R}$ having derivatives of all orders. This is a vector space over \mathbb{R} with respect to the usual addition and scalar multiplication of functions. Let $T: C^\infty(\mathbb{R}) \rightarrow C^\infty(\mathbb{R})$ be the linear transformation given by $T(f) = f'$, i.e. the derivative of f . Suppose that $f \in C^\infty(\mathbb{R})$ is an eigenvector of T with corresponding eigenvalue $\lambda \in \mathbb{R}$. Then

$$f' = T(f) = \lambda f .$$

This is a first-order differential equation whose solutions are of the form $f(t) = ce^{\lambda t}$ for some constant $c \in \mathbb{R}$. Consequently, every real number $\lambda \in \mathbb{R}$ is an eigenvalue of T , and λ corresponds to eigenvectors of the form $f(t) = ce^{\lambda t}$ for $c \neq 0$. Observe that for $\lambda = 0$ the eigenvectors are nonzero constant functions.

If $\lambda \in F$ is an eigenvalue of a matrix $A \in M_{n \times n}(F)$ with corresponding eigenvector $v \in F^n$, then v is a nonzero vector in the nullspace of L_B with $B = A - \lambda I_n$. This provides us with a method of computing the eigenvalues of a matrix.

Theorem 6.0.6. Let $A \in M_{n \times n}(F)$. Then a scalar $\lambda \in F$ is an eigenvalue of A if and only if $\det(A - \lambda I_n) = 0$.

Proof. A scalar $\lambda \in F$ is an eigenvalue of A if and only if there exists a nonzero vector $v \in F^n$ such that $Av = \lambda v$, that is, $(A - \lambda I_n)v = 0$. Since v is nonzero this condition is equivalent to the statement that $\lambda \in F$ is an eigenvalue of A if and only if L_B with $B = A - \lambda I_n$ is not injective. Since the domain and codomain of $L_B: F^n \rightarrow F^n$ have the same dimension, this is the same as saying that L_B and hence $B = A - \lambda I_n$ is not invertible, which is the case if and only if $\det(A - \lambda I_n) = 0$. \square

Definition 6.0.7. Let $A \in M_{n \times n}(F)$. The polynomial $p_A(t) = \det(A - tI_n)$ is called the **characteristic polynomial** of the matrix A .

Theorem 6.0.6 states that the eigenvalues of a matrix A are the zeroes of its characteristic polynomial. This allows us to compute the eigenvalues before we even have the eigenvectors. As we will see in the next example computing the eigenvectors then reduces to solving a system of linear equations.

Example 6.0.8. Let us consider the matrix $A \in M_{2 \times 2}(\mathbb{R})$ from Example 6.0.3, i.e.

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

The characteristic polynomial $p_A(t)$ of this matrix is given by

$$p_A(t) = \det \left(\begin{pmatrix} 1-t & 1 \\ 1 & -t \end{pmatrix} \right) = (1-t)(-t) - 1 = t^2 - t - 1.$$

The zeroes $\lambda_1, \lambda_2 \in \mathbb{R}$ of the quadratic polynomial $p_A(t)$ are given by

$$\lambda_1 = \frac{1 + \sqrt{5}}{2},$$

$$\lambda_2 = \frac{1 - \sqrt{5}}{2}.$$

Note that $\lambda_1 + \lambda_2 = 1$ and $\lambda_1 \lambda_2 = -1$. To find an eigenvector corresponding to the eigenvalue λ_1 we have to solve the following vector equation

$$\begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \lambda_1 \begin{pmatrix} x \\ y \end{pmatrix},$$

which yields the following system of equations

$$x + y = \lambda_1 x$$

$$x = \lambda_1 y$$

The first equation is equivalent to $y = (\lambda_1 - 1)x = -\lambda_2 x$. This agrees with the second equation, because $\lambda_1 \lambda_2 = -1$. Therefore the solution set of the above equations is given by the linear subspace

$$E_{\lambda_1} = \left\{ a \begin{pmatrix} \lambda_1 \\ 1 \end{pmatrix} \mid a \in \mathbb{R} \right\} \subset \mathbb{R}^2 .$$

Any vector in E_{λ_1} is an eigenvector of A corresponding to the eigenvalue λ_1 . Likewise, any vector in the subspace E_{λ_2} given by

$$E_{\lambda_2} = \left\{ a \begin{pmatrix} \lambda_2 \\ 1 \end{pmatrix} \mid a \in \mathbb{R} \right\}$$

is an eigenvector of A corresponding to the eigenvalue λ_2 . We pick one vector of each subspace, e.g.

$$v_1 = \begin{pmatrix} \lambda_1 \\ 1 \end{pmatrix} , \quad v_2 = \begin{pmatrix} \lambda_2 \\ 1 \end{pmatrix} .$$

We will see later that $\beta = \{v_1, v_2\}$ is indeed a basis for \mathbb{R}^2 and take that for granted for now. Observe that e_1 can be written as a linear combination of v_1 and v_2 as follows

$$e_1 = \frac{1}{\sqrt{5}}(v_1 - v_2) = \frac{1}{\sqrt{5}}v_1 + \left(-\frac{1}{\sqrt{5}}\right)v_2 ,$$

that is $e_1 = \alpha_1 v_1 + \alpha_2 v_2$ with $\alpha_1 = (\sqrt{5})^{-1}$ and $\alpha_2 = -(\sqrt{5})^{-1}$. Comparing this result with equation (15) we obtain the formula for $A^n e_1$ we were looking for

$$A^n e_1 = \alpha_1 \lambda_1^n v_1 + \alpha_2 \lambda_2^n v_2 = \frac{1}{\sqrt{5}} \begin{pmatrix} \lambda_1^{n+1} - \lambda_2^{n+1} \\ \lambda_1^n - \lambda_2^n \end{pmatrix}$$

From the first component of this vector we obtain the formula for f_n we were looking for

$$f_n = \frac{1}{\sqrt{5}} \left(\left(\frac{1 + \sqrt{5}}{2} \right)^{n+1} - \left(\frac{1 - \sqrt{5}}{2} \right)^{n+1} \right)$$

Considering the square roots appearing in this expression it is a little surprising that each f_n obtained in this way is in fact an integer. This is nevertheless true and can easily be verified for small values of n using a computer.

We have already seen above that the determinant also makes sense as a map that associates a scalar in F directly to a linear transformation $T: V \rightarrow V$ on a finite-dimensional vector space over F . Therefore the definition of the characteristic polynomial can be extended as well:

Definition 6.0.9. Let V be a finite-dimensional vector space over a field F and let $T: V \rightarrow V$ be a linear transformation. We define the **characteristic polynomial** $p_T(t)$ of T as follows

$$p_T(t) = \det(T - tI_V) ,$$

where $I_V: V \rightarrow V$ is the identity transformation.

The next theorem shows that eigenvectors corresponding to distinct eigenvalues are automatically linearly independent. In particular, if we can find enough of them, then they will automatically form a basis.

Theorem 6.0.10. *Let $T: V \rightarrow V$ be a linear transformation on a finite-dimensional vector space V over a field F and let $\lambda_1, \dots, \lambda_k$ be distinct eigenvalues of T . If v_1, \dots, v_k are eigenvectors of T such that λ_i corresponds to v_i for $i \in \{1, \dots, k\}$, then $\{v_1, \dots, v_k\}$ is linearly independent.*

Proof. We will prove this using induction. For $k = 1$ we only have one vector v_1 , which is nonzero, since it is an eigenvector of T . Therefore $\{v_1\}$ is linearly independent. Now suppose that the set $\{v_1, \dots, v_{k-1}\}$ is linearly independent. Suppose that $a_1, \dots, a_k \in F$ are scalars with the property that

$$a_1v_1 + a_2v_2 + \dots + a_kv_k = 0. \quad (16)$$

If we apply the linear transformation $(T - \lambda_k I_V)$ to both sides of the equation and use the fact that $(T - \lambda_k I_V)v_k = 0$ we obtain

$$a_1(\lambda_1 - \lambda_k)v_1 + a_2(\lambda_2 - \lambda_k)v_2 + \dots + a_{k-1}(\lambda_{k-1} - \lambda_k)v_{k-1} = 0.$$

By our induction hypothesis the set $\{v_1, \dots, v_{k-1}\}$ is linearly independent. Therefore the above computation implies $a_i(\lambda_i - \lambda_k) = 0$ for all $i \in \{1, \dots, k-1\}$. But since the eigenvalues are distinct, we have $\lambda_i - \lambda_k \neq 0$ and therefore $a_i = 0$ for all $i \in \{1, \dots, k-1\}$. Thus, the original equation (16) boils down to $a_kv_k = 0$, which implies that $a_k = 0$, since $v_k \neq 0$, because it is an eigenvector. Therefore $a_1 = \dots = a_k = 0$ and the set $\{v_1, \dots, v_k\}$ is indeed linearly independent. \square

Corollary 6.0.11. *Let V be a vector space over the field F with $\dim(V) = n$ and let $T: V \rightarrow V$ be a linear transformation. If T has n distinct eigenvalues, then T is diagonalizable.*

Proof. Suppose that T has n distinct eigenvalues $\lambda_1, \dots, \lambda_n$. Choose an eigenvector v_i corresponding to λ_i . By Thm. 6.0.10, the set $\{v_1, \dots, v_n\}$ is linearly independent, and since $\dim(V) = n$, this set is basis for V . Thus, we have found a basis of eigenvectors and T is diagonalizable. \square

It is a consequence of the last corollary that a matrix $A \in M_{n \times n}(F)$ with n distinct eigenvalues is diagonalizable. Let $T: V \rightarrow V$ be a linear transformation on a finite-dimensional vector space V over a field F . If λ is an eigenvalue of T , then the eigenvectors with corresponding eigenvalue λ are never unique. They lie in a linear subspace of V that we now define.

Definition 6.0.12. Let $T: V \rightarrow V$ be a linear transformation on a vector space V over a field F and let $\lambda \in F$ be an eigenvalue of T . Define

$$E_\lambda = \{v \in V \mid T(v) = \lambda v\} = \ker(T - \lambda I_V).$$

The linear subspace $E_\lambda \subset V$ is called the **eigenspace** of T corresponding to the eigenvalue λ . Analogously, we define the **eigenspace** of a matrix $A \in M_{n \times n}(F)$ to be the eigenspace of L_A .

Example 6.0.13. The two subspaces E_{λ_1} and E_{λ_2} in Example 6.0.8 are the eigenspaces of the matrix A in that example corresponding to λ_1 and λ_2 , respectively. Notice that in that example, both eigenspaces are one-dimensional. This does not always have to be the case! In this example, we will see a 3×3 -matrix with a two-dimensional eigenspace. Let $A \in M_{3 \times 3}(\mathbb{R})$ be the following matrix:

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

To compute the eigenvalues we look at the characteristic polynomial given by

$$\begin{aligned} p_A(t) &= \det \left(\begin{pmatrix} -t & 0 & -2 \\ 1 & 2-t & 1 \\ 1 & 0 & 3-t \end{pmatrix} \right) = (2-t)((-t)(3-t) + 2) \\ &= -(t-2)^2(t-1) \end{aligned}$$

where we used the cofactor expansion along the second column to compute the determinant. Thus, the eigenvalues of A are 2 and 1. The eigenspace of A corresponding to the eigenvalue 2 is by definition the nullspace of $L_B: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ with

$$B = A - 2I_3 = \begin{pmatrix} -2 & 0 & -2 \\ 1 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix} .$$

Thus, the eigenspace E_2 of A is given by

$$E_2 = \{(x, y, z) \in \mathbb{R}^3 \mid z = -x\} .$$

Any nonzero vector $v \in E_2$ is an eigenvector of A with corresponding eigenvalue 2. A basis for this vector space is given by $\{v_1, v_2\}$ with $v_1 = (1, 0, -1)$ and $v_2 = (0, 1, 0)$.

The eigenspace of A corresponding to the eigenvalue 1 is the nullspace of $L_C: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ with

$$C = \begin{pmatrix} -1 & 0 & -2 \\ 1 & 1 & 1 \\ 1 & 0 & 2 \end{pmatrix} .$$

A short computation of the nullspace shows that the eigenspace E_1 is given by

$$E_1 = \{a(-2, 1, 1) \mid a \in \mathbb{R}\} .$$

and a basis of E_1 is given by $\{v_3\}$ with $v_3 = (-2, 1, 1)$. In particular, $\beta = \{v_1, v_2, v_3\}$ is a basis of \mathbb{R}^3 which consists of eigenvectors of A . Therefore A is diagonalizable with

$$[L_A]_\beta^\beta = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix} .$$

Example 6.0.14. Not all matrices are diagonalizable! In this example, we will discuss two reasons why this can happen. Let

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \in M_{2 \times 2}(\mathbb{R})$$

The characteristic polynomial of this matrix is

$$p_A(t) = \det \left(\begin{pmatrix} 1-t & 1 \\ 0 & 1-t \end{pmatrix} \right) = (1-t)^2 .$$

Therefore the only eigenvalue of A is 1. If the matrix were diagonalizable, then there would be a basis β , such that

$$[L_A]_{\beta}^{\beta} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

because 1 is the only eigenvalue. But this would imply that L_A is the identity transformation, which is clearly not the case. What went wrong? The eigenspace E_1 of A for the eigenvalue 1 is the nullspace of L_C with

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

Therefore E_1 is the linear subspace given by

$$E_1 = \{ a(1, 0) \in \mathbb{R}^2 \mid a \in \mathbb{R} \} .$$

This is just a one-dimensional subspace of \mathbb{R}^2 . In particular, there can not be a basis of eigenvectors. Note that the factor $(1-t)$ appears in $(1-t)^2$ with power 2, whereas the dimension of E_1 is just one. We will see later that this discrepancy reveals that A is not diagonalizable.

Now let $B \in M_{2 \times 2}(\mathbb{R})$ be the following matrix:

$$B = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} .$$

with characteristic polynomial $p_B(t) = t^2 + 1$. Note that this polynomial has no zeroes on the real line! Therefore there are no eigenvalues in \mathbb{R} .

If $p \in P_n(F)$ is a polynomial, then a value $a \in F$ is a root of the polynomial p if and only if p can be factored as $p(t) = (t-a)q(t)$ for some polynomial $q \in P_{n-1}(F)$. This motivates the following definition:

Definition 6.0.15. A polynomial $p \in P_n(F)$ **splits over** F if there are $c, a_1, \dots, a_n \in F$ (not necessarily distinct) such that

$$p(t) = c(t-a_1)(t-a_2)\dots(t-a_n) .$$

Theorem 6.0.16. *The characteristic polynomial of any diagonalizable linear transformation $T: V \rightarrow V$ on a finite-dimensional vector space V splits.*

Proof. Let β be an ordered basis of V such that $[T]_{\beta}^{\beta} = D$, where D is a diagonal matrix, say

$$D = \begin{pmatrix} \lambda_1 & 0 & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{pmatrix}.$$

The matrix representation of $T - tI_V$ with respect to the basis β is then given by

$$\begin{pmatrix} \lambda_1 - t & 0 & \dots & 0 \\ 0 & \lambda_2 - t & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n - t \end{pmatrix}$$

and therefore $p_T(t) = \det(T - tI_V) = (\lambda_1 - t) \dots (\lambda_n - t) = (-1)^n (t - \lambda_1) \dots (t - \lambda_n)$. \square

As we have seen in Example 6.0.14 the converse of the above statement is false. Even if the characteristic polynomial splits, T need not be diagonalizable. However, the largest power with which the factor $(t - \lambda)$ appears in the characteristic polynomial will give us an upper bound on the dimension of E_{λ} as we will see shortly. First, we make the following definition:

Definition 6.0.17. Let $\lambda \in F$ be an eigenvalue of a linear transformation $T: V \rightarrow V$ or a matrix $A \in M_{n \times n}(F)$ with characteristic polynomial $p(t) \in P_n(F)$. The **(algebraic) multiplicity** of λ is the largest positive integer k for which $(t - \lambda)^k$ is a factor of p .

Theorem 6.0.18. *Let $T: V \rightarrow V$ be a linear transformation on a finite-dimensional vector space V over a field F and let $\lambda \in F$ be an eigenvalue of T having multiplicity m . Then*

$$1 \leq \dim(E_{\lambda}) \leq m.$$

Proof. Let $p = \dim(E_{\lambda})$, $n = \dim(V)$ and choose an ordered basis $\beta_{\lambda} = \{v_1, \dots, v_p\}$ for $E_{\lambda} \subseteq V$. Extend β_{λ} to an ordered basis $\beta = \{v_1, \dots, v_p, v_{p+1}, \dots, v_n\}$ for V . Let $A = [T]_{\beta}^{\beta}$. Observe that v_i is an eigenvector of T for each $i \in \{1, \dots, p\}$ with corresponding eigenvalue λ , and therefore

$$A = \begin{pmatrix} \lambda I_p & B \\ 0 & C \end{pmatrix}$$

with $(n - p) \times (n - p)$ -matrices B and C . Thus, the characteristic polynomial of T is

$$\begin{aligned} p_T(t) &= \det(A - tI_n) = \det \left(\begin{pmatrix} (\lambda - t)I_p & B \\ 0 & C - tI_{n-p} \end{pmatrix} \right) \\ &= (\lambda - t)^p \det(C - tI_{n-p}) \end{aligned}$$

where we used the cofactor expansion with respect to the first columns in the last equality. Thus, $(t - \lambda)^p = (-1)^p(\lambda - t)^p$ is a factor of $p_T(t)$ and hence the multiplicity m of λ is at least p , i.e. $m \geq p = \dim(E_\lambda)$. \square

Theorem 6.0.19. *Let V be a finite-dimensional vector space over a field F and let $T: V \rightarrow V$ be a linear transformation. Suppose that the characteristic polynomial $p_T(t)$ of T splits. Let $\lambda_1, \dots, \lambda_k \in F$ be the distinct eigenvalues of T . Then*

- a) *T is diagonalizable if and only if the multiplicity of λ_i is equal to $\dim(E_{\lambda_i})$ for all $i \in \{1, \dots, k\}$.*
- b) *If T is diagonalizable and β_i is an ordered basis for E_{λ_i} for each $i \in \{1, \dots, k\}$, then $\beta = \beta_1 \cup \dots \cup \beta_k$ is an ordered basis for V consisting of eigenvectors of T .*

Proof. For each $i \in \{1, \dots, k\}$ denote by m_i the multiplicity of λ_i , let $d_i = \dim(E_{\lambda_i})$ and let $n = \dim(V)$.

First suppose that T is diagonalizable and let β be a basis for V consisting of eigenvectors of T . For each $i \in \{1, \dots, k\}$ let $\beta_i = \beta \cap E_{\lambda_i}$ be the set of vectors in β that are eigenvectors with corresponding eigenvalue λ_i . Let $n_i \in \mathbb{N}$ be the number of vectors in β_i . Since β_i is a linearly independent subset of E_{λ_i} , which is a subspace of dimension d_i , we obtain $n_i \leq d_i$. Moreover, $d_i \leq m_i$ by Theorem 6.0.18. Note that the degree of the characteristic polynomial is n . Therefore the multiplicities m_i sum up to n . The n_i also sum up to n , since β contains n elements in total. Thus,

$$n = \sum_{i=1}^k n_i \leq \sum_{i=1}^k d_i \leq \sum_{i=1}^k m_i = n .$$

Therefore the above inequalities are actually equalities and

$$\sum_{i=1}^k (m_i - d_i) = 0 .$$

But $m_i - d_i \geq 0$ as we already observed above. This implies $m_i = d_i$ for all i .

