Linear Algebra II for Physics
Recap of Linear Algebra I

Alberto S. Cattaneo

Institute of Mathematics, University of Zurich

Last update: February 14, 2022

Contents

CHAPTER 1

Basic Concepts

This chapter is a summary of the basic concepts introduced in Linear Algebra I which will be used in this course.

1.1. Groups, rings, and fields

Definition 1.1 (Groups). A group is a set G with a distinguished element e, called the neutral element, and an operation, called "multiplication,"

$$\begin{array}{ccc} G \times G & \to & G \\ (a,b) & \mapsto & ab \end{array},$$

which is associative—i.e., $a(bc) = (ab)c$ for all $a, b, c \in G$—and which satisfies $ae = ea = a$. Moreover, for all $a \in G$, there is an element denoted by a^{-1}, and called the inverse of a, satisfying $aa^{-1} = a^{-1}a = e$ for all $a \in G$.

One can show that the inverse is unique, that $(a^{-1})^{-1} = a$ and that $(ab)^{-1} = b^{-1}a^{-1}$ for all $a, b \in G$.

Definition 1.2 (Abelian groups). A group G is called abelian if, in addition, $ab = ba$ for all $a, b \in G$. If G is abelian, one often uses the additive notation in which the neutral element is denoted by 0, the multiplication is called "addition" and denoted by

$$(a, b) \mapsto a + b,$$

and the inverse of an element a is denoted by $-a$ (and called the "additive inverse" or the "opposite" of a).

Examples 1.3. Here are some examples of groups:

(1) $\mathbb{Z}$ with the usual 0 and the usual addition is an abelian group (in additive notation).
(2) $\mathbb{Z}_{>1}$ with $e = 1$ and the usual multiplication is an abelian group (not in additive notation).

(3) Invertible $n \times n$ matrices form the group GL_n with e the identity matrix and the usual multiplication of matrices; this group is nonabelian for $n > 1$.[1]
(4) The set $\mathrm{Aut}(S)$ of bijective maps of a set S to itself form a group with multiplication given by the composition and neutral element given by the identity map. If $S = \{1, \dots, n\}$ this group is called the symmetric group on n elements and is denoted by S_n (or $\mathrm{Sym}(n)$); its elements are called permutations.

Definition 1.4 (Rings). A ring is an abelian group $(R, 0, +)$ together with a second associative operation, called "multiplication,"

$$\begin{array}{ccc} R \times R & \to & R \\ (a,b) & \mapsto & ab \end{array},$$

which is also distributive; i.e.,

$$a(b+c) = ab + ac \quad \text{and} \quad (a+b)c = ac + bc$$

for all $a, b, c \in R$. A ring R is called:

(1) a ring with one if it possesses a special element, denoted by 1, such that $a1 = 1a = a$ for all $a \in R$;
(2) a commutative ring if $ab = ba$ for all $a, b \in R$.

Examples 1.5. Here are some examples of rings:

(1) $\mathbb{Z}$—with the usual addition, multiplication, zero, and one—is a commutative ring with one.
(2) The set $2\mathbb{Z}$ of even numbers—with the usual addition, multiplication, and zero—is a commutative ring without one.
(3) The set $\mathrm{Mat}_{n\times n}$ of $n \times n$ matrices, with the usual addition and multiplication, with 0 the matrix whose all entries are 0, and with 1 the identity matrix, is a ring with one. It is noncommutative for $n > 1$.
(4) Polynomials form a commutative ring with one.
(5) If I is an open interval, the set $C^0(I)$ of continuous functions on I, the set $C^k(I)$ of k times continuously differentiable functions on I, and the set $C^\infty(I)$ of functions on I that are continuously differentiable any number of times are commutative rings with one. Recall that the operations are defined as

$$(f+g)(x) := f(x) + g(x), \quad (fg)(x) := f(x)\,g(x), \qquad x \in I.$$

[1] We assume, without explicitly recapping, the knowledge of matrices, including the notion of sum, product, and transposition. We will recap trace and determinant in Section 1.6.

The zero element is the function $0(x) = 0$ for all x and the one element is the function $1(x) = 1$ for all x.