Now suppose that $m_i = d_i$ for all $i \in \{1, \dots, k\}$. We will show that T is diagonalizable by constructing a basis for V that consists of eigenvectors as described in b). This will then also prove b). For each i let β_i be an ordered basis for the eigenspace E_{λ_i} . Let $\beta = \beta_1 \cup \dots \cup \beta_k$. Since $d_i = m_i$ for all i , the set β contains

$$\sum_{i=1}^k d_i = \sum_{i=1}^k m_i = n$$

elements. Thus, to see that β is a basis for V it suffices to prove that β is linearly independent. For all $i \in \{1, \dots, k\}$ let $\beta_i = \{v_1^i, \dots, v_{d_i}^i\}$ and let $a_j^i \in F$ be scalars such that

$$a_1^1 v_1^1 + \dots + a_{d_1}^1 v_{d_1}^1 + a_1^2 v_1^2 + \dots + a_{d_2}^2 v_{d_2}^2 + \dots + a_1^k v_1^k + \dots + a_{d_k}^k v_{d_k}^k = 0 . \quad (17)$$

For each $i \in \{1, \dots, k\}$ let $y_i = a_1^i v_1^i + \dots + a_{d_i}^i v_{d_i}^i$ and observe that each y_i is in E_{λ_i} , since it is a linear combination of basis vectors for that space. In particular, y_i is either an eigenvector of T with corresponding eigenvalue λ_i or it is zero. From (17) we obtain that

$$y_1 + \dots + y_k = 0 .$$

Suppose that at least two of the summands y_i are nonzero. Then we would have a linear combination of eigenvectors with corresponding distinct eigenvalues that sum up to zero. This contradicts the linear independence of such vectors that we proved in Theorem 6.0.10. Thus, all vectors y_i have to be zero. But if $y_i = 0$, then

$$a_1^i v_1^i + \dots + a_{d_i}^i v_{d_i}^i = 0$$

which implies $a_j^i = 0$, since the vectors $v_1^i, \dots, v_{d_i}^i$ form a basis for E_{λ_i} . This proves that β is linearly independent. Therefore β is an ordered basis for V consisting of eigenvectors of T and we conclude that T is diagonalizable. \square

6.1 Diagonalization: A Step by Step Instruction Manual

Let V be a finite-dimensional vector space with $n = \dim(V)$ over the field F and let $T: V \rightarrow V$ be a linear transformation. In this section we will give an instruction manual on how to check whether T is diagonalizable or not, and how to find a basis consisting of eigenvectors.

(Step 1) Find the characteristic polynomial of the linear transformation T (see Def. 6.0.9 or Def. 6.0.7, see also Sec. 5.3). This is a polynomial of degree n given by

$$p_T(t) = \det(T - tI_V) .$$

(Step 2) Calculate the roots $\lambda_1, \dots, \lambda_k \in F$ of the characteristic polynomial and check if the polynomial splits (see Def. 6.0.15). If $F = \mathbb{C}$, then any polynomial will split. If $F = \mathbb{R}$, then $t^2 + 1$ is an example of a polynomial that does not split. If $p_T(t)$ does not split, the linear transformation T is not diagonalizable (see Thm. 6.0.16). In any case the values $\lambda_1, \dots, \lambda_k$ are the eigenvalues of the linear transformation. If there are n distinct eigenvalues, then T is diagonalizable (Cor. 6.0.11) and (Step 4), (Step 5) provide a basis that consists of eigenvectors. Otherwise, proceed with (Step 3).

(Step 3) Find the multiplicity m_i of λ_i for each $i \in \{1, \dots, k\}$ (see Def. 6.0.17). For example, if $p_T(t) = (t - 2)(t - 3)^2$, then the multiplicity of $\lambda_1 = 2$ is $m_1 = 1$ and the multiplicity of $\lambda_2 = 3$ is $m_2 = 2$.

(Step 4) Find a basis β_i for the eigenspace $E_{\lambda_i} = \ker(T - \lambda_i I_V)$ for each $i \in \{1, \dots, k\}$. In case T is the left multiplication transformation with respect to a matrix $A \in M_{n \times n}(F)$ this will lead to a homogeneous system of linear equations. The solution set will be a subspace that is at least one-dimensional. The basis β_i will contain at most m_i elements (see Thm. 6.0.18).

(Step 5) Check if the dimension of E_{λ_i} , i.e. the number of elements in β_i , agrees with the multiplicity m_i . If this is true for all $i \in \{1, \dots, k\}$, then the linear transformation is diagonalizable and $\beta = \beta_1 \cup \dots \cup \beta_k$ is a basis for V that consists of eigenvectors (see Thm. 6.0.19). If there is an i for which this is not true, then T is not diagonalizable.

Exercise 6.1.1. Let $T: P_2(\mathbb{R}) \rightarrow P_2(\mathbb{R})$ be the linear transformation given as follows

$$T(ax^2 + bx + c) = cx^2 + bx + a .$$

- Determine the characteristic polynomial $p_T(t)$ of T and the eigenvalues of T .
- Find a basis β for the vector space $P_2(\mathbb{R})$ such that $[T]_{\beta}^{\beta}$ is a diagonal matrix.

Exercise 6.1.2. Let $T: M_{2 \times 2}(\mathbb{R}) \rightarrow M_{2 \times 2}(\mathbb{R})$ be the linear transformation given by transposition, i.e.

$$T(A) = A^t .$$

- Determine the characteristic polynomial $p_T(t)$ of T and the eigenvalues of T .
- Is T diagonalizable? If yes, find a basis β for the vector space $M_{2 \times 2}(\mathbb{R})$ such that $[T]_{\beta}^{\beta}$ is a diagonal matrix.

Exercise 6.1.3. Let V be a finite-dimensional vector space over a field F . Let $T: V \rightarrow V$ be an isomorphism.

- Let $\lambda \in F$ be an eigenvalue of T . Show that $\lambda \neq 0$ and $\lambda^{-1} \in F$ is an eigenvalue of the linear transformation T^{-1} .
- Prove that the eigenspace E_{λ} of T corresponding to λ is the same as the eigenspace of T^{-1} corresponding to λ^{-1} .
- Prove that if T is diagonalizable, then T^{-1} is diagonalizable.

7 The Jordan Normal Form

In this section we will discuss linear transformations that are not necessarily diagonalizable. In this case we can still ask the question: What is the best possible choice of basis, such that the corresponding matrix of the linear transformation becomes “as diagonal as possible”? We will see later what this best possible form will look like. We know from Thm. 6.0.16 that for a diagonalizable linear transformation $T: V \rightarrow V$ the characteristic polynomial splits. This depends of course on the field we are using: The polynomial $x^2 + 1 = (x + i)(x - i)$ splits over the complex numbers \mathbb{C} , but not over \mathbb{R} . To avoid these difficulties we will only consider vector spaces over \mathbb{C} in this section.

Let V be a vector space over \mathbb{C} and let $T: V \rightarrow V$ be a linear transformation. Suppose that T has only two eigenvalues $\lambda_1, \lambda_2 \in \mathbb{C}$, which are distinct. Note that if $v \in V$ satisfies $T(v) = \lambda_1 v$ and $T(v) = \lambda_2 v$, then we must have $v = 0$, since

$$(\lambda_1 - \lambda_2)v = T(v) - T(v) = 0$$

and $\lambda_1 - \lambda_2 \neq 0$. In particular, we must have $E_{\lambda_1} \cap E_{\lambda_2} = \{0\}$. This situation can be described using the following terminology:

Definition 7.0.1. Let V be a finite-dimensional vector space (over any field). Suppose that V_1, \dots, V_k are subspaces with the properties

- a) $V_i \cap \left(\sum_{i \neq j} V_j \right) = \{0\}$ for any $i \in \{1, \dots, k\}$,
- b) $V = V_1 + V_2 + \dots + V_k$.

In this situation we say that V is the **direct sum** of the vector spaces V_1, \dots, V_k and denote this in the following way:

$$V = V_1 \oplus V_2 \oplus \dots \oplus V_k .$$

Lemma 7.0.2. Suppose that V is the direct sum of the vector spaces V_1, \dots, V_k . Let $d_j = \dim(V_j)$ and suppose that $\beta_j = \{v_{j,1}, \dots, v_{j,d_j}\}$ is a basis for V_j . Then

$$\beta = \beta_1 \cup \dots \cup \beta_k = \bigcup_{j=1}^k \beta_j$$

is a basis for V .

Proof. Since $V = V_1 + \dots + V_k = \text{span}(V_1 \cup \dots \cup V_k)$ any vector $v \in V$ can be written as

$$v = v_1 + \dots + v_k$$

with $v_j \in V_j$ for all $j \in \{1, \dots, k\}$. But for each j the vector v_j is a linear combination of the vectors $v_{j,1}, \dots, v_{j,d_j}$. Therefore $\text{span}(\beta) = V$.

To see that β is linearly independent let $a_{j,i} \in F$ be scalars with the property that

$$\begin{aligned} & a_{1,1}v_{1,1} + \dots + a_{1,d_1}v_{1,d_1} \\ & + a_{2,1}v_{2,1} + \dots + a_{2,d_2}v_{2,d_2} \\ & + \dots \\ & + a_{k,1}v_{k,1} + \dots + a_{k,d_k}v_{k,d_k} = 0 . \end{aligned}$$

Let $v_j = a_{j,1}v_{j,1} + \dots + a_{j,d_j}v_{j,d_j}$. Then the above line boils down to $v_1 + \dots + v_k = 0$. If we fix $j \in \{1, \dots, k\}$, then we can rewrite this as

$$v_1 + \dots + v_{j-1} + v_{j+1} + \dots + v_k = -v_j .$$

The left hand side is in $V_1 + \dots + V_{j-1} + V_{j+1} + \dots + V_n$, whereas the right hand side is in V_j . Therefore both sides are contained in the intersection of the two spaces. But by Def. 7.0.1 a) this intersection just contains the zero vector, hence

$$v_j = a_{j,1}v_{j,1} + \dots + a_{j,d_j}v_{j,d_j} = 0$$

for all $j \in \{1, \dots, k\}$. This implies that all coefficients $a_{j,i}$ have to vanish, because β_j is a basis. \square

Let $T: V \rightarrow V$ be a diagonalisable linear transformation and let $E_{\lambda_1}, \dots, E_{\lambda_k}$ be the eigenspaces corresponding to the distinct eigenvalues $\lambda_1, \dots, \lambda_k$ of T . Then it is a consequence of Thm. 6.0.19 b) that

$$V = E_{\lambda_1} \oplus \dots \oplus E_{\lambda_k} .$$

Exercise 7.0.3. Let V be a vector space and let V_1, \dots, V_k be subspaces of V with the property that $V = V_1 \oplus \dots \oplus V_k$. Show that for any $v \in V$ there are vectors $v_i \in V_i$ for all $i \in \{1, \dots, k\}$ such that

$$v = v_1 + \dots + v_k .$$

Moreover, prove that this decomposition is unique, i.e. show that if $v'_i \in V_i$ for $i \in \{1, \dots, k\}$ is another set of vectors with $v = v'_1 + \dots + v'_k$, then $v'_i = v_i$.

7.1 Generalised eigenspaces

Suppose that V is a finite-dimensional vector space over \mathbb{C} and that $T: V \rightarrow V$ is a linear transformation. For an eigenvalue $\lambda \in \mathbb{C}$ of T let $m(\lambda)$ be the multiplicity of λ . We know from Thm. 6.0.18 that $\dim(E_\lambda) \leq m(\lambda)$ and that T is diagonalisable if and only if $\dim(E_\mu) = m(\mu)$ for all eigenvalues μ (Thm. 6.0.19 a)). Thus, for a non-diagonalisable matrix it could happen that some of the eigenspaces are too small, i.e. $\dim(E_\lambda) < m(\lambda)$ for some λ , so that Thm. 6.0.19 b) does not yield a basis of V , but only of a subspace. Remember that

$$E_\lambda = \ker(T - \lambda I_V) .$$

Definition 7.1.1. Let V be a finite-dimensional vector space over \mathbb{C} and let $T: V \rightarrow V$ be a linear transformation. Let $\lambda \in \mathbb{C}$ be an eigenvalue of T . The subspace $\bar{E}_\lambda \subset V$ defined by

$$\bar{E}_\lambda = \{v \in V \mid (T - \lambda I_V)^n(v) = 0 \text{ for some } n \in \mathbb{N}\}$$

is called the **generalised eigenspace** of T corresponding to the eigenvalue λ .

Any linear transformation $S: V \rightarrow V$ maps the zero vector to itself. Thus, if $v \in \ker(S^n)$, then $v \in \ker(S^{n+1})$. With $S = T - \lambda I_V$ we obtain by induction:

$$\ker((T - \lambda I_V)^n) \subset \ker((T - \lambda I_V)^{n+1}) \subset \dots \subset \ker((T - \lambda I_V)^{n+k})$$

for any $k \in \mathbb{N}_0$. This means that we could also write \bar{E}_λ as follows:

$$\bar{E}_\lambda = \bigcup_{n \in \mathbb{N}} \ker((T - \lambda I_V)^n) .$$

Observe that by definition $E_\lambda \subset \bar{E}_\lambda$.

Exercise 7.1.2. Check that \bar{E}_λ is a subspace of V .

Example 7.1.3. Consider the linear transformation $L_A: \mathbb{C}^2 \rightarrow \mathbb{C}^2$ on $V = \mathbb{C}^2$ given by the matrix

$$A = \begin{pmatrix} 3 & 1 \\ 0 & 3 \end{pmatrix}$$

The vector $e_1 = (1, 0)$ is an eigenvector of this transformation with eigenvalue $\lambda = 3$. The solution space of the matrix equation $Ax = 3x$ agrees with $\ker(L_A - 3I_V) = E_3$ and is given by

$$E_3 = \{(x, 0) \in \mathbb{C}^2 \mid x \in \mathbb{C}\} = \text{span}\{e_1\} .$$

The characteristic polynomial of A is $p_A(t) = (t - 3)^2$. Therefore the multiplicity $m(3)$ is 2, whereas $\dim(E_\lambda) = 1$. If I_2 denotes the identity matrix, then

$$(A - 3I_2)^2 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}^2 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

Therefore $\bar{E}_3 = \ker((L_A - 3I_V)^2) = \mathbb{C}^2$.

Example 7.1.4. Let us look at another example: This time take the linear transformation $L_A: \mathbb{C}^3 \rightarrow \mathbb{C}^3$ on $V = \mathbb{C}^3$ given by the matrix

$$A = \begin{pmatrix} 4 & 0 & 0 \\ 0 & 3 & 1 \\ 0 & 0 & 3 \end{pmatrix}$$

The characteristic polynomial of A is $p_A(t) = (t - 4)(t - 3)^2$. Therefore the eigenvalues of A are $\lambda_1 = 4$ and $\lambda_2 = 3$ with multiplicities $m(4) = 1$ and $m(3) = 2$. The vector $e_1 = (1, 0, 0)$ is an eigenvector of A with eigenvalue 4 and eigenspace

$$E_4 = \{(x, 0, 0) \in \mathbb{C}^3 \mid x \in \mathbb{C}\} = \text{span}\{e_1\} .$$

Likewise the vector $e_2 = (0, 1, 0)$ is an eigenvector of A corresponding to the eigenvalue 3 and a similar computation to the one in the last example shows that

$$E_3 = \{(0, x, 0) \in \mathbb{C}^3 \mid x \in \mathbb{C}\} = \text{span}\{e_2\} .$$

We obtain the generalised eigenspace \bar{E}_4 by looking at the null spaces of the linear transformations $(L_B)^n = L_{B^n}$ with

$$B = A - 4I_3 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 1 \\ 0 & 0 & -1 \end{pmatrix}$$

Exercise 7.1.5. Show by induction that

$$B^n = \begin{pmatrix} 0 & 0 & 0 \\ 0 & (-1)^n & (-1)^{n+1}n \\ 0 & 0 & (-1)^n \end{pmatrix}$$

and deduce that $\ker(L_{B^n}) = \text{span}\{e_1\}$.

The exercise shows that $E_4 = \bar{E}_4 = \text{span}\{e_1\}$. For the eigenvalue $\lambda_2 = 3$ we let $C = A - 3I_3$ and compute

$$C^2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}^2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

which implies

$$\bar{E}_3 = \ker((L_A - 3I_V)^2) = \ker(L_{C^2}) = \text{span}\{e_2, e_3\} .$$

Observe that $\dim(\bar{E}_3) = 2$ whereas $\dim(E_3) = 1$.

Exercise 7.1.6. Let $L_A: \mathbb{C}^4 \rightarrow \mathbb{C}^4$ be the linear transformation given by the matrix

$$\begin{pmatrix} 2 & 1 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 1 \\ 0 & 0 & 0 & 3 \end{pmatrix}$$

Show that the eigenvalues are $\lambda_1 = 2$ and $\lambda_2 = 3$ and compute the eigenspaces E_2 and E_3 and the generalised eigenspaces \bar{E}_2 and \bar{E}_3 .

Let V be a vector space over a field F and let $T: V \rightarrow V$ be a linear transformation. As already mentioned above we can make sense of powers of T by defining

$$T^n = \underbrace{T \circ \cdots \circ T}_{n \text{ times}} .$$

Likewise, since $\mathcal{L}(V, V)$ is a vector space, we can form linear combinations of linear maps. Therefore we can apply any polynomial p with coefficients in F to T . With $p(t) = a_n t^n + \cdots + a_1 t + a_0$ and $a_i \in F$ we have

$$p(T) = a_n T^n + \cdots + a_1 T + a_0 I_V .$$

In particular, it makes sense to apply the characteristic polynomial p_T to the linear transformation T itself. The following theorem that we are not going to prove here is a result by Cayley and Hamilton.

Theorem 7.1.7. *Let $T: V \rightarrow V$ be a linear transformation on a vector space V and let p be its characteristic polynomial. Then $p(T) = 0$.*

Exercise 7.1.8. Check the above theorem for the linear transformation $L_A: \mathbb{C}^2 \rightarrow \mathbb{C}^2$ given by the matrix

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

The next theorem gives a handy criterion to determine whether a linear transformation of the form $p(T)$ for some polynomial p is invertible. The proof, which we unfortunately have to omit, is an application of the so-called spectral mapping theorem.

Theorem 7.1.9. *Let V be a finite-dimensional vector space over \mathbb{C} . Let $T: V \rightarrow V$ be a linear transformation with eigenvalues $\lambda_1, \dots, \lambda_k$ and let p be a polynomial. Then $p(T)$ is invertible if and only if*

$$p(\lambda_i) \neq 0$$

for all $i \in \{1, \dots, k\}$.

Let $T: V \rightarrow V$ be a linear transformation as above and let λ_j be an eigenvalue of T . Observe that

$$(T - \lambda_j I_V) \circ T = T^2 - \lambda_j T = T \circ (T - \lambda_j I_V) .$$

By induction we obtain $(T - \lambda_j I_V)^m \circ T = T \circ (T - \lambda_j I_V)^m$ for any $m \in \mathbb{N}$. Let $v \in \bar{E}_{\lambda_j}$ and let $m \in \mathbb{N}$ be a positive integer with $(T - \lambda_j I_V)^m(v) = 0$ (such an m has to exist by the definition of \bar{E}_{λ_j}). Then we have

$$(T - \lambda_j I_V)^m(T(v)) = T((T - \lambda_j I_V)^m(v)) = 0 .$$

In particular, $T(v) \in \bar{E}_{\lambda_j}$ for any $v \in \bar{E}_{\lambda_j}$, i.e. T maps the subspace \bar{E}_{λ_j} into itself. Thus, $\bar{E}_{\lambda_j} \subset V$ is an invariant subspace in the sense of the following definition:

Definition 7.1.10. Let V be a vector space and let $T: V \rightarrow V$ be a linear transformation. A subspace $W \subset V$ is called an **invariant subspace for T** if $T(w) \in W$ for all $w \in W$. In this situation we will denote the restriction of T to W by

$$T|_W : W \rightarrow W$$

(note that we restrict the domain and the codomain).

The definition of \bar{E}_{λ} only implies that for each $v \in \bar{E}_{\lambda}$ there exists a positive integer $m \in \mathbb{N}$ with $(T - \lambda I_V)^m(v) = 0$. A priori this m will depend on the chosen vector v . However, the next lemma shows that there is an integer that works for all $v \in \bar{E}_{\lambda}$.

Lemma 7.1.11. *Let V be a finite-dimensional vector space over \mathbb{C} . Let $T: V \rightarrow V$ be a linear transformation with eigenvalues $\lambda_1, \dots, \lambda_r$, and let $\bar{E}_{\lambda_1}, \dots, \bar{E}_{\lambda_r}$ be the generalised eigenspaces. Let m_i be the multiplicity of the eigenvalue λ_i . Then*

$$(T - \lambda_j I_V)^{m_j}|_{\bar{E}_{\lambda_j}} = (T|_{\bar{E}_{\lambda_j}} - \lambda_j I_{\bar{E}_{\lambda_j}})^{m_j} = 0 .$$

for all $j \in \{1, \dots, r\}$.