DEFINITION 1.6 (Fields). A field $\mathbb{K}$ is a commutative ring with one in which every element different from zero is invertible. This is equivalent to saying that $\mathbb{K}^\times := \mathbb{K} \setminus \{0\}$ is a commutative group (not in additive notation).

The only fields we are going to consider in this course are the field $\mathbb{R}$ of real numbers and the field $\mathbb{C}$ of complex numbers.

Many of the results we present actually hold for any field, and most of the results hold for any field of characteristic zero, like $\mathbb{R}$ or $\mathbb{C}$, i.e., a field $\mathbb{K}$ such that there is no nonzero integer n satisfying $na = 0$ for a nonzero element a of $\mathbb{K}$.[2]

DEFINITION 1.7 (Subobjects). A subset of a group/ring/field which retains all the structures is called a subgroup/subring/subfield.

1.2. Vector spaces

A vector space over a field $\mathbb{K}$, whose elements are called scalars, is an abelian group $(V, +, 0)$—in additive notation—whose elements are called vectors, together with an operation

$$\begin{array}{ccc} \mathbb{K} \times V & \to & V \\ (\lambda, v) & \mapsto & \lambda v \end{array},$$

called multiplication by a scalar or scalar multiplication,[3] satisfying

$$\lambda(\mu v) = (\lambda\mu)v, \quad (\lambda + \mu)v = \lambda v + \mu v, \quad \lambda(v + w) = \lambda v + \lambda w,$$

for all $\lambda, \mu \in \mathbb{K}$ and all $v, w \in V$.

EXAMPLE 1.8 (Column vectors). The set $\mathbb{K}^n$ of n-tuples of scalars, conventionally arranged in a column and called column vectors, is a vector space with

$$\begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix} + \begin{pmatrix} w^1 \\ \vdots \\ w^n \end{pmatrix} := \begin{pmatrix} v^1 + w^1 \\ \vdots \\ v^n + w^n \end{pmatrix}, \quad 0 := \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix},$$

and

$$\lambda \begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix} := \begin{pmatrix} \lambda v^1 \\ \vdots \\ \lambda v^n \end{pmatrix}.$$

---

[2] By na one means the sum $a + \cdots + a$ with n summands.

[3] not to be confused with the scalar product a.k.a. the dot product or the inner product

The scalars v^i are called the components of the column vector, which is usually denoted by the corresponding boldface latter $\boldsymbol{v}$.[4]

EXAMPLE 1.9 (Row vectors). The set $(\mathbb{K}^n)^*$ of n-tuples of scalars, conventionally arranged in a row and called row vectors, is a vector space with

$$\big(v_1, \dots, v_n\big) + \big(w_1, \dots, w_n\big) := \big(v_1 + w_1, \dots, v_n + w_n\big)\,, \quad 0 := \big(0, \dots, 0\big)\,,$$

and

$$\lambda\,\big(v_1, \dots, v_n\big) = \big(\lambda v_1, \dots, \lambda v_n\big)\,.$$

The scalars v^i are called the components of the row vector, which is also usually denoted by the corresponding boldface latter $\boldsymbol{v}$.[4]

REMARK 1.10 (Components and indices). The scalars forming a row or column vectors are called its components. We will consistently denote the components of a column vectors with an upper index and the components of a row vectors with a lower index. This is a nowadays standard convention (especially in physics) that comes in handy with Einstein's convention for sums (which we will introduce in Definition 1.33).

EXAMPLE 1.11 (The trivial vector space). The vector space $V = \{0\}$ consisting only of the neutral element 0 is called the trivial vector space. It is denoted by 0, but also by $\mathbb{K}^0$ (and, if you wish, $(\mathbb{K}^0)^*$).

REMARK 1.12 (The zero notation). Observe that the symbol 0 is used for all of the following:

(1) The neutral element of an abelian group in additive notation.
(2) The zero element of a ring or a field.
(3) The zero element of a vector space.
(4) The vector space consiting only of the zero element ($0 = \{0\}$).
(5) A constant map having value 0 (e.g., the continuous real function $x \mapsto 0$ for all $x \in \mathbb{R}$, or the map $V \to W$, $v \mapsto 0$, where V and W are vector spaces).
(6) A matrix whose entries are all equal to 0 (even though we will prefer the notation $\mathbf{0}$).

EXAMPLE 1.13 (Polynomials). The ring $\mathbb{K}[x]$ of polynomials in an undetermined x with coefficients in $\mathbb{K}$ (i.e., expressions of the form $p = a_0 + a_1 x + \cdots + a_d x^d$, for some d, and $a_i \in \mathbb{K}$ for all i) is also a vector space over $\mathbb{K}$ with scalar multiplication $\lambda p := \lambda a_0 + \lambda a_1 x + \cdots + \lambda a_d x^d$ and the usual addition the addition of polynomials (i.e., addition of the coefficients).

[4] Other common notations are $\vec{v}$ and $\underline{v}$.

Example 1.14 (Functions). The rings of functions $C^k(I)$, $k \in \mathbb{N} \cup \{\infty\}$, of Example 1.5.(5) are also vector spaces over $\mathbb{R}$ with scalar multiplication $(\lambda f)(x) := \lambda\, f(x)$.

Definition 1.15 (Subspaces). A subset W of a vector space V that retains all the structures is called a (vector) subspace. Equivalently, $W \subseteq V$ is a subspace iff for every $w, \widetilde{w} \in W$ and for every $\lambda \in \mathbb{K}$ we have $w + \widetilde{w} \in W$ and $\lambda w \in W$.

Definition 1.16 ((Direct) sums of subspaces). If W_1 and W_2 are subspaces of V, we denote by $W_1 + W_2$ the subset of elements of V consisting of sums of elements of W_1 and W_2; i.e.,

$$W_1 + W_2 = \{w_1 + w_2,\ w_1 \in W_1,\ w_2 \in W_2\}.$$

It is also a subspace of V. If $W_1 \cap W_2 = \{0\}$, the sum is called the direct sum and is denoted by $W_1 \oplus W_2$.

Remark 1.17. A vector $v \in W_1 \oplus W_2$ uniquely decomposes as $v = w_1 + w_2$ with $w_1 \in W_1$ and $w_2 \in W_2$. The vectors w_1 and w_2 are called the components of v in the direct sum.

Proof. If $v = \widetilde{w}_1 + \widetilde{w}_2$ were another decomposition, by taking the difference we would get $w_1 - \widetilde{w}_1 = \widetilde{w}_2 - w_2$. Since the left hand side is in W_1, the right hand side is in W_2, and $W_1 \cap W_2 = \{0\}$, we have $w_1 - \widetilde{w}_1 = 0 = \widetilde{w}_2 - w_2$. □

Remark 1.18. Vice versa, if every vector in $W_1 + W_2$ has a unique decomposition $v = w_1 + w_2$ with $w_i \in W_i$, $i = 1, 2$, then this is a direct sum (i.e., $W_1 \cap W_2 = 0$).

Proof. Suppose $v \in W_1 \cap W_2$. From $0 = v - v$, we see that the first summand, v, is the component of 0 in W_1 and the second summand, $-v$, is its component in W_2. Since we can decompose the zero vector also as $0 = 0 + 0$ (and 0 belongs to both W_1 and W_2), by the assumed uniqueness of the decomposition, we then have $v = 0$. □

Definition 1.19. If W is a subspace of V, a susbpace W' such that $V = W \oplus W'$ is called a complement.

Every subspace admits a complement, see Lemma 1.27 and Proposition 1.55. This is elementary in the case of finite-dimensional vector spaces and requires the axiom of choice for infinite-dimensional ones (see Digression 1.56).

Definition 1.20. We may generalize Definition 1.16 to a sum of several subspaces $W_1, \dots, W_k$:

$$W_1 + \cdots + W_k := \{v_1 + \cdots + v_k,\ v_i \in W_i,\ i = 1, \dots, k\}.$$

We generalize the notion of direct sum via the property in Remark 1.18:

Definition 1.21. A sum $W_1 + \cdots + W_k$ is called a direct sum, and it is denoted by

$$W_1 \oplus \cdots \oplus W_k \quad \text{or} \quad \bigoplus_{i=1}^{k} W_i,$$

if every vector v in it has a unique decomposition $v = w_1 + \cdots + w_k$ with $w_i \in W_i$, $i = 1, \ldots, k$.