Proof. Let p be the characteristic polynomial of T . Over the complex numbers any polynomial splits into linear factors corresponding to the roots of p . The roots agree with the eigenvalues of T . Hence, we obtain

$$p(t) = (t - \lambda_1)^{m_1} (t - \lambda_2)^{m_2} \dots (t - \lambda_r)^{m_r}$$

Fix $j \in \{1, \dots, r\}$ and define the polynomial p_j to be

$$p_j(t) = \frac{p(t)}{(t - \lambda_j)^{m_j}} = \prod_{i \neq j} (t - \lambda_i)^{m_i} .$$

Let $W = \bar{E}_{\lambda_j}$ and consider the restriction $S = T|_W : W \rightarrow W$. Let $w \in W$ be an eigenvector of S with corresponding eigenvalue λ . Since $w \in \bar{E}_{\lambda_j}$ there exists $m \in \mathbb{N}$ with

$$0 = (S - \lambda_j I_W)^m(w) = (\lambda - \lambda_j)^m w .$$

Thus, $\lambda = \lambda_j$, which means that λ_j is the only eigenvalue of S . By definition we have $p_j(\lambda_j) \neq 0$. Therefore Thm. 7.1.9 implies that $p_j(S)$ is invertible. Now note that $p(t) = (t - \lambda_j)^{m_j} p_j(t)$. Hence, Thm. 7.1.7 shows that

$$0 = p(T) = (T - \lambda_j I_V)^{m_j} p_j(T) .$$

If we restrict both sides of this equation to $W \subset V$ we obtain

$$0 = (S - \lambda_j I_W)^{m_j} p_j(S) .$$

However, $p_j(S)$ is invertible. Therefore we must have

$$(T|_{\bar{E}_{\lambda_j}} - \lambda_j I_{\bar{E}_{\lambda_j}})^{m_j} = (S - \lambda_j I_W)^{m_j} = 0 . \quad \square$$

The main result of this section is the following:

Theorem 7.1.12. *Let V be a finite-dimensional vector space over \mathbb{C} and let $T : V \rightarrow V$ be a linear transformation with distinct eigenvalues $\lambda_1, \dots, \lambda_r$. Let $\bar{E}_{\lambda_1}, \dots, \bar{E}_{\lambda_r}$ be the generalised eigenspaces. Then*

$$V = \bar{E}_{\lambda_1} \oplus \dots \oplus \bar{E}_{\lambda_r} .$$

Proof. Let p be the characteristic polynomial of T and denote the multiplicity of the i th eigenvalue by m_i . As in the proof of Lem. 7.1.11 we obtain

$$p(t) = (t - \lambda_1)^{m_1} (t - \lambda_2)^{m_2} \dots (t - \lambda_r)^{m_r}$$

Consider again the polynomials p_k defined by

$$p_k(t) = \frac{p(t)}{(t - \lambda_k)^{m_k}} = \prod_{j \neq k} (t - \lambda_j)^{m_j} .$$

and let $q(t) = \sum_{k=1}^r p_k(t)$. Observe that $q(\lambda_j) \neq 0$ for all $j \in \{1, \dots, r\}$. By Thm. 7.1.9 the linear transformation $B = q(T)$ is invertible. Now we define linear maps $P_k : V \rightarrow V$ for $k \in \{1, \dots, r\}$ as follows:

$$P_k = B^{-1} p_k(T) .$$

These transformations have the following properties:

- a) $\sum_{k=1}^r P_k = I_V$,
- b) $P_k|_{\bar{E}_{\lambda_j}} = 0$ if $j \neq k$,
- c) $\text{Im}(P_k) = \bar{E}_{\lambda_k}$ and $P_k v = v$ for all $v \in \bar{E}_{\lambda_k}$.

Property a) is a consequence of the following short computation:

$$\sum_{k=1}^r P_k = B^{-1} \left(\sum_{k=1}^r p_k(T) \right) = B^{-1} q(T) = B^{-1} B = I_V .$$

To see that Property b) holds, note that for $j \neq k$ the polynomial p_k contains a linear factor of the form $(t - \lambda_j)^{m_j}$, i.e. $p_k(t) = (t - \lambda_j)^{m_j} p_{jk}(t)$ for some polynomial $p_{jk}(t)$. But this means that

$$p_k(T)|_{\bar{E}_{\lambda_j}} = (T - \lambda_j I_V)^{m_j}|_{\bar{E}_{\lambda_j}} \circ p_{jk}(T)|_{\bar{E}_{\lambda_j}} = 0 \quad (18)$$

by Lem. 7.1.11. Thus, we also have

$$P_k|_{\bar{E}_{\lambda_j}} = B^{-1} p_k(T)|_{\bar{E}_{\lambda_j}} = 0 .$$

To show c) we first show that $\text{Im}(p_k(T)) \subset \bar{E}_{\lambda_k}$. By definition $p(t) = (t - \lambda_k)^{m_k} p_k(t)$. Let $w = p_k(T)v \in \text{Im}(p_k(T))$. Then

$$(T - \lambda_k I_V)^{m_k} w = (T - \lambda_k I_V)^{m_k} p_k(T)v = p(T)v = 0$$

by Thm. 7.1.7. Thus, $w \in \bar{E}_{\lambda_k}$, since w is annihilated by some power of $(T - \lambda_k I_V)$. Now note that $Tv \in \bar{E}_{\lambda_j}$ for any $v \in \bar{E}_{\lambda_j}$ and any $j \in \{1, \dots, r\}$. Thus, we also have $Bv = q(T)v \in \bar{E}_{\lambda_j}$ for any $v \in \bar{E}_{\lambda_j}$. This implies $B^{-1}v \in \bar{E}_{\lambda_j}$ for any $v \in \bar{E}_{\lambda_j}$. Let $w \in \text{Im}(P_k)$. Then we have $w = B^{-1} p_k(T)v$ for some $v \in V$ and combining the last two ideas we see that $w \in \bar{E}_{\lambda_k}$. Since w was arbitrary, this shows $\text{Im}(P_k) \subset \bar{E}_{\lambda_k}$.

Now let $v \in \bar{E}_{\lambda_k}$. By (18) we obtain

$$p_k(T)v = \sum_{j=1}^r p_k(T)v = Bv ,$$

which implies $P_k(v) = B^{-1} Bv = v$ and therefore $\bar{E}_{\lambda_k} \subset \text{Im}(P_k)$, since v was arbitrary.

Now we are in position to complete the proof of the theorem. Take $v \in V$ and define $v_k = P_k v$. Then according to Property c), $v \in \bar{E}_{\lambda_k}$ and by a)

$$v = I_V(v) = \sum_{j=1}^r P_j(v) = \sum_{j=1}^r v_j$$

and therefore

$$V \subset \bar{E}_{\lambda_1} + \dots + \bar{E}_{\lambda_r} \subset V ,$$

which implies $V = \bar{E}_{\lambda_1} + \dots + \bar{E}_{\lambda_r}$. Now suppose that $j \neq k$ and $w \in \bar{E}_{\lambda_j} \cap \bar{E}_{\lambda_k}$. Using Properties c) and b) we then obtain

$$w = P_j(w) = 0 ,$$

since $w \in \bar{E}_{\lambda_k}$. This finishes the proof. \square

Let $T: V \rightarrow V$ be a linear transformation as in Thm. 7.1.12, let $d_j = \dim(\bar{E}_{\lambda_j})$ and choose a basis $\beta_j = \{v_{j,1}, \dots, v_{j,d_j}\}$ for each generalised eigenspace \bar{E}_{λ_j} . By Thm. 7.1.12 and Lem. 7.0.2 the set

$$\beta = \beta_1 \cup \dots \cup \beta_r \quad (19)$$

is a basis of V . We order this basis by ordering each β_j and then taking the elements of β_1 first, followed by the elements of β_2 and so on. For each $j \in \{1, \dots, r\}$ the subspace \bar{E}_{λ_j} is invariant. Thus, we see that for each $i \in \{1, \dots, d_j\}$ the vector $T(v_{j,i})$ can be written as a linear combination of basis elements from β_j again without using any other set β_k with $k \neq j$. Thus, if we determine the matrix of T with respect to the basis β it will have the following form

$$[T]_{\beta}^{\beta} = \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_k \end{pmatrix}$$

where $A_j \in M_{d_j \times d_j}(\mathbb{C})$ is given by

$$A_j = \left[T|_{\bar{E}_{\lambda_j}} \right]_{\beta_j}^{\beta_j}.$$

The full matrix $[T]_{\beta}^{\beta}$ is not diagonal, since every A_j is still a square matrix itself. Matrices of this form are sometimes called **block diagonal**.

7.2 Nilpotent linear transformations

Can we improve the block diagonal decomposition of $T: V \rightarrow V$? We know already that we can not always choose a basis, such that the matrix representation is diagonal. We will first study the structure of each block A_j separately. Let m_j be the multiplicity of the eigenvalue λ_j . As was pointed out in Lem. 7.1.11 we have

$$\left(T|_{\bar{E}_{\lambda_j}} - \lambda_j \cdot I_{\bar{E}_{\lambda_j}} \right)^{m_j} = 0,$$

Linear transformations which vanish when raised to a certain power are quite common in Linear Algebra and have a special name.

Definition 7.2.1. Let V be a vector space and let $N: V \rightarrow V$ be a linear transformation. If $N^k = 0$ for some power $k \in \mathbb{N}$, then N is called a **nilpotent linear transformation**.

Our observations from the last section can be summarised as follows:

Corollary 7.2.2. *Let V be a finite-dimensional vector space over \mathbb{C} . Let $T: V \rightarrow V$ be a linear transformation. Then there are transformations $D: V \rightarrow V$ and $N: V \rightarrow V$ such that D is diagonalisable, N is nilpotent, $ND = DN$ and $T = D + N$.*

Proof. By Thm. 7.1.12 the vector space V has a direct sum decomposition of the form

$$V = \bar{E}_{\lambda_1} \oplus \cdots \oplus \bar{E}_{\lambda_r} .$$

In particular, by Exercise 7.0.3 each $v \in V$ has a unique decomposition $v = v_1 + \cdots + v_r$ with $v_i \in \bar{E}_{\lambda_i}$. Therefore the map

$$D(v) = \lambda_1 v_1 + \lambda_2 v_2 + \cdots + \lambda_r v_r$$

is well-defined. Let $d_i = \dim(\bar{E}_{\lambda_i})$. If we choose a basis β as described in the paragraph before (19) then

$$[D]_{\beta}^{\beta} = \begin{pmatrix} \lambda_1 \cdot I_{d_1} & 0 & \cdots & 0 \\ 0 & \lambda_2 \cdot I_{d_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \lambda_r \cdot I_{d_r} \end{pmatrix}$$

where I_{d_i} is the identity matrix of dimension d_i . In particular, the right hand side is a diagonal matrix and D is diagonalisable.

Now let $N = T - D$ and note that each \bar{E}_{λ_j} is an invariant subspace for N , D and T . Thus, if $v = v_1 + \cdots + v_r$ with $v_i \in \bar{E}_{\lambda_i}$, then $T(v_i) \in \bar{E}_{\lambda_i}$ and $T(v_1) + \cdots + T(v_r)$ is the decomposition of $T(v)$. Therefore

$$TD(v) = T(\lambda_1 v_1 + \cdots + \lambda_r v_r) = \lambda_1 T(v_1) + \lambda_2 T(v_2) + \cdots + \lambda_r T(v_r) = DT(v) .$$

We obtain $DT = TD$, which implies $DN = DT - D^2 = TD - D^2 = ND$. Using the definition of D we get

$$N|_{\bar{E}_{\lambda_j}} = (T - \lambda_j I_V)|_{\bar{E}_{\lambda_j}} .$$

This is a nilpotent linear transformation by Lem. 7.1.11. Let m_i be the multiplicity of λ_i and note that $n = \dim(V) \geq m_i$. Thus,

$$\left(N|_{\bar{E}_{\lambda_j}} \right)^n = 0 .$$

Let $v \in V$ and let $v = v_1 + \cdots + v_r$ be its decomposition with $v_i \in \bar{E}_{\lambda_i}$. Then

$$N^n(v) = N^n(v_1) + N^n(v_2) + \cdots + N^n(v_r) = 0 .$$

Therefore $N: V \rightarrow V$ is nilpotent as well. □

Example 7.2.3. Let $V = \mathbb{R}^3$ and consider the linear transformation $N: V \rightarrow V$ given by

$$N(x, y, z) = (y, z - x - y, z - x)$$

With respect to the standard basis $\alpha = \{e_1, e_2, e_3\}$ the matrix representation of N is

$$[N]_{\alpha}^{\alpha} = \begin{pmatrix} 0 & 1 & 0 \\ -1 & -1 & 1 \\ -1 & 0 & 1 \end{pmatrix}$$

This helps us to compute the powers of N :

$$[N^2]_{\alpha}^{\alpha} = \begin{pmatrix} -1 & -1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & 1 \end{pmatrix} , \quad [N^3]_{\alpha}^{\alpha} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

In particular, $N^3 = 0$ and N is nilpotent. What is the most convenient basis for a matrix representation of N ? The only eigenvector of N is $(1, 0, 1)$ with corresponding eigenvalue 0 and its eigenspace is one-dimensional. Hence, N is not diagonalisable. However, the vector $(0, 1, 1)$ satisfies $N(0, 1, 1) = (1, 0, 1)$ and the vector $(0, 0, 1)$ has image $N(0, 0, 1) = (0, 1, 1)$. Thus, even though we can not construct a basis of vectors $w_i \in \mathbb{R}^3$ with $N(w_i) = \lambda_i w_i$, we can instead look at the vectors $v_1 = (1, 0, 1)$, $v_2 = (0, 1, 1)$ and $v_3 = (0, 0, 1)$ and consider the set $\beta = \{v_1, v_2, v_3\}$ instead. As mentioned above these vectors satisfy

$$N(v_1) = 0 \quad , \quad N(v_2) = v_1 \quad , \quad N(v_3) = v_2 .$$

Note that β is an ordered basis for \mathbb{R}^3 with respect to which the matrix representation of N takes the following form

$$[N]_{\beta}^{\beta} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} ,$$

i.e. it has only ones along the first line above the diagonal and zeroes everywhere else.

The above example tells us that we should look for sets of vectors that fit the following definition:

Definition 7.2.4. Let $N: V \rightarrow V$ be a nilpotent linear transformation. A chain of non-zero vectors v_1, \dots, v_p , which satisfies the following relations:

$$\begin{aligned} N(v_1) &= 0 , \\ N(v_2) &= v_1 , \\ &\vdots \\ N(v_p) &= v_{p-1} \end{aligned}$$

is called a **cycle of generalised eigenvectors** of N . The vector v_1 is called the **initial vector** of the cycle, the vector v_p is called the **end vector** of the cycle, and the number p is called the **length** of the cycle.

In Example 7.2.3 we found a *basis* consisting of a cycle of generalised eigenvectors of N . The next theorem hands us a tool to check if the union of multiple such cycles yields a linear independent set.

Theorem 7.2.5. *Let V be a finite-dimensional vector space. Let $N: V \rightarrow V$ be a nilpotent linear transformation, and let $\mathcal{C}_1, \dots, \mathcal{C}_r$ be cycles of generalised eigenvectors, $\mathcal{C}_k = \{v_1^{(k)}, \dots, v_{p_k}^{(k)}\}$, p_k being the length of the cycle \mathcal{C}_k . Assume that the set of initial vectors $\{v_1^{(1)}, \dots, v_1^{(r)}\}$ is linearly independent. Then no vectors belong to two cycles, and the union of all the vectors from all the cycles is a linearly independent set.*

Proof. Let $n = p_1 + p_2 + \dots + p_r$ be the total number of vectors in all cycles. We will prove the theorem by induction over n . In the case $n = 1$ we only have one cycle containing just one non-zero vector. The set containing this vector is linearly independent, so the theorem holds in this case.

Now let us assume that the theorem is true for all linear transformations and for all collections of cycles, as long as the total number of vectors in their union is strictly less than n . Note that the definition of a cycle of generalised eigenvectors implies that the linear span of all vectors $v_j^{(k)}$, i.e. the subspace $W = \text{span}(S)$ with

$$S = \{v_j^{(k)} \in V \mid j \in \{1, \dots, p_k\}, k \in \{1, \dots, r\}\}$$

is an invariant subspace for N . Therefore we can without loss of generality consider $N|_W: W \rightarrow W$ instead and assume that $V = W$.

Consider the subspace $\text{Im}(N) \subset V$. It follows from Def. 7.2.4 that this space is spanned by the set of vectors

$$S' = \{v_j^{(k)} \in V \mid j \in \{1, \dots, p_k - 1\}, k \in \{1, \dots, r\}\} .$$

If $p_k > 1$, then $\mathcal{C}'_k = \{v_1^{(k)}, \dots, v_{p_k-1}^{(k)}\}$ is a cycle of generalised eigenvectors. If $p_k = 1$, i.e. if \mathcal{C}_k contains only one vector, then it is annihilated by N . Therefore we have finitely many cycles, and the initial vectors of the cycles \mathcal{C}'_i are linearly independent. By our induction hypothesis S' is then also linearly independent. Since these vectors span $\text{Im}(N)$, they form a basis for it. Hence,

$$\text{rank}(N) = \dim(\text{Im}(N)) = n - r ,$$

since we removed the vectors $v_{p_k}^{(k)}$ for $k \in \{1, \dots, r\}$ from S to obtain S' . On the other hand $N(v_1^{(k)}) = 0$ for $k \in \{1, \dots, r\}$. Since these vectors are linearly independent by assumption, we have $\dim(\ker(N)) \geq r$. By the rank-nullity theorem

$$\dim(V) = \text{rank}(N) + \text{nullity}(N) = n - r + \dim(\ker(N)) \geq n - r + r = n .$$

But V is also spanned by the set S , which contains n vectors. Thus, S forms a basis for V . In particular, S is linearly independent. \square

The following theorem shows that a basis consisting of cycles of generalised eigenvectors (similar to the one we constructed in Example 7.2.3) exists for any nilpotent linear transformation.

Theorem 7.2.6. *Let V be a finite-dimensional vector space and let $N: V \rightarrow V$ be a nilpotent linear transformation. Then V has a basis that consists of a union of cycles of generalised eigenvectors for N .*

Proof. Let $n = \dim(V)$. We will prove the statement by induction over n . If $\dim(V) = 1$, then any non-zero linear transformation is invertible. Since N is nilpotent, we must have $N = 0$ in this case. But then any non-zero vector $v \in V$ is a basis for V and a cycle of generalised eigenvectors of length 1 for N .

Now assume that the statement is true for any linear transformation acting on a vector space V of dimension less than n . Since N is nilpotent, it is not invertible and $\text{rank}(N) < \dim(V)$. Therefore we can apply the induction hypothesis to $\text{Im}(N) \subset V$. This means there exist cycles $\mathcal{C}_1, \dots, \mathcal{C}_r$ of generalised eigenvectors for N such that their union is a basis for $\text{Im}(N)$. Let

$$\mathcal{C}_k = \{v_1^{(k)}, \dots, v_{p_k}^{(k)}\} \subset \text{Im}(N) ,$$

where p_k is the length of \mathcal{C}_k and $v_1^{(k)}$ is its initial vector. Note that $v_{p_k}^{(k)}$ is in the image of N . Hence, there exists $v_{p_k+1}^{(k)} \in V$ with the property $N(v_{p_k+1}^{(k)}) = v_{p_k}^{(k)}$. Thus, we can extend \mathcal{C}_k to a bigger cycle

$$\tilde{\mathcal{C}}_k = \{v_1^{(k)}, \dots, v_{p_k}^{(k)}, v_{p_k+1}^{(k)}\} .$$

The initial vector of $\tilde{\mathcal{C}}_k$ agrees with the initial vector of \mathcal{C}_k . In particular, the set of all initial vectors is still linearly independent. By Thm. 7.2.5 the union of all cycles $\tilde{\mathcal{C}}_1, \dots, \tilde{\mathcal{C}}_r$ is then also linearly independent.

Now consider the set $\{v_1^{(1)}, \dots, v_1^{(r)}\}$ consisting of all initial vectors of all cycles. By Def. 7.2.4 and our observations above this is a linearly independent subset of $\ker(N)$. Hence, we can extend it to a basis of $\ker(N)$ by adding vectors $u_1, \dots, u_q \in \ker(N)$ (the case $q = 0$ might happen and corresponds to the situation where the set of initial vectors is already a basis and we do not need to add anything). Each vector u_i might be considered as a cycle $\tilde{\mathcal{C}}_{r+i} = \{u_i\}$ of length 1. Since the set of initial vectors of all cycles $\tilde{\mathcal{C}}_1, \dots, \tilde{\mathcal{C}}_r, \tilde{\mathcal{C}}_{r+1}, \dots, \tilde{\mathcal{C}}_{r+q}$ is still linearly independent, Thm. 7.2.5 implies that the union of all these cycles is linearly independent as well.

To show that this union actually forms a basis, we count the dimensions of the spaces involved. We know that the cycles $\mathcal{C}_1, \dots, \mathcal{C}_r$ have $\text{rank}(N) = \dim(\text{Im}(N))$ vectors in total. Each cycle $\tilde{\mathcal{C}}_k$ was obtained by adding a single vector to \mathcal{C}_k . Therefore the total number of vectors in the union of all cycles $\tilde{\mathcal{C}}_1, \dots, \tilde{\mathcal{C}}_r$ is $\text{rank}(N) + r$.