Remark 1.22. An easier criterion is the following. A sum $W_1 + \cdots + W_k$ is a direct sum iff the zero vector has a unique decomposition.

Proof. If the sum is direct, the zero vector has a unique decomposition by definition, like every other vector.

Vice versa, suppose that the zero vector has a unique decomposition and that a vector v can be written both as $w_1 + \cdots + w_k$ and as $w'_1 + \cdots + w'_k$ with $w_i, w'_i \in W_i$. By taking the difference of these two decompositions, we get $(w_1 - w'_1) + \cdots + (w_k - w'_k) = 0$, which then implies $w_i = w'_i$ for every i. Therefore, the decomposition of every vector in the sum is unique. $\square$

Remark 1.23. A consequence of this is that for any $0 < r < k$ we have

$$\bigoplus_{i=1}^{k} W_i = \bigoplus_{i=1}^{r} W_i \oplus \bigoplus_{i=r+1}^{k} W_i.$$

Proof. We have to prove that $\bigoplus_{i=1}^{r} W_i \cap \bigoplus_{i=r}^{k} W_i = 0$. If v is in the intersection, we may uniquely decompose it as $w_1 + \cdots + w_r$ and as $w_{r+1} + \cdots + w_k$ with $w_i \in W_i$. Taking the difference, we get $0 = w_1 + \cdots + w_r - w_{r+1} - \cdots - w_k$. By uniqueness of the decomposition in $\bigoplus_{i=1}^{k} W_i$, we get $w_i = 0$ for every i. $\square$

Remark 1.24. Note that, by definition, $W_1 \oplus W_2 = W_2 \oplus W_1$ and that, by the last remark, $(W_1 \oplus W_2) \oplus W_3 = W_1 \oplus (W_2 \oplus W_3)$, where W_1, W_2 and W_3 are subspaces of V. Therefore, the direct sum is commutative and associative. It also have a "neutral element," namely the zero subspace $0 := \{0\}$.

Remark 1.25 (Infinite sums). Let $(W_i)_{i \in S}$ be a, possibly infinite, collection of subspaces of a vector space V. Their sum is the subspace of V consisting of all vectors of the form $w_{i_1} + \cdots + w_{i_k}$, $w_{i_j} \in W_{i_j}$ for $j = 1, \ldots, k$ and some integer k. This sum is direct, denoted by

$$\bigoplus_{i \in S} W_i,$$

if each vector in it has a unique decomposition (or, equivalently, if $w_{i_1} + \cdots + w_{i_k} = 0$, $w_{i_j} \in W_{i_j}$, $i_j \neq i_{j'}$ for $j \neq j'$, implies $w_{i_j} = 0$ for all j).

Definition 1.26 (Direct sums of vector spaces). If V_1 and V_2 are vector spaces over the same field, we denote by $V_1 \oplus V_2$—and call it the direct sum of V_1 and V_2—the Cartesian product $V_1 \times V_2$ of pairs of elements of V_1 and V_2 with the following vector space structure:

$$(v_1, v_2)+(\widetilde{v}_1, \widetilde{v}_2) := (v_1+\widetilde{v}_1, v_2+\widetilde{v}_2),\ 0 := (0,0),\ \lambda(v_1, v_2) := (\lambda v_1, \lambda v_2).$$

The spaces V_1 and V_2 are identified with the subspaces $\{(v, 0), v \in V_1\}$ and $\{(0, v), v \in V_2\}$ of $V := V_1 \oplus V_2$. Under this identification, $V_1 \cap V_2 = \{0\}$, so the notation for the direct sum of the vector spaces V_1 and V_2 fits with that of the direct sum of the subspaces V_1 and V_2 of V. This generalizes to a collection $V_1, \ldots, V_k$ of vector spaces over the same field. By $\bigoplus_{i=1}^k V_i$, we denote the Cartesian product $V_1 \times \cdots \times V_k$ with

$$\begin{aligned} (v_1, \ldots, v_k) + (\widetilde{v}_1, \ldots \widetilde{v}_k) &:= (v_1 + \widetilde{v}_1, \ldots, v_k + \widetilde{v}_k), \\ 0 &:= (0, \ldots, 0), \\ \lambda(v_1, \ldots, v_k) &:= (\lambda v_1, \ldots, \lambda v_k). \end{aligned}$$

Again, we may regard V_i as the subspace of $\bigoplus_{i=1}^k V_i$ consisting of $k$tuples with 0 in each position but the ith.