We know that $\dim(\ker(N)) = r + q$, since $\{v_1^{(1)}, \dots, v_1^{(r)}, u_1, \dots, u_q\}$ is a basis for the kernel. Likewise, the union of all $\tilde{\mathcal{C}}_1, \dots, \tilde{\mathcal{C}}_r, \tilde{\mathcal{C}}_{r+1}, \dots, \tilde{\mathcal{C}}_{r+q}$ contains $\text{rank}(N) + r + q$ vectors in total. But this number of vectors agrees with $\dim(V)$, since

$$\text{rank}(N) + r + q = \text{rank}(N) + \text{nullity}(N) = \dim(V) ,$$

which implies that the union of $\tilde{\mathcal{C}}_1, \dots, \tilde{\mathcal{C}}_r, \tilde{\mathcal{C}}_{r+1}, \dots, \tilde{\mathcal{C}}_{r+q}$ is actually a basis. \square

Remark 7.2.7. Let $N: V \rightarrow V$ be a nilpotent linear transformation and let \mathcal{C}_1 be a cycle of generalised eigenvectors of length 3, \mathcal{C}_2 be a cycle of length 1 and \mathcal{C}_3 be another one of length 2. Let $\beta = \mathcal{C}_1 \cup \mathcal{C}_2 \cup \mathcal{C}_3$. Assume that this is a basis for V . Each cycle is already ordered and we order β by taking the elements of \mathcal{C}_1 first, then \mathcal{C}_2 and so on. The corresponding matrix representation of N is

$$[N]_{\beta}^{\beta} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Note that the size of the coloured blocks agrees with the length of the cycles: \mathcal{C}_1 yields a block of size 3, \mathcal{C}_2 one of size 1 and \mathcal{C}_3 another one of size 2.

The **main theorem** of this section, shown below, states that for any linear transformation $T: V \rightarrow V$ we can always find a basis β with the property that the only non-zero entries of the matrix $[T]_{\beta}^{\beta}$ are either on the diagonal or on the line above the diagonal. Moreover, β can be chosen in such a way that the entries on the line above the diagonal are either 0 or 1.

Theorem 7.2.8. *Let V be a finite-dimensional vector space over \mathbb{C} and let $T: V \rightarrow V$ be a linear transformation. There exists an ordered basis β (called a **Jordan basis**) such that $[T]_{\beta}^{\beta}$ is a block diagonal matrix consisting of blocks of the form*

$$J_k(\lambda) = \begin{pmatrix} \lambda & 1 & 0 & \cdots & 0 & 0 \\ 0 & \lambda & 1 & \cdots & 0 & 0 \\ 0 & 0 & \lambda & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \lambda & 1 \\ 0 & 0 & 0 & \cdots & 0 & \lambda \end{pmatrix} \in M_{k \times k}(\mathbb{C}) ,$$

where λ is an eigenvalue of T . By definition a block of size 1 is just λ . The blocks $J_k(\lambda)$ are called **Jordan blocks**.

Proof. Let $\lambda_1, \dots, \lambda_r \in \mathbb{C}$ be the eigenvalues of T and let \bar{E}_{λ_i} be the corresponding generalised eigenspaces. Let

$$N_k = (T - \lambda_k I_V)|_{\bar{E}_{\lambda_k}} .$$

By Lem. 7.1.11 each N_k is nilpotent and using Thm. 7.2.6 we can find a basis β_k for \bar{E}_{λ_k} that consists of a union of cycles of generalised eigenvectors for N_k . Note that

$$T|_{\bar{E}_{\lambda_k}} = N_k + \lambda_k I_{\bar{E}_{\lambda_k}} .$$

Therefore the matrix representation of $T|_{\bar{E}_{\lambda_k}}$ is given by

$$\left[T|_{\bar{E}_{\lambda_k}} \right]_{\beta_k}^{\beta_k} = [N_k]_{\beta_k}^{\beta_k} + \lambda_k \cdot \left[I_{\bar{E}_{\lambda_k}} \right]_{\beta_k}^{\beta_k} = [N_k]_{\beta_k}^{\beta_k} + \lambda_k \cdot I_{d_k} ,$$

where $d_k = \dim(\bar{E}_{\lambda_k})$. By our observation in the previous remark the matrix $[N_k]_{\beta_k}^{\beta_k}$ is only non-zero on the line above the diagonal, where it may contain zeroes and ones. If we add $\lambda_k \cdot I_{d_k}$ to it, then the resulting matrix will consist of blocks of the form $J_\ell(\lambda_k)$ for various values of ℓ . If we now choose $\beta = \beta_1 \cup \dots \cup \beta_r$ as in (19), then β is a basis for V and $[T]_\beta^\beta$ will have the desired form. \square

Definition 7.2.9. Let F be a field. If a matrix $A \in M_{n \times n}(F)$ has the form described in Thm. 7.2.8, then we say that A is in **Jordan normal form**. In particular, if $T: V \rightarrow V$ is a linear transformation on a vector space V and β is a Jordan basis, then $[T]_\beta^\beta$ is in Jordan normal form.

Corollary 7.2.10. Let V be a finite-dimensional vector space over \mathbb{C} and let $T: V \rightarrow V$ be a linear transformation with distinct eigenvalues $\lambda_1, \dots, \lambda_r$. Denote by $m_i \in \mathbb{N}$ the algebraic multiplicity of the eigenvalue λ_i . Then we have

$$\dim(\bar{E}_{\lambda_i}) = m_i$$

for all $i \in \{1, \dots, r\}$.

Proof. Let β_i be the basis for the generalised eigenspace \bar{E}_{λ_i} constructed in Thm. 7.2.8. Then $\beta = \beta_1 \cup \dots \cup \beta_r$ is a basis for V with the property that $[T]_\beta^\beta$ is in Jordan normal form. We make the following observations: The characteristic polynomial of a Jordan block $J_k(\lambda)$ of size $k \in \mathbb{N}$ is given by

$$p_{J_k(\lambda)}(t) = (-1)^k (t - \lambda)^k .$$

Moreover, for a block diagonal matrix A with

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

we have $p_A(t) = p_{A_1}(t) \cdots p_{A_k}(t)$. Hence, if we compute the characteristic polynomial of T using the basis β , we see that each Jordan block $J_k(\lambda)$ contributes a factor $(t - \lambda)^k$ to $p_T(t)$. Thus, the algebraic multiplicity m_i of λ_i is the sum $s(\lambda_i)$ of all the sizes of the Jordan blocks with eigenvalue λ_i . But these are exactly the blocks that arise from the basis β_i of \bar{E}_{λ_i} . Each basis element in β_i is responsible for one column and row in a Jordan block for λ_i . Thus, we have

$$\dim(\bar{E}_{\lambda_i}) = s(\lambda_i) = m_i . \quad \square$$

Example 7.2.11. The matrix $A \in M_{3 \times 3}(\mathbb{C})$ given below is in Jordan normal form:

$$A = \begin{pmatrix} 3 & 1 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix} .$$

Note that this matrix consists of **two** Jordan blocks. The first one is

$$J_2(3) = \begin{pmatrix} 3 & 1 \\ 0 & 3 \end{pmatrix} \in M_{2 \times 2}(\mathbb{C})$$

and the second one is

$$J_1(3) = 3 \in M_{1 \times 1}(\mathbb{C}) .$$

This also illustrates that the eigenvalues of two blocks can agree.

Example 7.2.12. The Jordan normal form of a linear transformation is not unique. Let A be the matrix from Example 7.2.11, let $\alpha = \{e_1, e_2, e_3\}$ be the ordered standard basis and let $\beta = \{e_3, e_1, e_2\}$ be the basis that agrees with α as a set, but has a different order. Then we have

$$A = [L_A]_{\alpha}^{\alpha} = \begin{pmatrix} 3 & 1 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix} , \quad [L_A]_{\beta}^{\beta} = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 3 & 1 \\ 0 & 0 & 3 \end{pmatrix}$$

Both of these matrices are in Jordan normal form. Note that they agree up to a permutation of the blocks from Thm. 7.2.8. It turns out that the Jordan normal form is unique up to such a permutation of the Jordan blocks, but we will not prove this statement in this lecture.

7.3 Low dimensional examples

In this section we will discuss how to compute Jordan normal forms for linear transformations $T: V \rightarrow V$ in the cases where $V = \mathbb{C}^2$ or $V = \mathbb{C}^3$. The first step in all cases discussed below is determining the characteristic polynomial p_T and the eigenvalues $\lambda_1, \dots, \lambda_r$, i.e. the roots of p_T . We will assume this has been done. Then we are left with the following cases:

7.3.1 Dimension 2, two distinct eigenvalues $\lambda_1 \neq \lambda_2$

In this case all eigenvalues of T are distinct. Therefore by Cor. 6.0.11 the linear transformation T is diagonalisable. To find the basis β such that $[T]_{\beta}^{\beta}$ is a diagonal matrix determine the eigenspaces E_{λ_1} and E_{λ_2} . Any pair $\beta = \{v_1, v_2\}$ of non-zero vectors with $v_i \in E_{\lambda_i}$ will form a basis of \mathbb{C}^2 and

$$[T]_{\beta}^{\beta} = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}$$

is in Jordan normal form with Jordan blocks $J_1(\lambda_1) = \lambda_1$ and $J_1(\lambda_2) = \lambda_2$. Our computation of the formula for the elements in the Fibonacci sequence (in Example 6.0.8) is of this type.

7.3.2 Dimension 2, one eigenvalue $\lambda = \lambda_1 = \lambda_2$

If the two eigenvalues of a linear transformation $T: \mathbb{C}^2 \rightarrow \mathbb{C}^2$ agree, then we have to determine E_λ and compute its dimension to see whether T is diagonalisable or not. If $\dim(E_\lambda) = 2$, then T is diagonalisable and any basis β for E_λ will have the property that

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix} .$$

If $\dim(E_{\lambda_1}) = 1$, then T is not diagonalisable by Thm. 6.0.19. Since there is just one entry above the diagonal in a 2×2 -matrix, the Jordan normal form is completely determined. In particular, there has to be a basis β of $\bar{E}_\lambda = \mathbb{C}^2$ such that

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} .$$

To construct β we first have to find a basis $\beta' = \{v_1\}$ for the one-dimensional subspace E_λ . Then we have to extend β' to a cycle of generalised eigenvectors for the nilpotent operator

$$N = T - \lambda I_V .$$

Hence, according to Def. 7.2.4 we have to find $v_2 \in \mathbb{C}^2$ such that

$$N(v_2) = T(v_2) - \lambda v_2 = v_1 .$$

By Thm. 7.2.6 the set $\beta = \{v_1, v_2\}$ is a basis for V and

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix} .$$

Example 7.3.1. Consider the matrix $A \in M_{2 \times 2}(\mathbb{C})$ given by

$$A = \begin{pmatrix} -5 & -4 \\ 1 & -1 \end{pmatrix} .$$

Its characteristic polynomial $p_A(t)$ computes to

$$p_A(t) = \det \begin{pmatrix} -5-t & -4 \\ 1 & -1-t \end{pmatrix} = (5+t)(1+t) + 4 = t^2 + 6t + 9 = (t+3)^2$$

From this we see that $\lambda = -3$ is the only eigenvalue of A . The eigenspace E_{-3} is $\ker(L_B)$ with

$$B = A + 3 \cdot I_2 = \begin{pmatrix} -2 & -4 \\ 1 & 2 \end{pmatrix}$$

Solving the corresponding system of linear equations yields

$$E_{-3} = \{(x, y) \in \mathbb{C}^2 \mid x = -2y\} = \{a(-2, 1) \in \mathbb{C}^2 \mid a \in \mathbb{C}\} .$$

A basis for E_{-3} is given by the set $\beta' = \{v_1\}$ with $v_1 = (-2, 1)$. To extend this to a cycle of generalised eigenvectors for the nilpotent linear transformation $N = L_A + 3 \cdot I_{\mathbb{C}^2} = L_B$ we have to solve the linear equation $Bv_2 = v_1$, which yields with $v_2 = (x, y)$

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

The solution set of this equation is

$$S = \{(x, y) \in \mathbb{C}^2 \mid x + 2y = 1\}.$$

Let $v_2 = (-1, 1)$. This is an element of S . The set $\beta = \{v_1, v_2\}$ is the basis we were looking for. Indeed, we can compute

$$\begin{aligned} L_A(v_1) &= Av_1 = \begin{pmatrix} -5 & -4 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} -2 \\ 1 \end{pmatrix} = \begin{pmatrix} 6 \\ -3 \end{pmatrix} = (-3) \cdot v_1 + 0 \cdot v_2, \\ L_A(v_2) &= Av_2 = \begin{pmatrix} -5 & -4 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} -1 \\ 1 \end{pmatrix} = \begin{pmatrix} 1 \\ -2 \end{pmatrix} = v_1 + (-3) \cdot v_2. \end{aligned}$$

Hence, the matrix representation with respect to the ordered basis $\beta = \{v_1, v_2\}$ is

$$[L_A]_{\beta}^{\beta} = \begin{pmatrix} -3 & 1 \\ 0 & -3 \end{pmatrix}.$$

7.3.3 Dimension 3, three distinct eigenvalues

Just as in the 2-dimensional case Cor. 6.0.11 applies and shows that T is diagonalisable. To find a corresponding basis β , with respect to which the matrix of T is diagonal, choose basis vectors v_i for the three one-dimensional eigenspaces E_{λ_i} . The set $\beta = \{v_1, v_2, v_3\}$ is then a basis for \mathbb{C}^3 by Cor. 6.0.11 and

$$[T]_{\beta}^{\beta} = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}.$$

Since the setting is very similar to the 2-dimensional case, we will not give another example here.

7.3.4 Dimension 3, two distinct eigenvalues

Let $T: \mathbb{C}^3 \rightarrow \mathbb{C}^3$ be a linear transformation with two distinct eigenvalues λ_1 and λ_2 . Without loss of generality we can assume that the algebraic multiplicity m_1 of λ_1 is 2 and, likewise, that $m_2 = 1$. This means that the characteristic polynomial has the form

$$p_T(t) = c(t - \lambda_1)^2(t - \lambda_2)$$

for some constant $c \in \mathbb{C}$. If the eigenspace E_{λ_1} has dimension 2, then T is diagonalisable by Thm. 6.0.19. Moreover, if $\beta_1 = \{v_1, v_2\}$ is a basis for E_{λ_1} and $\beta_2 = \{v_3\}$ is a basis for E_{λ_2} , then $\beta = \beta_1 \cup \beta_2 = \{v_1, v_2, v_3\}$ is a basis for \mathbb{C}^3 with the property that

$$[T]_{\beta}^{\beta} = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix} .$$

The last remaining case occurs when the eigenspace E_{λ_1} only has dimension 1. Let v_1 be a basis vector for E_{λ_1} . To extend this to a cycle of generalised eigenvectors for the nilpotent linear transformation

$$N = T - \lambda_1 \cdot I_V$$

we have to find a vector $v_2 \in \mathbb{C}^3$ that satisfies the equation

$$N(v_2) = T(v_2) - \lambda_1 v_2 = v_1$$

just as in the 2-dimensional case. In particular, note that

$$(T - \lambda_1 \cdot I_V)^2(v_2) = N^2(v_2) = N(v_1) = 0 .$$

Hence, $v_2 \in \bar{E}_{\lambda_1}$. Let v_3 be a basis vector for the one-dimensional eigenspace E_{λ_2} . By Thm. 7.2.5 and Thm. 7.1.12 the set $\beta = \{v_1, v_2, v_3\}$ is a basis for \mathbb{C}^3 and

$$[T]_{\beta}^{\beta} = \begin{pmatrix} \lambda_1 & 1 & 0 \\ 0 & \lambda_1 & 0 \\ 0 & 0 & \lambda_2 \end{pmatrix} .$$

Example 7.3.2. Let $L_A: \mathbb{C}^3 \rightarrow \mathbb{C}^3$ be the linear transformation given by left multiplication by the matrix

$$A = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 3 & -1 \\ -1 & 0 & 2 \end{pmatrix} .$$

To find a basis β such that $[L_A]_{\beta}^{\beta}$ is in Jordan normal form we first have to determine the characteristic polynomial. We have

$$p_A(t) = \det \begin{pmatrix} 3-t & 0 & 0 \\ 0 & 3-t & -1 \\ -1 & 0 & 2-t \end{pmatrix} = (3-t)^2(2-t) = -(t-3)^2(t-2) ,$$

where we used cofactor expansion along the first row to compute the determinant. Therefore the eigenvalues are $\lambda_1 = 3$ with algebraic multiplicity $m_1 = 2$ and $\lambda_2 = 2$ with multiplicity $m_2 = 1$. The eigenspace E_3 is the kernel of L_B with

$$B = A - 3I_3 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ -1 & 0 & -1 \end{pmatrix}$$

Thus, E_3 is given by

$$E_3 = \{(x, y, z) \in \mathbb{C}^3 \mid z = x = 0\} = \{a(0, 1, 0) \in \mathbb{C}^3 \mid a \in \mathbb{C}\}.$$

Let $v_1 = (0, 1, 0)$. Then $\{v_1\}$ is a basis for E_3 . The equation $Av_2 - 3v_2 = Bv_2 = v_1$ for $v_2 = (x, y, z)$ leads to

$$\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ -1 & 0 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}.$$

The solution set of this system is given by

$$S = \{(1, y, -1) \in \mathbb{C}^3 \mid y \in \mathbb{C}\}$$

Let $v_2 = (1, 0, -1)$. Then $v_2 \in S$ and $\beta_1 = \{v_1, v_2\}$ is a basis for the generalised eigenspace \bar{E}_3 . The eigenspace E_2 is the kernel of L_C with

$$C = A - 2 \cdot I_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & -1 \\ -1 & 0 & 0 \end{pmatrix}$$

Thus, we have

$$E_2 = \{(x, y, z) \in \mathbb{C}^3 \mid x = 0, y = z\} = \{a(0, 1, 1) \in \mathbb{C}^3 \mid a \in \mathbb{C}\}$$

and with $v_3 = (0, 1, 1)$ the set $\beta_2 = \{v_3\}$ is a basis for E_2 . Altogether we obtain a basis $\beta = \beta_1 \cup \beta_2 = \{v_1, v_2, v_3\}$ and with respect to the β the matrix representation of L_A takes the form

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

7.3.5 Dimension 3, one eigenvalue

Now let $T: \mathbb{C}^3 \rightarrow \mathbb{C}^3$ be a linear transformation with a single eigenvalue $\lambda \in \mathbb{C}$. In this case the characteristic polynomial of T takes the form $p_T(t) = c(t - \lambda)^3$ and the multiplicity m of λ is 3. Depending on the dimension of the eigenspace E_λ we have to distinguish three cases: If $\dim(E_\lambda) = 3 = m$, then $E_\lambda = V$ and by Thm. 6.0.19 the transformation T is diagonalisable. Let β be any basis for E_λ . Then we have

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{pmatrix} = \lambda \cdot I_3.$$

Observe that this case can only happen if T is already of the form $\lambda \cdot I_V$, since this is the only linear transformation that has a matrix representation which is a multiple of the identity matrix.