Note that $\mathbb{K}^n = \mathbb{K} \oplus \cdots \oplus \mathbb{K}$ with n summands. As a subspace of $\mathbb{K}^n$, the kth summand consists of vectors in which only the kth component may be different from zero.

Moreover, $\mathbb{K}^n = \mathbb{K}^r \oplus \mathbb{K}^l$ for all nonnegative integers r and l with $r + l = n$. In this case, as a subspace of $\mathbb{K}^n$, the first summand consists of vectors whose last l components are zero, and the second summand consists of vectors whose first r components are zero. We use this decomposition (with $l = 1$) for the following

Lemma 1.27. *Every subspace of $\mathbb{K}^n$ has a complement.*

Proof. Let W be a subspace of $\mathbb{K}^n$. We want to show that we can always find a subspace W' of $\mathbb{K}^n$ such that $W \oplus W' = \mathbb{K}^n$ (i.e., $W \cap W' = \{0\}$ and $W + W' = \mathbb{K}^n$).

If $n = 0$, there is nothing to prove, since necessarily $W = \{0\}$ and $W' = \{0\}$.

Otherwise, we prove the lemma by induction on $n > 0$. If $n = 1$, the proof is immediate: In case $W = \{0\}$, we take $W' = \mathbb{K}$. If, otherwise, W contains a nonzero vector $\boldsymbol{v}$, then $W = \mathbb{K}$, since every vector in $\mathbb{K}$ can be written as $\lambda \boldsymbol{v}$; therefore, $W' = \{0\}$.

Now assume we have proved the lemma for $\mathbb{K}^n$, and let W be a subspace of $\mathbb{K}^{n+1}$. Let W_1 be the subspace of vectors in W whose last component is zero and let W_2 be the subspace of vectors of W whose first n components are zero. We can view W_1 as a subspace of the first summand, $\mathbb{K}^n$, in the decomposition $\mathbb{K}^{n+1} = \mathbb{K}^n \oplus \mathbb{K}$ and W_2 as a subspace of the second summand, $\mathbb{K}$. By the induction assumption, there is a complement W_1' of W_1 in the first summand and a complement W_2' of W_2 in the second. Then $W_1' \oplus W_2'$ is a complement of W in $\mathbb{K}^n$. □

1.3. Linear maps

A map $F\colon V \to W$ between $\mathbb{K}$-vector spaces is called a linear map if

$$F(\lambda v + \mu \widetilde{v}) = \lambda F(v) + \mu F(\tilde{v})$$

for all $\lambda, \mu \in \mathbb{K}$ and all $v, \widetilde{v} \in V$.

EXAMPLES 1.28. Here are some examples of linear maps:

(1) The inclusion map of a subspace is linear.
(2) If V is the direct sum of vector spaces V_1 and V_2—i.e., $V = V_1 \oplus V_2$ as in Definition 1.26—then we have the linear maps, called canonical projections, $\pi_i\colon V \to V_i$, $i = 1, 2$, given by
$$\pi_i(v_1, v_2) = v_i.$$
If we regard V_1 and V_2 as subspaces of V, we may also regard the projections as linear maps $P_i\colon V \to V$:
$$P_1(v_1, v_2) = (v_1, 0) \quad \text{and} \quad P_2(v_1, v_2) = (0, v_2).$$
More precisely, $P_i = \iota_i \circ \pi_i$ where ι_i is the inclusion of V_i into V.
(3) Multiplication, from the left, by an $m \times n$ matrix defines a linear map $\mathbb{K}^n \to \mathbb{K}^m$.
(4) Multiplication, from the right, by an $m \times n$ matrix defines a linear map $(\mathbb{K}^m)^* \to (\mathbb{K}^n)^*$.
(5) The derivative defines a linear map $C^k(I) \to C^{k-1}(I)$, $f \mapsto f'$ (we assume $k \in \mathbb{N}_{>0} \cup \{\infty\}$).

REMARK 1.29. Here are some facts and notations.

(1) If F is linear, one often writes Fv instead of $F(v)$.
(2) The image of a linear map $F\colon V \to W$[5] is denoted by $\operatorname{im} F$ or $F(V)$ and is a subspace of W.

[5] i.e., the set of vectors $w \in W$ for which there is a $v \in V$ with $w = F(v)$

(3) The subset of elements of V mapped to 0 by a linear map $F\colon V \to W$ is denoted by $\ker F$ and is called its kernel. It is a subspace of V. A linear map F turns out to be injective iff $\ker F = \{0\}$.
(4) The composition of linear maps, say, $F\colon V \to W$ and $G\colon W \to Z$, is automatically linear. Instead of $G \circ F$ one often writes GF.
(5) If a linear map F is linear and invertible, its inverse map F^{-1} is automatically linear.
(6) A linear map $F\colon V \to W$ is also called a homomorphism from V to W.
(7) An invertible linear map $F\colon V \to W$ is also called an isomorphism from V to W.
(8) If an isomorphism from V to W exists, then V and W are called isomorphic and one writes $V \cong W$.
(9) A linear map $F\colon V \to V$ is also called an endomorphism of V or a linear operator (or just an operator) on V. If it is invertible, it is also called an automorphism. The identity map, denoted by Id or Id_V or 1, is an automorphism.
(10) If F is an endomorphism of V, a subspace W of V is called F-invariant if $F(W) \subseteq W$ (i.e., $F(w) \in W$ for every $w \in W$). The restriction of F to an invariant subspace W then yields an endomorphism of W.

We introduce the following sets of linear maps and their additional structures:

(1) $\mathrm{Hom}(V, W)$ is the set of all homomorphisms from V to W. If $F, G \in \mathrm{Hom}(V, W)$, we define $F + G \in \mathrm{Hom}(V, W)$ by

$$(F + G)(v) := F(v) + G(v).$$

We denote by 0 the zero homomorphism $0(v) = 0$ for all $v \in V$. With the scalar multiplication $(\lambda F)(v) := \lambda F(v)$, the set $\mathrm{Hom}(V, W)$ is a vector space over $\mathbb{K}$.
(2) $\mathrm{End}(V) = \mathrm{Hom}(V, V)$ is the set of all endomorphisms of V. As a particular case of the above, it is a vector space over $\mathbb{K}$. It is also a ring with one, where the multiplication is given by the composition and the one element is the identity map.
(3) $\mathrm{Aut}(V) \subset \mathrm{End}(V)$ is the set of all automorphisms of V. It is a group with multiplication given by composition.

Remark 1.30 (Injective linear maps). Note that a linear map F is injective iff $\ker F = 0$. In fact, if F is injective, then $Fv = 0 = F0$

implies $v = 0$. On the other hand, the equality $Fv = Fv'$ implies, by linearity, that $v - v' \in \ker F$, so if $\ker F = 0$ then we have $v = v'$.

DEFINITION 1.31 (Dual space). The vector space $\mathrm{Hom}(V, \mathbb{K})$ is usually denoted by V^* and called the dual space of V. Note that V^*, like every Hom space, is itself a vector space. An element α of V^* is a linear map $V \to \mathbb{K}$ and is usually called a linear functional. In addition to the notation $\alpha(v)$ to indicate $\alpha \in V^*$ evaluated on $v \in V$, one often writes (α, v).[6]

EXAMPLE 1.32. A row vector $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n) \in (\mathbb{K}^n)^*$ defines a linear functional on $\mathbb{K}^n$ via

$$\begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix} \mapsto (\boldsymbol{\alpha}, \boldsymbol{v}) := \sum_{i=1}^n \alpha_i v^i. \tag{1.1}$$

One can show that these are all possible linear functionals on $\mathbb{K}^n$, so $(\mathbb{K}^n)^*$ is the dual of $\mathbb{K}^n$, which justifies the notation.

DEFINITION 1.33 (Einstein's convention). In these notes we will follow a very handy convention introduced by A. EINSTEIN according to which, whenever an index appears twice in an expression, once as a lower and once as an upper index, then a sum over that index is understood.