Suppose that $\dim(E_\lambda) = 2$. Let $\{v'_1, v'_2\}$ be a basis for E_λ . Each of the eigenvectors v'_1, v'_2 can be seen as a cycle of generalised eigenvectors of length 1. But we only know that a linear combination of v'_1 and v'_2 can be extended to a cycle of length 2. Therefore we have to replace the basis $\{v'_1, v'_2\}$ by one that is adapted to this case. The Jordan normal form for T with respect to the basis β , that we still need to construct, will look like this:

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 1 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{pmatrix} .$$

In particular, we know that for $N = T - \lambda I_V$ we have $N^2 = 0$, since

$$[N^2]_\beta^\beta = ([N]_\beta^\beta)^2 = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}^2 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} .$$

Let $v_2 \in V$ be any vector that is not in E_λ and let $v_1 = N(v_2)$. Then we have $v_1 \neq 0$ and $N(v_1) = N^2(v_2) = 0$. Therefore $v_2 \in \bar{E}_\lambda$ and $v_1 \in E_\lambda$. Extend $\{v_1\}$ by a vector $v_3 \in E_\lambda$ to a basis for E_λ . Then $\beta = \{v_1, v_2, v_3\}$ will be basis for \mathbb{C}^3 and the matrix representation will be

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 1 & 0 \\ 0 & \lambda & 0 \\ 0 & 0 & \lambda \end{pmatrix} .$$

The last case to discuss is $\dim(E_\lambda) = 1$. Let v_1 be a basis vector for E_λ . To extend this to a cycle of generalised eigenvectors we now have to solve two equations. We define v_2 as above to be the solution of

$$T(v_2) - \lambda v_2 = v_1 .$$

Likewise, v_3 is obtained as the solution of

$$T(v_3) - \lambda v_3 = v_2 .$$

Note that these two equations ensure

$$\begin{aligned} T(v_2) &= v_1 + \lambda v_2 , \\ T(v_3) &= v_2 + \lambda v_3 . \end{aligned}$$

Moreover, $v_2 \in \ker((T - \lambda I_V)^2) \subset \bar{E}_\lambda$ and $v_3 \in \ker((T - \lambda I_V)^3) \subset \bar{E}_\lambda$. In particular, $\beta = \{v_1, v_2, v_3\}$ is a cycle of generalised eigenvectors for $N = T - \lambda I_V$. Hence, by Thm. 7.2.6 it is a basis. The matrix representation with respect to β yields

$$[T]_\beta^\beta = \begin{pmatrix} \lambda & 1 & 0 \\ 0 & \lambda & 1 \\ 0 & 0 & \lambda \end{pmatrix} .$$

Example 7.3.3. Consider the linear transformation L_A with $A \in M_{3 \times 3}(\mathbb{C})$ given by

$$A = \begin{pmatrix} 5 & -4 & 0 \\ 1 & 1 & 0 \\ 2 & -3 & 3 \end{pmatrix} .$$

The characteristic polynomial $p_A(t)$ is

$$\begin{aligned} p_A(t) &= \det \begin{pmatrix} 5-t & -4 & 0 \\ 1 & 1-t & 0 \\ 2 & -3 & 3-t \end{pmatrix} = (3-t) \cdot ((5-t)(1-t) + 4) \\ &= (3-t)(t^2 - 6t + 9) = -(t-3)^3 \end{aligned}$$

Hence, A has just one eigenvalue $\lambda = 3$. The eigenspace E_3 is the kernel of L_B with

$$B = \begin{pmatrix} 2 & -4 & 0 \\ 1 & -2 & 0 \\ 2 & -3 & 0 \end{pmatrix} .$$

To determine $\ker(L_B)$ we have to solve the following system of linear equations:

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

We leave it as an exercise to check that the solution set is given by

$$E_3 = \ker(L_B) = \{a(0, 0, 1) \mid a \in \mathbb{C}\} .$$

Let $v_1 = (0, 0, 1)$. The set $\{v_1\}$ is a basis for E_3 . Note that $N = T - 3I_V = L_B$. Therefore to construct v_2 we need to find the solution set of

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

Since the first two rows are equivalent, this leads to the system of equations

$$\begin{aligned} x - 2y &= 0 , \\ 2x - 3y &= 1 . \end{aligned}$$

The vector $v_2 = (2, 1, 0)$ is a solution. We can construct v_3 in a similar way by finding a solution of the system

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

This leads to the two equations:

$$\begin{aligned}x - 2y &= 1, \\2x - 3y &= 0,\end{aligned}$$

which is solved for example by the vector $v_3 = (-3, -2, 0)$. By the above observations the set $\beta = \{v_1, v_2, v_3\}$ is a basis for \mathbb{C}^3 and

$$[L_A]_{\beta}^{\beta} = \begin{pmatrix} 3 & 1 & 0 \\ 0 & 3 & 1 \\ 0 & 0 & 3 \end{pmatrix}.$$

8 Duals, symmetric bilinear forms and inner product spaces

Consider the vector space \mathbb{R}^2 over \mathbb{R} . The dot product (or standard inner product) of two vectors $(a_1, a_2), (b_1, b_2) \in \mathbb{R}^2$ is defined as

$$\langle (a_1, a_2), (b_1, b_2) \rangle = a_1 b_1 + a_2 b_2.$$

In fact, we can extend this definition to \mathbb{R}^n for arbitrary $n \in \mathbb{N}$ by defining

$$\langle (a_1, \dots, a_n), (b_1, \dots, b_n) \rangle = a_1 b_1 + \dots + a_n b_n$$

for two vectors $(a_1, \dots, a_n), (b_1, \dots, b_n) \in \mathbb{R}^n$. The dot product has a lot of applications in geometry: For example, it can be used to define the length $\|v\|$ of a vector $v \in \mathbb{R}^n$ via

$$\|v\| = \sqrt{\langle v, v \rangle}.$$

The Pythagorean theorem justifies this definition as shown for $n = 2$ in Figure 4.




Figure 4: The length of the vector $v = (x, y)$ is $\sqrt{x^2 + y^2} = \sqrt{\langle v, v \rangle}$.

The dot product also provides us with a way of measuring the distance between two points described by two vectors $v, w \in \mathbb{R}^n$, which turns out to be $\|v - w\|$. If φ denotes the angle between two vectors $v, w \in \mathbb{R}^n$, then it satisfies the equation

$$\langle v, w \rangle = \|v\| \|w\| \cos(\varphi).$$

Last, but not least, the dot product also allows us to determine whether two vectors $v, w \in \mathbb{R}^n$ are perpendicular (or orthogonal) to each other, which is the case if and only if

$$\langle v, w \rangle = 0 .$$

In this section we will discuss the essential properties of this product, which will lead to a definition that works for arbitrary vector spaces. Note that the dot product on \mathbb{R}^n is a map

$$\langle \cdot, \cdot \rangle: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R} .$$

However, even though $\mathbb{R}^n \times \mathbb{R}^n = \mathbb{R}^{2n}$ is a vector space, this map is **not linear** as can be easily checked from the definition. Nevertheless, if we fix $v \in \mathbb{R}^n$, then the map

$$\langle v, \cdot \rangle: \mathbb{R}^n \rightarrow \mathbb{R}$$

does satisfy $\langle v, cw_1 + w_2 \rangle = c\langle v, w_1 \rangle + \langle v, w_2 \rangle$ and hence is a linear map. Therefore we will begin our journey there and study linear transformations from a vector space to its field in the next section.

8.1 Dual spaces

Given two vector spaces V and W over the same field F we have already seen in Def. 2.0.8 that the set $\mathcal{L}(V, W)$ of all linear transformations $T: V \rightarrow W$ is again a vector space over F . A special case of this definition is particularly interesting.

Definition 8.1.1. Let V be a vector space over the field F . The vector space $\mathcal{L}(V, F)$ of linear transformations $\varphi: V \rightarrow F$ is called the **dual space of V** and is denoted by V^* . Its elements are sometimes called **linear forms on V** .

Example 8.1.2. Let $V = \mathbb{R}^3$. The linear transformations $\varphi: \mathbb{R}^3 \rightarrow \mathbb{R}$ and $\psi: \mathbb{R}^3 \rightarrow \mathbb{R}$ given by $\varphi(x, y, z) = x$ and $\psi(x, y, z) = x - y$ are both elements of V^* .

Example 8.1.3. Let $V = C([0, 1])$ be the vector space of continuous functions on the unit interval $[0, 1] \subset \mathbb{R}$. The linear transformation $\varphi: V \rightarrow \mathbb{R}$ given by

$$\varphi(f) = \int_0^1 f(x) dx$$

is an element in V^* .

Let V be a finite dimensional vector space over the field F with $n = \dim(V)$ and let $\beta = \{v_1, \dots, v_n\}$ be a basis for V . Let $\varphi_\beta: V \rightarrow F^n$ be the isomorphism from Thm. 3.0.19, which sends a vector v to the unique vector $(a_1, \dots, a_n) \in F^n$ such that

$$v = a_1 v_1 + \dots + a_n v_n .$$

Let $p_i: F^n \rightarrow F$ be defined by $p_i(a_1, \dots, a_n) = a_i$. Note that p_i is also a linear transformation. For $i \in \{1, \dots, n\}$ we define a map $v_i^*: V \rightarrow F$ as follows

$$v_i^* = p_i \circ \varphi_\beta .$$

i.e. v_i^* maps a vector $v \in V$ to its i th coordinate with respect to the basis β . In particular, observe that

$$v_i^*(v_j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases} . \quad (20)$$

Since each v_i^* is a linear transformation, we have $v_i^* \in V^*$.

Lemma 8.1.4. *Let V be a finite-dimensional vector space over the field F of dimension $n = \dim(V)$ and let $\beta = \{v_1, \dots, v_n\}$ be a basis for V . The set $\beta^* = \{v_1^*, \dots, v_n^*\} \subset V^*$ with v_i^* as defined in the last paragraph is a basis for the vector space V^* .*

Proof. The dimension of $V^* = \mathcal{L}(V, F)$ is equal to n by Cor. 3.0.18. Hence, it suffices to show that β^* is linearly independent. To see this, let $a_1, \dots, a_n \in F$ be scalars with the property that

$$a_1 v_1^* + \dots + a_n v_n^* = 0 .$$

If we apply the left hand side to a basis vector v_j , equation (20) shows that

$$(a_1 v_1^* + \dots + a_n v_n^*)(v_j) = a_1 v_1^*(v_j) + \dots + a_n v_n^*(v_j) = a_j .$$

This is equal to the right hand side of the first equation applied to v_j , but the right hand side is the zero map. Therefore $a_j = 0$ for all $j \in \{1, \dots, n\}$ and β^* is linearly independent. \square

Definition 8.1.5. Let V be a finite-dimensional vector space over the field F of dimension $n = \dim(V)$ and let $\beta = \{v_1, \dots, v_n\}$ be a basis for V . The basis β^* of V^* associated to V and β is called the **basis dual to β** or just the **dual basis** if it is clear from the context which β is meant.

Example 8.1.6. Let $V = \mathbb{R}^3$, let $\alpha = \{e_1, e_2, e_3\}$ be the standard basis and consider the vectors $f_i \in \mathbb{R}^3$ defined by

$$f_1 = (1, 1, 0) \quad , \quad f_2 = (0, 1, 1) \quad , \quad f_3 = (0, 0, 1) .$$

The set $\beta = \{f_1, f_2, f_3\}$ forms another basis for \mathbb{R}^3 . We will use (20) to compute the basis β^* dual to β . By Lemma 8.1.4 each f_i can be written as a linear combination of the vectors e_1^* , e_2^* and e_3^* . With $f_i^* = a_i e_1^* + b_i e_2^* + c_i e_3^*$ we have

$$f_i^*((x, y, z)) = a_i e_1^*((x, y, z)) + b_i e_2^*((x, y, z)) + c_i e_3^*((x, y, z)) = a_i x + b_i y + c_i z .$$

To determine the coefficients $a_i, b_i, c_i \in \mathbb{R}$ we use (20). In the case of f_1^* we obtain, for example

$$\begin{aligned} 1 &= f_1^*(f_1) = f_1^*((1, 1, 0)) = a_1 + b_1 , \\ 0 &= f_1^*(f_2) = f_1^*((0, 1, 1)) = b_1 + c_1 , \\ 0 &= f_1^*(f_3) = f_1^*((0, 0, 1)) = c_1 . \end{aligned}$$

This system of linear equations has a unique solution given by $a_1 = 1, b_1 = 0, c_1 = 0$. Therefore $f_1^*((x, y, z)) = x$.

We repeat these steps for f_2^* and f_3^* . In case of f_2^* the system of linear equations is the following one:

$$\begin{aligned} 0 &= f_2^*(f_1) = f_2^*((1, 1, 0)) = a_2 + b_2 , \\ 1 &= f_2^*(f_2) = f_2^*((0, 1, 1)) = b_2 + c_2 , \\ 0 &= f_2^*(f_3) = f_2^*((0, 0, 1)) = c_2 . \end{aligned}$$

This time the solution is $a_2 = -1, b_2 = 1, c_2 = 0$ and hence $f_2^*((x, y, z)) = -x + y$. The same procedure applied to f_3^* gives

$$\begin{aligned} 0 &= f_3^*(f_1) = f_3^*((1, 1, 0)) = a_3 + b_3 , \\ 0 &= f_3^*(f_2) = f_3^*((0, 1, 1)) = b_3 + c_3 , \\ 1 &= f_3^*(f_3) = f_3^*((0, 0, 1)) = c_3 \end{aligned}$$

and the solution is $a_3 = 1, b_3 = -1, c_3 = 1$. Thus, $f_3^*((x, y, z)) = x - y + z$. In summary we have

$$\begin{aligned} f_1^*((x, y, z)) &= x , \\ f_2^*((x, y, z)) &= -x + y , \\ f_3^*((x, y, z)) &= x - y + z . \end{aligned}$$

Every vector space V over the field F has a dual space V^* . In particular, we can iterate this construction and look at $V^{**} = (V^*)^*$. This is called the **bidual** of V and by definition it is the space of linear forms $\chi: V^* \rightarrow F$. Now note that there is a map

$$\Theta_V: V \rightarrow V^{**} \quad , \quad v \mapsto e_v ,$$

where $e_v: V^* \rightarrow F$ is the evaluation map at $v \in V$. It maps a linear form $\varphi \in V^*$ to the scalar $e_v(\varphi) = \varphi(v)$.

Theorem 8.1.7. *Let V be a finite-dimensional vector space over the field F . Then the map $\Theta_V: V \rightarrow V^{**}$ is an isomorphism.*

Proof. We will first show that Θ_V is in fact a linear transformation. Let $v, w \in V$ and let $c \in F$. Then we have $\Theta_V(cv + w) = e_{cv+w}$ and

$$e_{cv+w}(\varphi) = \varphi(cv + w) = c\varphi(v) + \varphi(w) = ce_v(\varphi) + e_w(\varphi) = (ce_v + e_w)(\varphi)$$

for all $\varphi \in V^*$, which implies that

$$\Theta_V(cv + w) = e_{cv+w} = ce_v + e_w = c\Theta_V(v) + \Theta_V(w) .$$

Now note that $\dim(V) = \dim(V^*) = \dim(V^{**})$ by Cor. 3.0.18. Hence, it suffices to prove that Θ_V is injective, which is equivalent to $\ker(\Theta_V) = \{0\}$. Let $n = \dim(V)$,

let $\beta = \{v_1, \dots, v_n\}$ be a basis for V and denote by $\beta^* = \{v_1^*, \dots, v_n^*\}$ the dual basis for V^* . Let $v \in \ker(\Theta_V)$. This means that $e_v(\varphi) = \varphi(v) = 0$ for all $\varphi \in V^*$. Let $v = a_1v_1 + \dots + a_nv_n$ be the decomposition of v with respect to the basis β . For each $i \in \{1, \dots, n\}$ we have

$$0 = v_i^*(v) = v_i^*(a_1v_1 + \dots + a_nv_n) = a_1v_i^*(v_1) + \dots + a_nv_i^*(v_n) = a_i$$

and therefore $v = 0$. In other words, $\ker(\Theta_V) = \{0\}$ and Θ_V is an isomorphism. \square

Remark 8.1.8. For infinite-dimensional vector spaces V it is no longer true that Θ_V is an isomorphism. Nevertheless, it can still be shown to be injective in this case. Note also that we did not have to make any choices in the construction of $\Theta_V: V \rightarrow V^{**}$. Mathematicians like to call structures that you get “for free”, i.e. without making any choices, **canonical**. Therefore Θ_V is the canonical map between V and V^{**} .

8.2 Bilinear forms and inner products on real vector spaces

Let $V = \mathbb{R}^n$ and consider the dot product discussed in the introduction, i.e. the map defined as follows:

$$\langle (a_1, \dots, a_n), (b_1, \dots, b_n) \rangle = \sum_{i=1}^n a_i b_i. \quad (21)$$

If we keep the first entry fixed and replace the second entry by a linear combination of two vectors, then we see that for $v = (a_1, \dots, a_n)$, $w = (b_1, \dots, b_n)$ and $w' = (b'_1, \dots, b'_n)$ we have

$$\begin{aligned} \langle v, cw + dw' \rangle &= \sum_{i=1}^n a_i (cb_i + db'_i) = c \left(\sum_{i=1}^n a_i b_i \right) + d \left(\sum_{i=1}^n a_i b'_i \right) \\ &= c \langle v, w \rangle + d \langle v, w' \rangle. \end{aligned}$$

The analogous statement is also true if we keep the second entry fixed and replace the first one by a linear combination. Thus, even though $\langle \cdot, \cdot \rangle$ is not a linear transformation, it is still linear if we keep one of the two entries fixed and vary the other one.

Definition 8.2.1. Let V be a vector space over a field F . A **bilinear form** on V is a map $b: V \times V \rightarrow F$ such that

$$\begin{aligned} b(v, cw_1 + dw_2) &= cb(v, w_1) + db(v, w_2), \\ b(cv_1 + dv_2, w) &= cb(v_1, w) + db(v_2, w) \end{aligned}$$

for all $c, d \in F$ and $v, v_1, v_2, w, w_1, w_2 \in V$. A bilinear form $b: V \times V \rightarrow F$ is called **symmetric** if $b(v, w) = b(w, v)$ for all $v, w \in V$.

Remark 8.2.2. Just like linear forms, i.e. the elements $\varphi \in V^*$ are special cases of linear maps $T: V \rightarrow W$, bilinear forms are also instances of a more general notion: Let V, W and X be three vector spaces over the same field F . A **bilinear map**

$$B: V \times W \rightarrow X$$

is a map with the following properties:

$$\begin{aligned} B(v, cw_1 + dw_2) &= cB(v, w_1) + dB(v, w_2) , \\ B(cv_1 + dv_2, w) &= cB(v_1, w) + dB(v_2, w) \end{aligned}$$

for all $c, d \in F$, $v, v_1, v_2 \in V$ and $w, w_1, w_2 \in W$. We will not discuss these maps any further here, but they lead to other important definitions in Linear Algebra, such as the tensor product of two vector spaces.

The dot product on \mathbb{R}^n has the remarkable property that a vector $v \in \mathbb{R}^n$ is completely determined by knowing its products with all other vectors $w \in \mathbb{R}^n$. This property has important consequences in Linear Algebra. Therefore we make the following definition:

Definition 8.2.3. Let V be a finite-dimensional vector space over a field F and let $b: V \times V \rightarrow F$ be a bilinear form on V . We say that b is **non-degenerate** if $b(v, w) = 0$ for all $w \in V$ implies that $v = 0$.

Example 8.2.4. We have already seen above that the dot product $\langle \cdot, \cdot \rangle$ on \mathbb{R}^n defined in (21) is a bilinear form. Let $v = (a_1, \dots, a_n)$, $w = (b_1, \dots, b_n)$. Then we have

$$\langle v, w \rangle = \sum_{i=1}^n a_i b_i = \sum_{i=1}^n b_i a_i = \langle w, v \rangle$$

which shows that the dot product is also symmetric. Let $v = (a_1, \dots, a_n) \in \mathbb{R}^n$ and suppose that $\langle v, w \rangle = 0$ for all $w \in V$. In particular, we have

$$0 = \langle v, v \rangle = \sum_{i=1}^n a_i^2$$

but all summands on the right hand side are non-negative. Hence, we must have $a_i = 0$ for all $i \in \{1, \dots, n\}$ or, in other words, $v = 0$. This means that the dot product is also non-degenerate.

Example 8.2.5. To show that the dot product is non-degenerate we used that $\langle v, v \rangle = 0$ implies $v = 0$ in the last example. This gives rise to the question whether there are non-degenerate bilinear forms with $b(v, v) = 0$ for some vector $v \in V$, but $v \neq 0$. Consider $b: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ defined for vectors $v = (a_0, \dots, a_{n-1})$ and $w = (b_0, \dots, b_{n-1})$ by

$$b(v, w) = -a_0 b_0 + a_1 b_1 + \dots + a_{n-1} b_{n-1} .$$

The vector $v = (1, 1, 0, \dots, 0)$ satisfies $b(v, v) = 0$ and $v \neq 0$. Nevertheless b is non-degenerate. In fact, assume that $v = (a_0, \dots, a_{n-1})$ is a vector with the property that $b(v, w) = 0$ for all $w \in \mathbb{R}^n$. Choosing $w = (0, a_1, \dots, a_{n-1})$ we see that

$$0 = b(v, w) = \sum_{i=1}^{n-1} a_i^2 ,$$

which implies $a_i = 0$ for all $i \in \{1, \dots, n-1\}$. Using $w = (a_0, 0, \dots, 0)$ we get

$$0 = b(v, w) = -a_0^2$$

which can only hold if a_0 vanishes as well. Therefore we get $v = 0$ and b is indeed non-degenerate.