EXAMPLE 1.34. According to Einstein's convention, the evaluation of a row vector $\boldsymbol{\alpha}$ on a column vector $\boldsymbol{v}$, as in equation (1.1), is simply written as

$$(\boldsymbol{\alpha}, \boldsymbol{v}) = \alpha_i v^i.$$

DEFINITION 1.35 (Bidual space). The dual space of the dual space V^* of a vector space V is denoted by V^{**} and is called the bidual space of V.

REMARK 1.36. Note that every $v \in V$ defines a linear functional on V^* by $V^* \ni \alpha \mapsto v(\alpha) := \alpha(v)$. Therefore, we may regard V as a subspace of V^{**}. We will see (Proposition 1.62) that, if V is finite-dimensional, one actually has $V = V^{**}$.

REMARK 1.37 (Direct sum). Suppose $V = V_1 \oplus V_2$ as in Example 1.28.(2). Then we have the following relations among the projections:

$$P_1 + P_2 = 1, \quad P_1^2 = P_2, \quad P_2^2 = P_2, \quad P_1 P_2 = P_2 P_1 = 0.$$

[6] This notation actually indicates the induced bilinear map $V^* \times V \to \mathbb{K}$.

Remark 1.38. More generally, an endomorphism P of V is called a projection if

$$P^2 = P.$$

Note that $Q := 1-P$ is also a projection and that we have $PQ = QP = 0$. This is related to the previous remark as follows: define $V_1 := \operatorname{im} P$ and $V_2 := \operatorname{im} Q$. Then we have $V = V_1 \oplus V_2$ and $P_1 = P$ and $P_2 = Q$.

Remark 1.39 (The dual map). A linear map $F\colon V \to W$ induces a dual map $F^*\colon W^* \to V^*$ as follows: an element α of W^*—i.e., a linear functional on W—is mapped to $F^*\alpha \in V^*$ defined by

$$(F^*\alpha)(v) := \alpha(Fv).$$

Note that F^* is also linear. Moreover, F^{**} restricted to V is the map F again. We will see in Section 1.5.2 that the dual of a map is related to the transposition of matrices.

1.4. Bases

A basis of a $\mathbb{K}$-vector space V is a collection $(e_i)_{i\in S}$ of elements of V such that for every vector $v \in V$ there are are uniquely determined scalars $v^i \in \mathbb{K}$, only finitely many of which are different from zero, such that

$$v = \sum_i v^i e_i.$$

Note that, by omitting the zero summands, this is a sum of finitely many terms. The scalars v^i are called the components of v in the given basis. Using Einstein's convention (see Definition 1.33), we write the expansion of v in the basis $(e_i)_{i\in S}$ as

$$v = v^i e_i.$$

Remark 1.40. In order to use Einstein's convention, one has to be consistent with the positioning of the indices. Typically we will use lower indices for basis elements and, consequently, upper indices for components of vectors. In some cases, see below, we use upper indices for basis elements and, consequently, lower indices for components of vectors.[7]

Example 1.41 (The standard bases). The space $\mathbb{K}^n$ has the standard basis $(\boldsymbol{e}_1, \dots, \boldsymbol{e}_n)$ where $\boldsymbol{e}_i$ denotes the column vector that has a 1

[7]As will be explained later, we use upper indices for a basis of a dual space.

in the ith position and a 0 otherwise. The space $(\mathbb{K}^n)^*$ also has a standard basis, now with upper indices $(\boldsymbol{e}^1, \dots, \boldsymbol{e}^n)$, where $\boldsymbol{e}^i$ denotes the row vector that has a 1 in the ith position and a 0 otherwise. Namely:

$$\boldsymbol{v} = \begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix} \qquad \Longrightarrow \qquad \boldsymbol{v} = v^i \boldsymbol{e}_i$$

$$\boldsymbol{v} = \big(v_1, \dots, v_n\big) \qquad \Longrightarrow \qquad \boldsymbol{v} = v_i \boldsymbol{e}^i$$

Remark 1.42 (Related concepts). Some important concepts are related to the notion of basis. All the vectors mentioned in the following list belong to a fixed $\mathbb{K}$-vector space V.