Bilinear forms of this type play a crucial role in physics: Einstein's special theory of relativity takes place in Minkowski space. This is the vector space \mathbb{R}^4 where the first coordinate is interpreted as the time direction and the other three components label spatial directions. It comes equipped with the following bilinear form, which is defined on two vectors $v = (t_1, x_1, y_1, z_1)$ and $w = (t_2, x_2, y_2, z_2)$ by

$$b(v, w) = -c^2 t_1 t_2 + x_1 x_2 + y_1 y_2 + z_1 z_2 ,$$

where c denotes the speed of light. The time units can be chosen in such a way that $c = 1$ holds. Therefore we can and will forget about this additional constant. Points in Minkowski space label events in spacetime and for a vector $v = (t, x, y, z)$ the sign of

$$b(v, v) = -t^2 + x^2 + y^2 + z^2$$

has a nice interpretation: If $b(v, v) > 0$, then the event happening at the origin can not affect the event taking place at the point labelled by $v \in \mathbb{R}^4$ in spacetime, since this would require signals that travel faster than light. These vectors are called *spacelike to the origin*. On the other hand if $b(v, v) < 0$, then the two events can influence each other. In this case the vector is called *timelike*. If $b(v, v) = 0$, the vector is called *lightlike*.




Figure 5: Sketch of Minkowski space in two dimensions, vectors pointing into the red regions are spacelike, those pointing into the blue regions are timelike, vectors pointing along the diagonals are lightlike.

Example 8.2.6. Let $V = \mathbb{R}^3$ and consider the bilinear form $b: V \times V \rightarrow \mathbb{R}$ defined for vectors $v = (a_1, a_2, a_3)$ and $w = (b_1, b_2, b_3)$ by

$$b(v, w) = a_1 b_1 + a_2 b_2 .$$

The vector $v = (0, 0, 1)$ satisfies $b(v, w) = 0 \cdot b_1 + 0 \cdot b_2 = 0$ for all $w \in \mathbb{R}^3$. Since $v \neq 0$, this shows that this bilinear form is degenerate.

Definition 8.2.7. Let F be a field and let $n \in \mathbb{N}$. The **trace** of $A = (a_{ij})_{i,j} \in M_{n \times n}(F)$ is defined by

$$\operatorname{tr}(A) = \sum_{i=1}^n a_{ii} .$$

It is a straightforward to check that $\operatorname{tr}: M_{n \times n}(F) \rightarrow F$ is a linear transformation.

Exercise 8.2.8. Let $A, B \in M_{n \times n}(F)$. Prove that the trace satisfies the equation

$$\operatorname{tr}(AB) = \operatorname{tr}(BA) .$$

Example 8.2.9. Let $V = M_{n \times n}(\mathbb{R})$. Consider the map $b: M_{n \times n}(\mathbb{R}) \times M_{n \times n}(\mathbb{R}) \rightarrow \mathbb{R}$ defined by

$$b(A, B) = \operatorname{tr}(A^t B) .$$

Let $B, B' \in M_{n \times n}(F)$ and $c, c' \in F$. Then we have

$$b(A, cB + c'B') = \operatorname{tr}(A^t(cB + c'B')) = \operatorname{tr}(cA^t B + c'A^t B') = c \operatorname{tr}(A^t B) + c' \operatorname{tr}(A^t B')$$

using the linearity of the trace, i.e. b is linear in the second argument. A similar argument shows that it is also linear in the first argument. Therefore b is a bilinear form on $M_{n \times n}(\mathbb{R})$. Since the diagonal entries of a matrix $X \in M_{n \times n}(\mathbb{R})$ and its transpose X^t agree, we have $\operatorname{tr}(X) = \operatorname{tr}(X^t)$. Therefore

$$b(A, B) = \operatorname{tr}(A^t B) = \operatorname{tr}((A^t B)^t) = \operatorname{tr}(B^t A) = b(B, A) ,$$

i.e. b is symmetric. Suppose that $A \in M_{n \times n}(\mathbb{R})$ satisfies $b(A, B) = 0$ for all $B \in M_{n \times n}(\mathbb{R})$. The case $B = A$ leads to

$$0 = b(A, A) = \operatorname{tr}(A^t A) = \sum_{i,j=1}^n a_{ij}^2 ,$$

which can only be true if all $a_{ij} = 0$, since $a_{ij}^2 \geq 0$. It follows that b is non-degenerate as well.

Let V be a finite-dimensional vector space over the field F and let $b: V \times V \rightarrow F$ be a bilinear form. Then it gives rise to two linear transformations

$$\begin{aligned} b_L: V &\rightarrow V^* & , & & v &\mapsto b(v, \cdot) , \\ b_R: V &\rightarrow V^* & , & & v &\mapsto b(\cdot, v) , \end{aligned}$$

where $b(v, \cdot): V \rightarrow F$ is given by $w \mapsto b(v, w)$ and $b(\cdot, v)$ is defined analogously. As we will see below, if b is non-degenerate, then knowing one of the two maps will allow us to reconstruct b from it. We start by showing that b_L and b_R tell us whether b is non-degenerate.

Lemma 8.2.10. *Let V be a finite-dimensional vector space over the field F and let $b: V \times V \rightarrow F$ be a bilinear form. Then the following are equivalent:*

- 1) b is non-degenerate,
- 2) b_L is an isomorphism,
- 3) if $b(v, w) = 0$ for all $v \in V$, then $w = 0$.
- 4) b_R is an isomorphism,

Proof. Assume that b is non-degenerate and let $v \in \ker(b_L)$. This means that $b_L(v) \in V^*$ is the zero map and therefore

$$0 = b_L(v)(w) = b(v, w)$$

for all $w \in V$, which implies $v = 0$. Hence, b_L is injective. But since $\dim(V) = \dim(V^*)$, it must be an isomorphism. This proves 1) \Rightarrow 2). Likewise, if b_L is an isomorphism, then it is injective. Let $v \in V$ be a vector with $b(v, w) = 0$ for all $w \in W$. Then we must have $b_L(v) = 0$ and hence $v = 0$. This shows 2) \Rightarrow 1). The equivalence of 3) and 4) can be proven in a similar way by replacing b_L by b_R .

It remains to be seen that 2) and 3) are equivalent. Suppose that b_L is an isomorphism, let $n = \dim(V)$ and choose a basis $\beta = \{v_1, \dots, v_n\}$ for V . We claim that the set $\{b_L(v_1), \dots, b_L(v_n)\}$ is linearly independent. To see this, let $a_1, \dots, a_n \in F$ be such that

$$a_1 b_L(v_1) + \dots + a_n b_L(v_n) = b_L(a_1 v_1 + \dots + a_n v_n) = 0 .$$

But then $a_1 v_1 + \dots + a_n v_n \in \ker(b_L) = \{0\}$, which gives $a_1 = \dots = a_n = 0$, because β is a basis. Since $\dim(V) = \dim(V^*)$, it follows that $\{b_L(v_1), \dots, b_L(v_n)\}$ is a basis for V^* . Let $\{\chi_1, \dots, \chi_n\} \subset V^{**}$ be its dual basis and observe that it satisfies

$$\chi_j(b_L(v_i)) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases} .$$

By Thm. 8.1.7 the canonical map $\Theta_V: V \rightarrow V^{**}$ is an isomorphism. Let $w_i = \Theta_V^{-1}(\chi_i)$ and note that a similar argument as the one above shows that $\gamma = \{w_1, \dots, w_n\}$ is another basis for V . By construction we have

$$b(v_i, w_j) = b_L(v_i)(w_j) = \Theta_V(w_j)(b_L(v_i)) = \chi_j(b_L(v_i)) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases} .$$

Let $w = c_1 w_1 + \dots + c_n w_n$ be such that $b(v, w) = 0$ for all $v \in V$. By the observations above

$$0 = b(v_i, w) = \sum_{j=1}^n c_j b(v_i, w_j) = c_i$$

for all $i \in \{1, \dots, n\}$. This proves 2) \Rightarrow 3). Exchanging the roles of b_L and b_R the statement 4) \Rightarrow 1) follows in a similar way. \square

Theorem 8.2.11. *Let V be a finite-dimensional vector space over the field F . There is a 1 : 1-correspondence between the set of all non-degenerate bilinear forms $b: V \times V \rightarrow F$ and the set of isomorphisms $\psi: V \rightarrow V^*$. This correspondence maps b to b_L .*

Proof. We know from Lemma 8.2.10 that $b_L: V \rightarrow V^*$ is an isomorphism if and only if b is non-degenerate. Therefore the map from forms to isomorphisms is well-defined. To construct a map in the other direction let $\psi: V \rightarrow V^*$ be an arbitrary isomorphism and define

$$b_\psi(v, w) = \psi(v)(w) .$$

By the linearity of ψ this is linear in the first entry, and since it maps to the vector space of linear forms on V , it is also linear in the second entry. Therefore b_ψ is a bilinear form. Now suppose that we have $v \in V$ with the property that $b_\psi(v, w) = 0$ for all $w \in V$. This means that $\psi(v) = 0 \in V^*$ and hence that $v \in \ker(\psi) = \{0\}$. But then $v = 0$ and b_ψ is non-degenerate.

If we start with a bilinear form b and with $\psi = b_L$, then the above construction retrieves b . If we start with an isomorphism $\psi: V \rightarrow V^*$, then the map b_L associated to b_ψ agrees with ψ . This shows that the two constructions are inverse to each other. \square

As mentioned in the introduction the dot product on \mathbb{R}^n allows us to determine whether two vector are orthogonal to each other. In fact, the notion of orthogonality also makes sense for all non-degenerate symmetric bilinear forms:

Definition 8.2.12. Let V be a finite-dimensional vector space over a field F and let $b: V \times V \rightarrow F$ be a non-degenerate symmetric bilinear form on V . We say that $v \in V$ is **orthogonal to** $w \in V$ if

$$b(v, w) = 0 .$$

Let $U \subset V$ be a subset. The **orthogonal complement** U^\perp of U in V is defined as

$$U^\perp = \{v \in V \mid b(v, u) = 0 \text{ for all } u \in U\} ,$$

i.e. it consists of all vectors $v \in V$ that are orthogonal to all vectors in U .

Remark 8.2.13. If b is non-degenerate, then $V^\perp = \{0\}$. If we drop the non-degeneracy condition and define orthogonal complements in the same way as above, then there might be vectors that are orthogonal to the whole space without being zero.

To define orthogonal complements for bilinear forms that are not symmetric we would have to distinguish between left and right complement, i.e. the two sets

$$\begin{aligned} U^\perp &= \{v \in V \mid b(v, u) = 0 \text{ for all } u \in U\} , \\ {}^\perp U &= \{v \in V \mid b(u, v) = 0 \text{ for all } u \in U\} \end{aligned}$$

would not necessarily agree. However, since we will not need these constructions, we will not dig deeper into the technicalities here.

In the next three lemmata we prove some of the basic properties of the orthogonal complement.

Lemma 8.2.14. *Let V be a finite-dimensional vector space over the field F and let $b: V \times V \rightarrow F$ be a non-degenerate symmetric bilinear form on V . The orthogonal complement U^\perp of any subset $U \subset V$ is a linear subspace of V .*

Proof. Let $v, w \in U^\perp$. By definition we have $b(v, u) = 0$ and $b(w, u) = 0$ for all $u \in U$. Using the linearity of b in its first argument we get

$$b(v + w, u) = b(v, u) + b(w, u) = 0$$

for all $u \in U$. Let $v \in U^\perp$ and $a \in F$. Using linearity again we have

$$b(av, u) = a b(v, u) = 0$$

for all $u \in U$. This shows that U^\perp is a subspace of V . □

Lemma 8.2.15. *Let V be a finite-dimensional vector space over the field F and let $b: V \times V \rightarrow F$ be a non-degenerate symmetric bilinear form on V . Let $U \subset V$ be any subset. Then we have*

$$U^\perp = \text{span}(U)^\perp .$$

Proof. Suppose that $U_1 \subset U_2$. By definition U_2^\perp is the set of vectors orthogonal to all $u \in U_2$. Since $U_1 \subset U_2$ the vectors in U_2^\perp are in particular orthogonal to vectors in U_1 and therefore $U_2^\perp \subset U_1^\perp$. Since $U \subset \text{span}(U)$, we have

$$\text{span}(U)^\perp \subset U^\perp .$$

Let $v \in U^\perp$ and let $u = a_1 u_1 + \dots + a_k u_k \in \text{span}(U)$ be a linear combination of vectors $u_i \in U$ with $a_i \in F$ for $i \in \{1, \dots, k\}$. Linearity of b in the second argument implies

$$b(v, u) = b(v, a_1 u_1 + \dots + a_k u_k) = a_1 b(v, u_1) + \dots + a_k b(v, u_k) = 0 .$$

This means $v \in \text{span}(U)^\perp$ and therefore $U^\perp \subset \text{span}(U)^\perp$. □

Lemma 8.2.16. *Let V be a finite-dimensional vector space over the field F and let $b: V \times V \rightarrow F$ be a non-degenerate symmetric bilinear form on V . For any linear subspace $W \subset V$ we have*

a) $\dim(W) + \dim(W^\perp) = \dim(V)$,

b) $W = W^{\perp\perp}$.

Proof. Let $r = \dim(W)$, $n = \dim(V)$ and choose a basis $\beta = \{w_1, \dots, w_r\}$ for W . Let $\theta: V \rightarrow F^r$ be the linear transformation given by

$$\theta(v) = (b(w_1, v), \dots, b(w_r, v)) .$$

By Lem. 8.2.15 we have

$$\ker(\theta) = \{w_1, \dots, w_r\}^\perp = \text{span}(\{w_1, \dots, w_r\})^\perp = W^\perp .$$

If we can show that θ is also surjective, then a) will follow from the rank-nullity theorem. Just as in the proof of Lem. 8.2.10, the subset $\{b_L(w_1), \dots, b_L(w_r)\} \subset V^*$ is linearly independent. Extend it to a basis γ of V^* and let $\gamma^* = \{\chi_1, \dots, \chi_n\} \subset V^{**}$ be its dual basis. Now define $w'_i = \Theta_V^{-1}(\chi_i)$ for all $i \in \{1, \dots, r\} \subset V$. Then we have

$$b(w_j, w'_k) = b_L(w_j)(w'_k) = \chi_k(b_L(w_j)) = \begin{cases} 1 & \text{if } j = k \\ 0 & \text{else} \end{cases} .$$

In particular, $\theta(w'_k) = e_k$. Let $(a_1, \dots, a_r) \in F^r$ and let $w' = a_1 w'_1 + \dots + a_r w'_r \in V$. If we apply θ to this vector we see that

$$\theta(w') = a_1 \theta(w'_1) + \dots + a_r \theta(w'_r) = (a_1, \dots, a_r) .$$

Hence, θ is surjective and by the rank-nullity theorem

$$\dim(V) = \dim(\text{Im}(\theta)) + \dim(\ker(\theta)) = r + \dim(W^\perp) = \dim(W) + \dim(W^\perp) ,$$

which finishes the proof of a). To show b) let $w \in W$ and note that $b(w, w') = 0$ for all $w' \in W^\perp$. Therefore $w \in W^{\perp\perp}$, which shows that $W \subset W^{\perp\perp}$ is a linear subspace. From a) we deduce the following two equations

$$\begin{aligned} \dim(W) + \dim(W^\perp) &= \dim(V) , \\ \dim(W^\perp) + \dim(W^{\perp\perp}) &= \dim(V) \end{aligned}$$

and subtracting the second from the first we see that $\dim(W) = \dim(W^{\perp\perp})$. This statement together with $W \subset W^{\perp\perp}$ implies that $W = W^{\perp\perp}$ proving b). \square

Example 8.2.17. In this section we will discuss some geometric examples of orthogonal complements of subsets in \mathbb{R}^3 equipped with the dot product as the non-degenerate symmetric bilinear form, i.e. $b(v, w) = \langle v, w \rangle$.

Fix a vector $v \in \mathbb{R}^3$ and consider the orthogonal complement $\{v\}^\perp$. This coincides with the two-dimensional plane through the origin, which is perpendicular to v .




Figure 6: The plane orthogonal to the vector $(1, 0, 0)$ is the y - z -plane.

Similarly, if $v_1, v_2 \in \mathbb{R}^3$ are two linearly independent vectors, then the vector space $W = \{v_1, v_2\}^\perp$ is the line through the origin, which is perpendicular to the plane spanned by v_1 and v_2 .

Fix a vector $v \in \mathbb{R}^3$. By Lem. 8.2.15 and Lem. 8.2.16 we have

$$L = \{v\}^{\perp\perp} = (\{v\}^\perp)^\perp = \text{span}(\{v\})^{\perp\perp} = \text{span}(\{v\}) .$$

Therefore the vector space L agrees with the line spanned by v . More generally, if $U \subset \mathbb{R}^3$ is any subset, then by the same reasoning as above:

$$U^{\perp\perp} = (U^\perp)^\perp = \text{span}(U)^{\perp\perp} = \text{span}(U) .$$

For example, for a linearly independent subset $U = \{v_1, v_2\} \subset \mathbb{R}^3$ the vector space $U^{\perp\perp}$ is the plane spanned by the two vectors.

Example 8.2.18. Let $V = \mathbb{R}^2$ and consider the non-degenerate symmetric bilinear form $b: V \times V \rightarrow \mathbb{R}$ given by

$$b((t_1, x_1), (t_2, x_2)) = -t_1 t_2 + x_1 x_2 .$$

This is the two-dimensional version of Minkowski space as discussed in Example 8.2.5. The orthogonal complement of the vector $v = (2, 1)$ is the vector space

$$\{v\}^\perp = \{a(1, 2) \in \mathbb{R}^2 \mid a \in \mathbb{R}\} .$$

The line $L = \{a(1, 1) \in \mathbb{R}^2 \mid a \in \mathbb{R}\}$ is its own orthogonal complement with respect to the bilinear form b , i.e. $L^\perp = L$.




Figure 7: The line ℓ is the orthogonal complement of the vector v with respect to the bilinear form given above.

As pointed out above there are vectors $v \in \mathbb{R}^4$ with the property that $b(v, v) < 0$ for the bilinear form from Example 8.2.5. In contrast to this, such vectors do not exist for the form from Example 8.2.4. We capture this difference in the following definition:

Definition 8.2.19. Let V be a finite-dimensional vector space over \mathbb{R} . A bilinear form $b: V \times V \rightarrow \mathbb{R}$ on V is called **positive definite** if

$$b(v, v) > 0$$

for all $v \neq 0$. A positive definite, symmetric bilinear form on V is also called an **inner product** on V . In this case the pair (V, b) is called an **inner product space**.

Suppose that $b: V \times V \rightarrow \mathbb{R}$ is a positive definite bilinear form and let $v \in V$ be a vector such that $b(v, w) = 0$ for all $w \in V$. Then we must also have $b(v, v) = 0$ and positive definiteness implies that $v = 0$. Therefore any positive definite bilinear form is non-degenerate.

Whenever a vector space V comes equipped with an inner product b , we can define the length of a vector $v \in V$ by

$$\|v\| = \sqrt{b(v, v)},$$

which is a non-negative real number, since $b(v, v) \geq 0$.

Lemma 8.2.20. Let V be a finite-dimensional vector space over \mathbb{R} . A map $b: V \times V \rightarrow \mathbb{R}$ is an inner product on V if and only if it satisfies the following conditions:

a) It is linear in the second argument, i.e. for all $v, w_1, w_2 \in V$ and $c \in F$ we have

$$b(v, cw_1 + w_2) = cb(v, w_1) + b(v, w_2).$$

b) It is symmetric, i.e. $b(v, w) = b(w, v)$ for all $v, w \in V$.

c) It satisfies $b(v, v) \geq 0$ for all $v \in V$ and

$$b(v, v) = 0 \quad \Leftrightarrow \quad v = 0.$$

Proof. If b is a positive definite, symmetric bilinear form, then it satisfies a), b) and c). Therefore we will only prove the converse. Suppose that b satisfies a), b) and c). By a) the map b is linear in the second argument. Let $v_1, v_2, w \in V$ and let $c \in F$. Using b) we get

$$b(cv_1 + v_2, w) = b(w, cv_1 + v_2) = cb(w, v_1) + b(w, v_2) = cb(v_1, w) + b(v_2, w).$$

Therefore b is a bilinear form and by b) it is symmetric. By c) it is also positive definite and therefore defines an inner product on V . \square

Example 8.2.21. The dot product defined in Example 8.2.4 is positive definite, since $\langle v, v \rangle > 0$ for all $v \neq 0$ and therefore an inner product in the sense of Def. 8.2.19.

Example 8.2.22. Let $V = M_n(\mathbb{R})$ and let $b: M_n(\mathbb{R}) \times M_n(\mathbb{R}) \rightarrow \mathbb{R}$ be the bilinear form given by

$$b(A, B) = \text{tr}(A^t B)$$

as in Example 8.2.9. We have already seen there that for $A = (a_{ij})_{i,j=1}^n$ we have

$$b(A, A) = \sum_{i,j=1}^n a_{ij}^2$$

Therefore $b(A, A) > 0$ if $A \neq 0$ and b is positive definite. Hence, it defines an inner product on $M_n(\mathbb{R})$.

Example 8.2.23. Let $V = \mathbb{R}^2$ and let $b: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}$ be the map

$$b((a_1, a_2), (b_1, b_2)) = 2a_1b_1 - a_1b_2 - a_2b_1 + 2a_2b_2 .$$

Let $v = (a_1, a_2), w_1 = (b_1, b_2), w_2 = (b'_1, b'_2)$, let $c \in \mathbb{R}$. Then we have

$$\begin{aligned} b(v, cw_1 + w_2) &= 2a_1(cb_1 + b'_1) - a_1(cb_2 + b'_2) - a_2(cb_1 + b'_1) + 2a_2(cb_2 + b'_2) \\ &= c(2a_1b_1 - a_1b_2 - a_2b_1 + 2a_2b_2) + 2a_1b'_1 - a_1b'_2 - a_2b'_1 + 2a_2b'_2 \\ &= cb(v, w_1) + b(v, w_2) . \end{aligned}$$

Therefore b satisfies condition a) in Lem. 8.2.20. Now let $v = (a_1, a_2)$ and $w = (b_1, b_2)$ and note that

$$b(v, w) = 2a_1b_1 - a_1b_2 - a_2b_1 + 2a_2b_2 = 2b_1a_1 - b_1a_2 - b_2a_1 + 2b_2a_2 = b(w, v) .$$

Hence, b also satisfies condition b) in Lem. 8.2.20. To see that condition c) also holds for b observe that for $v = (a_1, a_2)$ we have

$$b(v, v) = 2a_1^2 - 2a_1a_2 + 2a_2^2 = a_1^2 + (a_1 - a_2)^2 + a_2^2 .$$

This is a sum of non-negative real numbers. Hence, we have that $b(v, v) = 0$ if and only if $a_1 = a_2 = 0$, i.e. if and only if $v = 0$. By Lem. 8.2.20 the bilinear form b defines an inner product on $V = \mathbb{R}^2$.

8.3 Hermitian forms and inner products on complex vector spaces

Observe that bilinear forms on vector spaces over \mathbb{C} do not behave in quite the same way as those on vector spaces over \mathbb{R} . For example, if we want the length of a vector to be a positive real number, then the definition

$$\|v\| = \sqrt{b(v, v)}$$

does not work, since $b(v, v) \in \mathbb{C}$ might not be a real number. To compensate this we need to modify the definition of bilinear form slightly, which leads to the notion of hermitian forms:

Definition 8.3.1. Let V be a vector space over \mathbb{C} . A **hermitian form** on V is a map $b: V \times V \rightarrow \mathbb{C}$ with the following properties:

a) For all $v, w_1, w_2 \in V$ and $c \in \mathbb{C}$ we have

$$b(v, cw_1 + w_2) = cb(v, w_1) + b(v, w_2) ,$$

i.e. b is linear in the second argument.

b) For all $v, w \in V$ the map b satisfies

$$b(v, w) = \overline{b(w, v)} .$$

Remark 8.3.2. Let $b: V \times V \rightarrow \mathbb{C}$ be a hermitian form on V , let $v_1, v_2, w \in V$ and let $c \in \mathbb{C}$. The properties a) and b) in Def. 8.3.1 imply that

$$\begin{aligned} b(cv_1 + v_2, w) &= \overline{b(w, cv_1 + v_2)} = \overline{cb(w, v_1) + b(w, v_2)} \\ &= \overline{c} \overline{b(w, v_1)} + \overline{b(w, v_2)} = \overline{c} b(v_1, w) + b(v_2, w) . \end{aligned}$$

Therefore b is not linear in the first argument, but antilinear (a map $T: V \rightarrow W$ is antilinear if $T(cv + w) = \overline{c}T(v) + T(w)$). We also say that b is a sesquilinear form.

Replacing bilinear forms by hermitian ones we can now restate all definitions made in the last section:

Definition 8.3.3. Let V be a finite-dimensional vector space over \mathbb{C} . Let $b: V \times V \rightarrow \mathbb{C}$ be a hermitian form on V .

a) We say that b is **non-degenerate** if $b(v, w) = 0$ for all $w \in V$ implies that $v = 0$.

b) We call b **positive definite** if $b(v, v) > 0$ for all $v \neq 0$. Note that this definition makes sense, since $b(v, v) = \overline{b(v, v)}$ implies that $b(v, v) \in \mathbb{R}$.

A positive definite hermitian form on V is called an **inner product on V** .

Just as in the case of real vector spaces, a positive definite hermitian form is automatically non-degenerate.

Example 8.3.4. Let $V = \mathbb{C}^n$ and define $\langle \cdot, \cdot \rangle: V \times V \rightarrow \mathbb{C}$ on vectors of the form $v = (a_1, \dots, a_n)$ and $w = (b_1, \dots, b_n)$ by

$$\langle v, w \rangle = \sum_{i=1}^n \overline{a_i} b_i .$$

The proof that this is linear in the second argument is similar to the one in Example 8.2.4 and we omit it. Note that

$$\langle v, w \rangle = \sum_{i=1}^n \overline{a_i} b_i = \overline{\sum_{i=1}^n b_i a_i} = \overline{\langle w, v \rangle}$$

shows that $\langle \cdot, \cdot \rangle$ is a hermitian form on \mathbb{C}^n . We have

$$\langle v, v \rangle = \sum_{i=1}^n \overline{a_i} a_i = \sum_{i=1}^n |a_i|^2 .$$

Therefore $\langle \cdot, \cdot \rangle$ is positive definite and defines an inner product on \mathbb{C}^n , which is often called the **standard inner product**.

Example 8.3.5. In this example we will discuss the hermitian counterpart of the symmetric bilinear form from Ex. 8.2.9. We need the following definition:

Definition 8.3.6. Let $A \in M_{n \times n}(\mathbb{C})$. The **adjoint matrix** A^* of A is defined to be

$$A^* = \overline{A^t} ,$$

i.e. if A has entries a_{ij} , then the entry in the i th row and the j th column of A^* is $\overline{a_{ji}}$.

If $X \in M_{n \times n}(\mathbb{C})$ is a matrix with entries x_{ij} , then the diagonal entries of X^* are $\overline{x_{ii}}$. In particular, we have $\text{tr}(X^*) = \overline{\text{tr}(X)}$. Now let $b: M_{n \times n}(\mathbb{C}) \times M_{n \times n}(\mathbb{C}) \rightarrow \mathbb{C}$ be given by

$$b(A, B) = \text{tr}(A^* B) .$$

By a similar argument as the one given in Ex. 8.2.9 this form is linear in the second argument. But by our observation about adjoint matrices and traces we also have

$$b(A, B) = \text{tr}(A^* B) = \overline{\text{tr}((A^* B)^*)} = \overline{\text{tr}(B^* A)} = \overline{b(B, A)} ,$$

which shows that b is a hermitian form. We leave the proof that this is positive definite as an exercise.

Exercise 8.3.7. Prove that the hermitian form from Example 8.3.5 is positive definite and therefore defines an inner product on $M_{n \times n}(\mathbb{C})$.

8.4 Orthonormal sets and orthonormal bases

Definition 8.4.1. Let V be a finite-dimensional vector space over $F \in \{\mathbb{R}, \mathbb{C}\}$ and let $b: V \times V \rightarrow F$ be an inner product on V . A finite set of vectors $S = \{v_1, \dots, v_k\}$ is called an **orthonormal set** if for any two indices $i, j \in \{1, \dots, k\}$ we have

$$b(v_i, v_j) = \begin{cases} 1 & \text{if } i = j , \\ 0 & \text{else} \end{cases}$$

Let basis β for V that is also an orthonormal set is called an **orthonormal basis**.

Observe that an orthonormal basis for a finite-dimensional inner product space consists of vectors $\beta = \{v_1, \dots, v_n\}$ that are pairwise orthogonal and satisfy

$$\|v_i\| = \sqrt{b(v_i, v_i)} = 1 ,$$

i.e. the length of each vector is equal to 1. We will see in the next theorem that orthonormal bases for vector spaces are quite common. In fact, we can turn any basis into an orthonormal one.

Theorem 8.4.2. Let V be a finite-dimensional vector space over $F \in \{\mathbb{R}, \mathbb{C}\}$ equipped with an inner product $b: V \times V \rightarrow F$. Let $n = \dim(V)$ and suppose that $\beta = \{v_1, \dots, v_n\}$ is a basis for V . Then there exists an orthonormal basis $\gamma = \{f_1, \dots, f_n\}$ for V with the property that $\text{span}(\{f_1, \dots, f_k\}) = \text{span}(\{v_1, \dots, v_k\})$ for all $1 \leq k \leq n$.

Proof. We will prove the statement by induction over $n = \dim(V)$. For $n = 1$ the basis β contains one element $v_1 \in V$. Since $v_1 \neq 0$, we have $b(v_1, v_1) > 0$ by positive definiteness of b . Let

$$f_1 = \frac{v_1}{\|v_1\|}$$

satisfies $b(f_1, f_1) = 1$ and $\text{span}(\{f_1\}) = \text{span}(\{v_1\})$. Now suppose that $\{f_1, \dots, f_{m-1}\}$ exists and has the properties stated in the theorem. Define

$$f'_m = v_m - \sum_{i=1}^{m-1} b(f_i, v_m) f_i$$

For each $j \in \{1, \dots, m-1\}$ we have

$$\begin{aligned} b(f_j, f'_m) &= b(f_j, v_m) - \sum_{i=1}^{m-1} b(f_j, b(f_i, v_m) f_i) = b(f_j, v_m) - \sum_{i=1}^{m-1} b(f_i, v_m) b(f_j, f_i) \\ &= b(f_j, v_m) - b(f_j, v_m) = 0. \end{aligned}$$

If $f'_m = 0$, then $v_m = \sum_{i=1}^{m-1} b(f_i, v_m) f_i$ and therefore $v_m \in \text{span}(\{f_1, \dots, f_{m-1}\})$. But since

$$\text{span}(\{f_1, \dots, f_{m-1}\}) = \text{span}(\{v_1, \dots, v_{m-1}\})$$

by the induction hypothesis, this would mean that $\{v_1, \dots, v_m\}$ is linearly dependent, which is a contradiction. Hence, we must have $f'_m \neq 0$ and $b(f'_m, f'_m) > 0$. Let

$$f_m = \frac{f'_m}{\|f'_m\|},$$

where $\|f'_m\| = \sqrt{b(f'_m, f'_m)}$. For each $j \in \{1, \dots, m-1\}$ we still get

$$b(f_j, f_m) = \frac{1}{\|f'_m\|} b(f_j, f'_m) = 0.$$

To see that $\{f_1, \dots, f_m\}$ is linearly independent let $a_1, \dots, a_m \in F$ satisfy the equation

$$a_1 f_1 + \dots + a_m f_m = 0. \quad (22)$$

If we apply $b(f_j, \cdot)$ to the left hand side we get

$$b(f_j, a_1 f_1 + \dots + a_m f_m) = a_1 b(f_j, f_1) + \dots + a_m b(f_j, f_m) = a_j.$$

Comparing this to the right hand side we have $a_j = 0$ for all $j \in \{1, \dots, m-1\}$. Hence, equation (22) turns into $a_m f_m = 0$. But since $f_m \neq 0$, we have $a_m = 0$ as well.

It remains to be shown that $\text{span}(\{f_1, \dots, f_m\}) = \text{span}(\{v_1, \dots, v_m\})$: First note that

$$f_m \in \text{span}(\{f_1, \dots, f_{m-1}, v_m\}) .$$

By induction we know that $\text{span}(\{f_1, \dots, f_{m-1}\}) = \text{span}(\{v_1, \dots, v_{m-1}\})$, which gives

$$\text{span}(\{f_1, \dots, f_m\}) \subset \text{span}(\{v_1, \dots, v_m\}) .$$

Resorting the terms in the definition f'_m we obtain $v_m = \|f'_m\| f_m + \sum_{i=1}^{m-1} b(f_i, v_m) f_i$ and therefore $v_m \in \text{span}(\{f_1, \dots, f_m\})$. Using $\text{span}(\{f_1, \dots, f_{m-1}\}) = \text{span}(\{v_1, \dots, v_{m-1}\})$ once more, we have

$$\text{span}(\{v_1, \dots, v_m\}) \subset \text{span}(\{f_1, \dots, f_m\}) . \quad \square$$

Remark 8.4.3. The procedure described in the proof of the last theorem is called the **Gram-Schmidt orthonormalisation process**.

Example 8.4.4. Consider the vector space $V = \mathbb{R}^3$ equipped with the dot product, i.e. with $b(v, w) = \langle v, w \rangle$. The set $\beta = \{v_1, v_2, v_3\}$ with $v_1 = (1, 0, -1)$, $v_2 = (2, -1, 0)$ and $v_3 = (1, 2, 1)$ is a basis for V . We will turn it into an orthonormal basis using Gram-Schmidt orthonormalisation. The first step is to compute the length of v_1 . We have

$$\|v_1\| = \sqrt{2} .$$

Hence, we have $f_1 = \frac{1}{\sqrt{2}}v_1$. Now we can compute f'_2 as follows

$$f'_2 = v_2 - \langle f_1, v_2 \rangle f_1 .$$

The inner product is $\langle f_1, v_2 \rangle = \frac{1}{\sqrt{2}}\langle v_1, v_2 \rangle = \sqrt{2}$. Hence, we have

$$f'_2 = v_2 - \sqrt{2}f_1 = v_2 - v_1 = (1, -1, 1)$$

and with $\langle f'_2, f'_2 \rangle = 3$ we have

$$f_2 = \frac{1}{\sqrt{3}} f'_2 = \frac{1}{\sqrt{3}} (1, -1, 1) .$$

To determine f'_3 we need the inner products $\langle f_1, v_3 \rangle = \frac{1}{\sqrt{2}}\langle v_1, v_3 \rangle = 0$ and $\langle f_2, v_3 \rangle = 0$. Therefore

$$f'_3 = v_3 - \langle f_1, v_3 \rangle f_1 - \langle f_2, v_3 \rangle f_2 = v_3 .$$

Hence, the vector f_3 is

$$f_3 = \frac{1}{\sqrt{6}} v_3 .$$

The set $\beta' = \{f_1, f_2, f_3\}$ is an orthonormal basis for \mathbb{R}^3 .

The terms in the construction of the orthonormal basis have a geometric meaning: Consider two vectors $w, v \in \mathbb{R}^2$ as shown in Fig. 8 and define

$$w' = \frac{\langle w, v \rangle}{\langle v, v \rangle} v$$

This vector has two properties:

- a) $w' \in \text{span}(\{v\})$,
- b) $(w - w') \perp v$.

We say that w' is the orthogonal projection of w onto v .




Figure 8: The orthogonal projection of w onto v .

There is a generalisation of the notion of orthogonal projection to inner product spaces, which works as follows: Suppose that (V, b) is a finite-dimensional inner product space. Let $E \subset V$ be a subspace and let $\{f_1, \dots, f_r\}$ be an orthonormal basis for E . For any vector $w \in V$ the vector

$$w' = \sum_{i=1}^r \frac{b(f_i, w)}{b(f_i, f_i)} f_i = \sum_{i=1}^r b(f_i, w) f_i$$

satisfies $w' \in \text{span}\{f_1, \dots, f_r\} = E$ and for any $j \in \{1, \dots, r\}$ we have

$$\begin{aligned} b(f_j, w - w') &= b(f_j, w) - \sum_{i=1}^r b(f_j, b(f_i, w) f_i) = b(f_j, w) - \sum_{i=1}^r b(f_i, w) b(f_j, f_i) \\ &= b(f_j, w) - b(f_j, w) = 0. \end{aligned}$$

Therefore $(w - w') \in E^\perp$. The vector $w' \in E$ is called the orthogonal projection of $w \in V$ onto the subspace E .

Hence, the first step in the Gram-Schmidt orthonormalisation process can be seen as subtracting from the given vector v_m the orthogonal projection of v_m onto the subspace $E_{m-1} = \text{span}\{f_1, \dots, f_{m-1}\}$.

8.5 Adjoint and Isometries

Consider $V = \mathbb{R}^3$ as a vector space over \mathbb{R} equipped with the dot product $\langle \cdot, \cdot \rangle$. There are many natural linear transformations $T: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that preserve the length of all

vectors. For example, the rotations around a fixed axis have this property. So does the reflection about a plane in \mathbb{R}^3 . In this section we will study the analogous maps on arbitrary finite-dimensional vector spaces over \mathbb{R} or \mathbb{C} equipped with arbitrary inner products. In this context these linear maps are called *isometries*.

Another interesting class of linear transformations are the so-called *self-adjoint* ones. It turns out that a self-adjoint linear map $T: V \rightarrow V$ is always diagonalisable and has real eigenvalues (even though V might be a complex vector space with a complex-valued inner product). These operators play a central role in quantum mechanics. To define them we first need to introduce adjoint operators.

Theorem 8.5.1. *Let V and W be finite-dimensional vector spaces over \mathbb{R} , let*

$$\begin{aligned} b_V &: V \times V \rightarrow \mathbb{R} , \\ b_W &: W \times W \rightarrow \mathbb{R} \end{aligned}$$

be inner products on V and W , respectively. Let $T: V \rightarrow W$ be a linear map. Then there is a unique linear map $T^: W \rightarrow V$ with the property that*

$$b_W(w, T(v)) = b_V(T^*(w), v) \quad (23)$$

for all $v \in V$ and $w \in W$. Moreover, if $\beta = \{v_1, \dots, v_n\}$ is an orthonormal basis for V , $\gamma = \{w_1, \dots, w_m\}$ is an orthonormal basis for W , then we have

$$[T^*]_{\gamma}^{\beta} = \left([T]_{\beta}^{\gamma} \right)^t ,$$

i.e. the matrix representation of T^ is the transpose of the matrix representation of T .*

Proof. Since b_V and b_W are inner products, they are in particular non-degenerate. By Lem. 8.2.10 the map $b_L^V: V \rightarrow V^*$ with $b_L^V(v_1)(v_2) = b_V(v_1, v_2)$ is an isomorphism. Let b_L^W be the analogous map for W and the inner product b_W . Observe that the linear map $T: V \rightarrow W$ induces a linear transformation $\widehat{T}: W^* \rightarrow V^*$ as follows: For $\varphi \in W^*$ the map $\widehat{T}(\varphi): V \rightarrow \mathbb{R}$ is given by

$$\widehat{T}(\varphi)(v) = \varphi(T(v)) .$$

Now note that (24) can be reformulated as follows:

$$\widehat{T} \circ b_L^W = b_L^V \circ T^* .$$

In fact, we have for $v \in V$ and $w \in W$ that

$$\begin{aligned} \left[(\widehat{T} \circ b_L^W)(w) \right] (v) &= \widehat{T}(b_L^W(w))(v) = b_L^W(w)(T(v)) = b_W(w, T(v)) , \\ \left[(b_L^V \circ T^*)(w) \right] (v) &= b_L^V(T^*(w))(v) = b_V(T^*(w), v) \end{aligned}$$

and equality of the two right hand sides is exactly the statement in (24). Since b_L^V is an isomorphism, T^* is uniquely defined by

$$(b_L^V)^{-1} \circ \widehat{T} \circ b_L^W = T^* .$$

Let $A = [T]_\beta^\gamma$ with entries (a_{ij}) and let $B = [T^*]_\gamma^\beta$ with entries (b_{kl}) be the matrix representations of T and T^* respectively. Since γ is an orthonormal basis, we have $b_W(w_k, w_i) = 0$ if $k \neq i$ and $b_W(w_k, w_k) = 1$ and therefore

$$b_W(w_k, T(v_j)) = \sum_{i=1}^m b_W(w_k, a_{ij}w_i) = \sum_{i=1}^m a_{ij}b_W(w_k, w_i) = a_{kj} .$$

Likewise,

$$b_V(T^*(w_k), v_j) = \sum_{l=1}^n b_V(b_{lk}v_l, v_j) = \sum_{l=1}^n b_{lk}b_V(v_l, v_j) = b_{jk}$$

Hence, (24) implies that $b_{jk} = a_{kj}$, or in other words, $B = A^t$. \square

For complex vector spaces equipped with inner products the situation only changes slightly:

Theorem 8.5.2. *Let V and W be finite-dimensional vector spaces over \mathbb{C} , let*

$$\begin{aligned} b_V &: V \times V \rightarrow \mathbb{C} , \\ b_W &: W \times W \rightarrow \mathbb{C} \end{aligned}$$

be inner products on V and W , respectively. Let $T: V \rightarrow W$ be a linear map. Then there is a unique linear map $T^: W \rightarrow V$ with the property that*

$$b_W(w, T(v)) = b_V(T^*(w), v) \tag{24}$$

for all $v \in V$ and $w \in W$. Moreover, if $\beta = \{v_1, \dots, v_n\}$ is an orthonormal basis for V , $\gamma = \{w_1, \dots, w_m\}$ is an orthonormal basis for W , then we have

$$[T^*]_\gamma^\beta = \overline{([T]_\beta^\gamma)^t} ,$$

i.e. the matrix representation of T^ is the conjugate transpose of the matrix representation of T .*

Proof. The proof for existence and uniqueness of T^* is exactly the same as for Thm. 8.5.1. Let $A = [T]_\beta^\gamma$ with entries (a_{ij}) and let $B = [T^*]_\gamma^\beta$ with entries (b_{kl}) be the matrix representations of T and T^* respectively. The only thing that changes in the complex case are the two computations:

$$\begin{aligned} b_W(w_k, T(v_j)) &= \sum_{i=1}^m b_W(w_k, a_{ij}w_i) = \sum_{i=1}^m a_{ij}b_W(w_k, w_i) = a_{kj} , \\ b_V(T^*(w_k), v_j) &= \sum_{l=1}^n b_V(b_{lk}v_l, v_j) = \sum_{l=1}^n \overline{b_{lk}}b_V(v_l, v_j) = b_{jk} , \end{aligned}$$

where we used that the inner product is conjugate linear in the first entry. This implies $b_{jk} = \overline{a_{kj}}$, or in other words, $B = \overline{A^t}$. \square

In both cases, real and complex, we make the following definition:

Definition 8.5.3. Let V and W be finite-dimensional vector spaces with inner products b_V and b_W , respectively. Let $T: V \rightarrow W$ be a linear map. Then the associated linear transformation $T^*: W \rightarrow V$ defined in Thm. 8.5.1 is called the **adjoint of T** . A linear transformation $S: V \rightarrow V$ that satisfies $S^* = S$ is called **self-adjoint**.

Remark 8.5.4. For a matrix $A \in M_{m \times n}(\mathbb{C})$ taking the conjugate transpose is often also denoted by A^* , i.e.

$$A^* = \overline{A^t}.$$

This is consistent with taking the associated linear map and we have $(L_A)^* = L_{A^*}$. Since conjugation maps every real number to itself, we have $A^* = A^t$ for $A \in M_{m \times n}(\mathbb{R})$, which is consistent with Thm. 8.5.1 for the real case, i.e. $(L_A)^* = L_{A^*}$ is still true for real matrices.

Corollary 8.5.5. Let V be a finite dimensional vector space over \mathbb{R} or over \mathbb{C} and let $T: V \rightarrow V$ be a linear transformation. Let β be an orthonormal basis for V . Then T is self-adjoint if and only if

$$[T]_{\beta}^{\beta} = \left([T]_{\beta}^{\beta}\right)^*.$$

Proof. The equation $T = T^*$ is equivalent to $[T]_{\beta}^{\beta} = [T^*]_{\beta}^{\beta}$, but $[T^*]_{\beta}^{\beta} = \left([T]_{\beta}^{\beta}\right)^*$ by Thm. 8.5.1 in the real case and Thm. 8.5.2 in the complex case. \square

Example 8.5.6. Let $V = \mathbb{R}^2$, let $b_V = \langle \cdot, \cdot \rangle$ be the dot product and consider the linear transformation

$$T: V \rightarrow V, \quad (x, y) \mapsto (x, 2x + y).$$

Let $\beta = \{e_1, e_2\}$ be the standard basis for \mathbb{R}^2 . To find the adjoint of T we first compute its matrix representation with respect to β . We have

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

We know from Thm. 8.5.1 that $[T^*]_{\beta}^{\beta} = \left([T]_{\beta}^{\beta}\right)^*$, where the $*$ in this case is just transposition of the matrix, since V is a vector space over \mathbb{R} . Hence,

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

We can use the matrix representation to write down the adjoint linear map. It turns out to be

$$T^*: V \rightarrow V, \quad (x, y) \mapsto (x + 2y, y)$$

Since $T \neq T^*$, the linear map T is not self-adjoint. We can also see this directly from the above matrices. Let us check that the linear transformation T^* indeed satisfies the

defining property of the adjoint, i.e. that $\langle w, T(v) \rangle = \langle T^*(w), v \rangle$. With $w = (x_1, y_1)$ and $v = (x_2, y_2)$ we get

$$\begin{aligned}\langle (x_1, y_1), T(x_2, y_2) \rangle &= \langle (x_1, y_1), (x_2, 2x_2 + y_2) \rangle = x_1x_2 + 2y_1x_2 + y_1y_2, \\ \langle T^*(x_1, y_1), (x_2, y_2) \rangle &= \langle (x_1 + 2y_1, y_1), (x_2, y_2) \rangle = x_1x_2 + 2y_1x_2 + y_1y_2.\end{aligned}$$

Example 8.5.7. Let $V = \mathbb{R}^2$, let $b_V = \langle \cdot, \cdot \rangle$ be the dot product and consider the linear transformation

$$T: V \rightarrow V, \quad (x, y) \mapsto (x + 2y, 2x + y).$$

Let $\beta = \{e_1, e_2\}$ be the standard basis for \mathbb{R}^2 . The matrix representation of T with respect to β is

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

Observe that $[T]_{\beta}^{\beta} = \left([T]_{\beta}^{\beta}\right)^*$ and therefore $T = T^*$ by Cor. 8.5.5. In particular, T is self-adjoint.

Example 8.5.8. Let $n \in \mathbb{N}$ and consider $V = P_n(\mathbb{C})$. Let

$$b_V(p, q) = \int_{-1}^1 \overline{p(x)} q(x) dx.$$

This defines an inner product on V (It is a good exercise to check this!). Let $T: V \rightarrow V$ be the following linear transformation: Let $g(x) = 1 - x^2$ and define

$$T(p) = (g \cdot p')',$$

where f' denotes the first derivative of f . Since differentiation lowers the degree of a polynomial by 1 and g is a polynomial of degree 2, the degree of the polynomial $T(p)$ is at most n . Therefore the map is well-defined and it is also linear. Note that $g(x) = \overline{g(x)}$ and $g(1) = g(-1) = 0$. Using integration by parts we obtain for any $p, q \in V$

$$\begin{aligned}b_V(p, T(q)) &= \int_{-1}^1 \overline{p(x)} (g \cdot q')'(x) dx = \overline{p(x)} g(x) q'(x) \Big|_{-1}^1 - \int_{-1}^1 \overline{p'(x)} g(x) q'(x) dx \\ &= - \int_{-1}^1 \overline{(g \cdot p')(x)} q'(x) dx = \int_{-1}^1 \overline{(g \cdot p')'(x)} q(x) dx = b_V(T(p), q).\end{aligned}$$

Since T^* is the unique linear transformation, which satisfies $b_V(p, T(q)) = b_V(T^*(p), q)$ for all $p, q \in V$ we must have $T = T^*$ and therefore T is self-adjoint.

We have seen in the section about eigenvalues and eigenvectors that not all linear transformations are diagonalisable and that the best we can hope for in the general situation is to find the Jordan normal form. Self-adjoint linear transformations behave very nicely in this respect: It turns out that such a map is always diagonalisable, all its eigenvalues are real numbers (even in the case of complex vector spaces) and that there is an orthonormal basis consisting of eigenvectors. We will prove this in the next theorem:

Theorem 8.5.9. *Let V be a finite dimensional vector space over \mathbb{R} or over \mathbb{C} and let b_V be an inner product on V . Let $T: V \rightarrow V$ be a self-adjoint linear transformation.*

- a) *The eigenvalues of T are real numbers.*
- b) *Eigenvectors corresponding to distinct eigenvalues are orthogonal.*
- c) *There is an orthonormal basis of V consisting of eigenvectors for T . In particular, T is diagonalisable.*

Proof. We only need to prove a) for vector spaces over the complex numbers. Let $\lambda \in \mathbb{C}$ be an eigenvalue of T . This means that there exists a non-zero vector $v \in V$ with the property that $T(v) = \lambda v$. Now note that

$$\lambda b_V(v, v) = b_V(v, \lambda v) = b_V(v, T(v)) = b_V(T(v), v) = b_V(\lambda v, v) = \bar{\lambda} b_V(v, v) , \quad (25)$$

where we used that T is self-adjoint to get $b_V(v, T(v)) = b_V(T(v), v)$. Since $v \neq 0$ and b_V is non-degenerate, we must have $b_V(v, v) \neq 0$. Hence, (25) implies that $\lambda = \bar{\lambda}$ or in other words $\lambda \in \mathbb{R}$.

Let $\lambda_1, \lambda_2 \in \mathbb{R}$ be two distinct eigenvalues of T (as shown in a) this means $\lambda_i \in \mathbb{R}$). Let $v_i \in V$ be an eigenvector corresponding to λ_i for $i \in \{1, 2\}$. Then we have

$$\begin{aligned} \lambda_2 b_V(v_1, v_2) &= b_V(v_1, \lambda_2 v_2) = b_V(v_1, T(v_2)) = b_V(T(v_1), v_2) \\ &= b_V(\lambda_1 v_1, v_2) = \lambda_1 b_V(v_1, v_2) . \end{aligned}$$

But since $\lambda_1 \neq \lambda_2$, this means that $b_V(v_1, v_2) = 0$ or, in other words, the eigenvectors are orthogonal. This proves statement b).

We will prove c) for vector spaces over \mathbb{C} first and proceed by induction over the dimension of V . If we assume that V is one-dimensional, then $T: V \rightarrow V$ is just multiplication by a scalar $\lambda \in \mathbb{R}$ and any non-zero vector $v \in V$ is an eigenvector of T corresponding to the eigenvalue λ . Let

$$c = \frac{1}{\sqrt{b_V(v, v)}} .$$

The vector $w = cv$ has length 1. Hence, $\beta = \{w\}$ is an orthonormal basis for V consisting of eigenvectors.

Let $n = \dim(V)$ and assume that the statement is true for any self-adjoint linear transformation $S: W \rightarrow W$ on a vector space W with $\dim(W) < n$. Let $p_T(t)$ be the characteristic polynomial of T . Any polynomial over \mathbb{C} splits. Therefore T has at least one eigenvalue $\lambda \in \mathbb{C}$. By part a) of the theorem λ is a real value. Let $v_1 \in V$ be an eigenvector with corresponding eigenvalue λ . By rescaling we may assume that $b_V(v_1, v_1) = 1$. Now consider $W = \{v_1\}^\perp$, i.e. the orthogonal complement of v_1 . We claim that T restricts to a linear map $T: W \rightarrow W$. To see this we have to show that $T(w)$ is orthogonal to v_1 for all $w \in W$. This is true, since $b_V(w, v_1) = 0$ implies that

$$b_V(T(w), v_1) = b_V(w, T(v_1)) = b_V(w, \lambda v_1) = \lambda b_V(w, v_1) = 0 .$$

The restricted linear map $T: W \rightarrow W$ is still self-adjoint. Since $\dim(W) = n - 1$, the vector space W has an orthonormal basis $\{v_2, \dots, v_n\}$ consisting of eigenvectors of T by our induction hypothesis. But then $\beta = \{v_1, v_2, \dots, v_n\}$ is an orthonormal basis for V , which consists of eigenvectors of the original map T .

The argument is very similar for vector spaces over \mathbb{R} . However, we need to prove that $p_T(t)$ splits. This can be seen as follows: Fix an arbitrary orthonormal basis γ of V and let $A = [T]_\gamma^\gamma$. We may consider the real entries of A as complex numbers and look at $L_A: \mathbb{C}^n \rightarrow \mathbb{C}^n$. Note that $A^* = A$ and therefore L_A is self-adjoint. Moreover, $p_T(t) = p_A(t)$. Let $\lambda \in \mathbb{C}$ be a root of $p_A(t)$. Then it is an eigenvalue of L_A and hence a real number by part a) of the theorem. Since this is true for all roots and $p_A(t)$ splits over \mathbb{C} , we have that $p_T(t)$ splits over \mathbb{R} . \square

The next class of maps we will look at are the linear transformations preserving a given inner product. As we will see below the set of all these transformations forms a group.

Definition 8.5.10. Let V be a vector space over \mathbb{C} and let b be an inner product on V . The **group of isometries** or **unitary group** $U(V)$ of V is defined to be

$$U(V, b) = \text{Isom}(V) = \{T: V \rightarrow V \mid b(T(v), T(w)) = b(v, w) \forall v, w \in V\} .$$

The terminology is slightly different for real vector spaces: Let V be a vector space over \mathbb{R} and let b be an inner product on V . The **group of isometries** or **orthogonal group** $O(V)$ of V is defined to be

$$O(V, b) = \text{Isom}(V) = \{T: V \rightarrow V \mid b(T(v), T(w)) = b(v, w) \forall v, w \in V\} .$$

Lemma 8.5.11. Let V be a vector space over \mathbb{C} or over \mathbb{R} and let b be an inner product on V . Let $T: V \rightarrow V$ be an isometry. Then T is invertible and

$$T^{-1} = T^* .$$

Proof. To prove that T is invertible it suffices to show that it is injective. Let $v \in \ker(T)$. For all $w \in V$ we have

$$b(v, w) = b(T(v), T(w)) = b(0, T(w)) = 0 ,$$

where we used that T is an isometry to obtain the first equality. Since b is non-degenerate, this implies $v = 0$ and hence T is indeed injective. Let $v, w \in V$ now be arbitrary. Since T is isometry, the following holds

$$b(v, w) = b(T(v), T(w)) = b(T^*(T(v)), w) .$$

But because b is non-degenerate, this implies $T^* \circ T = I_V$ and therefore $T^* = T^{-1}$. \square

Remark 8.5.12. The previous lemma shows that $U(V, b)$ and $O(V, b)$ are actually groups with respect to the matrix multiplication. If the inner product on V is clear from the context, these groups are also often denoted just by $U(V)$ and $O(V)$. If $V = \mathbb{C}^n$ and $b = \langle \cdot, \cdot \rangle$, then $U(V)$ is denoted by $U(n)$. Notice that

$$U(n) = \{U \in M_n(\mathbb{C}) \mid UU^* = U^*U = I_n\} .$$

If $V = \mathbb{R}^n$ and b is again the dot product, then $O(V)$ is also denoted by $O(n)$ and

$$O(n) = \{A \in M_n(\mathbb{R}) \mid AA^t = A^tA = I_n\} .$$

Theorem 8.5.13. *Let V be a vector space over \mathbb{C} and let b be an inner product on V . Let $T \in U(V, b)$.*

- a) *All eigenvalues $\lambda \in \mathbb{C}$ of T satisfy $|\lambda| = 1$ (i.e. they lie on the unit circle).*
- b) *Eigenvectors of distinct eigenvalues are orthogonal.*
- c) *There exists an orthonormal basis of eigenvectors for T . In particular, T is diagonalisable.*

Proof. First note that if v is an eigenvector of the invertible linear map T corresponding to the eigenvalue λ , then $\lambda \neq 0$ and v is also an eigenvector of T^{-1} corresponding to the eigenvalue λ^{-1} . With this observation the proof of a) is very similar to part a) in Thm. 8.5.9. Let $v \in V$ be an eigenvector of the unitary linear transformation T . Then

$$\lambda b(v, v) = b(v, T(v)) = b(T^*(v), v) = b(T^{-1}(v), v) = b(\lambda^{-1}v, v) = \overline{\lambda^{-1}} b(v, v) ,$$

where we used that T is unitary to get the third equality. Since $v \neq 0$ and b is non-degenerate, we must have $b(v, v) \neq 0$ and therefore $\lambda = \overline{\lambda^{-1}}$ or equivalently

$$|\lambda|^2 = \overline{\lambda} \lambda = 1 .$$

Let $\lambda_1, \lambda_2 \in \mathbb{C}$ be distinct eigenvalues of T . Let $v_i \in V$ be eigenvectors with corresponding eigenvalue λ_i for $i \in \{1, 2\}$. Then

$$\lambda_2 b(v_1, v_2) = b(v_1, T(v_2)) = b(T^*(v_1), v_2) = b(T^{-1}(v_1), v_2) = \overline{\lambda_1^{-1}} b(v_1, v_2) = \lambda_1 b(v_1, v_2) ,$$

where we used that $\overline{\lambda_1^{-1}} \lambda_1 = 1$ (by part a) to get the last equality. But since $\lambda_1 \neq \lambda_2$ this implies that $b(v_1, v_2) = 0$, which proves b).

The last statement is again shown by induction. If $\dim(V) = 1$, any non-zero vector of length 1 gives an orthonormal basis of V that consists of eigenvectors (see the proof of Thm. 8.5.9 for details). Now let $n = \dim(V)$ and assume that the statement holds for all $S \in U(W, b_W)$ with $\dim(W) < n$. Since $p_T(t)$ splits over \mathbb{C} , the linear map T has at least one eigenvalue $\lambda \in \mathbb{C}$. Let $v_1 \in V$ be an eigenvector of T corresponding to λ . Without loss of generality we may assume $b(v_1, v_1) = 1$. Let $W = \{v_1\}^\perp \subset V$. Let $w \in W$. Then we have

$$0 = b(w, v_1) = b(T(w), T(v_1)) = \lambda b(T(w), v_1) .$$

But since $|\lambda| = 1$ by part a), we have $\lambda \neq 0$ and therefore the above equality gives $b(T(w), v_1) = 0$ or equivalently: $T(w) \in W$. Hence, we can restrict T to a linear transformation $T: W \rightarrow W$, which is an isometry, since $b(w_1, w_2) = b(T(w_1), T(w_2))$ still holds for all $w_1, w_2 \in W$. Since $\dim(W) = n - 1$ there exists an orthonormal basis $\{v_2, \dots, v_n\}$ consisting of eigenvectors for the restricted linear transformation $T: W \rightarrow W$. But then $\beta = \{v_1, v_2, \dots, v_n\}$ is an orthonormal basis for V that consists of eigenvectors for the original linear map T . \square

9 What's next?

Linear algebra has an abundance of applications in other areas, and I would like to close these notes by just mentioning a few of them. The idea of vector spaces can be generalised in several directions. A vector space is equipped with a scalar multiplication by elements in a field F . The integers \mathbb{Z} do not form a field, since not every non-zero element has a multiplicative inverse (i.e. the element $\frac{1}{n}$ is not in \mathbb{Z} if $n \notin \{-1, 1\}$). However, they still have an addition and a multiplication, which satisfy the usual distributive laws. We say that the integers form a **ring**. There are many interesting rings out there, some of which you have already seen. The matrices $M_{n \times n}(F)$ also form a (non-commutative) ring for example. This gives rise to the question whether there are “vector spaces over rings”. These exist, they are called **modules** over rings and they are studied in algebra. Many of the constructions that we employed for vector spaces have to be modified or do not work at all for modules. Hence, their theory is more complicated, but also more flexible.

Another big topic, in which the concepts outlined in this lecture will resurface, is **functional analysis**. In short, this is the study of linear transformations on infinite-dimensional vector spaces. These linear transformations are usually called **operators**. Many infinite-dimensional vector spaces come equipped with natural inner products. If a vector space is complete with respect to the norm induced by an inner product, then it is called a **Hilbert space**. Operators on Hilbert spaces play a crucial role not only in Mathematics, but also in Physics, where they appear in **quantum mechanics**.

We have already seen that certain bilinear forms on real vector spaces have applications in Physics as well. To understand the **theory of special relativity**, physicists study the symmetries preserving the bilinear form on Minkowski space defined above. In particular, they look at linear transformations $U: \mathbb{R}^4 \rightarrow \mathbb{R}^4$ with the property that $b(Uv, Uw) = b(v, w)$ for all $v, w \in \mathbb{R}^4$. The set of all transformations that preserve a given bilinear form is a group and there is a rich mathematical theory about these symmetry groups. The keyword to look up here is **Lie groups**.

Approximation problems (for example in computer science) provide another application of linear algebra. For example, solving a non-linear differential equation numerically is sometimes possible by linearising it first, which effectively turns it into a linear algebra problem. To find the main directions in a point cloud of data computer scientists use the so-called **principal component analysis**, which boils down to finding the eigenvectors and the eigenvalues of a certain matrix.

This list provides just some of the applications of linear algebra in “real life” and I hope that these notes will be helpful for your journeys into some of these directions.