(1) A linear combination of a finite collection $v_1, \dots, v_k$ of vectors is a vector of the form $\lambda^i v_i$ with $\lambda^i \in \mathbb{K}$. The set

$$\operatorname{Span}\{v_1, \dots, v_k\} := \{\lambda^i v_i,\ \lambda_i \in \mathbb{K}\}$$

of all linear combinations of $v_1, \dots, v_k$ is called their span and is a subspace of V. If the set contains a single vector v, instead of $\operatorname{Span}\{v\}$ we also use the notation $\mathbb{K}v$, so

$$\mathbb{K}v = \{\lambda v,\ \lambda \in \mathbb{K}\}.$$

(2) A linear combination of a collection $(e_i)_{i\in S}$ of vectors is by definition a linear combination of a finite subcollection. When writing $\lambda^i v_i$, it is then assumed that only finitely many λ_is are different from zero. The set $\operatorname{Span}_{i\in S} e_i$ of linear combinations of the vectors in the collection is also a subspace of V.

(3) A collection $(e_i)_{i\in S}$ of vectors is called a system of generators for V if every vector of V can be expressed as a linear combination of the e_is (but we do not require uniqueness of this expression). In other words, $\operatorname{Span}_{i\in S} e_i = V$. Each e_i is called a generator.

(4) A collection $(e_i)_{i\in S}$ of vectors is called linearly independent if a linear combination can be zero only if all the coefficients are zero. That is, if

$$\lambda^i e_i = 0 \implies \lambda_i = 0\ \forall i.$$

The collection is called linearly dependent otherwise.

(5) A basis is then the same as a linearly independent system of generators.

One can esily see that the following hold:

(1) If $F\colon V \to W$ is injective and $(e_i)_{i\in S}$ is a linearly independent family of vectors in V, then $(Fe_i)_{i\in S}$ is a linearly independent family of vectors in W.

(2) If $F\colon V \to W$ is surjective and $(e_i)_{i\in S}$ is a system of generators in V, then $(Fe_i)_{i\in S}$ is a system of generators in W.

Therefore,

PROPOSITION 1.43. *If $F\colon V \to W$ is an isomorphism and $(e_i)_{i\in S}$ is a basis of V, then $(Fe_i)_{i\in S}$ is a basis of W.*

Moreover, one has the

THEOREM 1.44. *Any two bases of the same vector space have the same cardinality.*

If V admits a finite basis (i.e., a basis $(e_i)_{i\in S}$ with S a finite set), then it is called a finite-dimensional vector space; otherwise, it is called an infinite-dimensional vector space.

DIGRESSION 1.45 (Existence of bases). By definition a finite-dimensional vector space has a basis (actually, a finite one). By the axiom of choice one can prove that every vector space has a basis (actually, the existence of bases for all vector spaces is equivalent to the axiom of choice).

DEFINITION 1.46 (Dimension). If V is finite-dimensional with a basis of cardinality n (i.e., $|S| = n$), then we set

$$\dim V = n$$

and call it the dimension of V. Usually we then choose $S = \{1, \dots, n\}$ and denote the basis by $(e_1, \dots, e_n)$. Note that the trivial vector space $V = \{0\}$ is zero-dimensional. In particular, we have

$$\dim \mathbb{K}^n = \dim(\mathbb{K}^n)^* = n \quad \forall n \in \mathbb{N}.$$

If V is infinite-dimensional, we set

$$\dim V = \infty.$$

REMARK 1.47 (Dimension over a field). The same abelian group V may sometimes be regarded as a vector space over different fields $\mathbb{K}$. In this case, it is convenient to to remember which field we are considering when computing the dimension: we will write $\dim_{\mathbb{K}} V$ for the dimension of V as a $\mathbb{K}$-vector space.

REMARK 1.48 (Complex spaces as real spaces). In particular, a case we will often encounter is that of a complex vector space V (i.e., a vector space over $\mathbb{C}$). For every real λ, we still have the scalar multiplication $v \mapsto \lambda v$, so V may also be regarded as a real vector space (i.e., a vector space over $\mathbb{R}$). If $\mathcal{B}_\mathbb{C} = (e_1, \dots, e_n)$ is a basis of V as a complex vector space, then[8]

$$\mathcal{B}_\mathbb{R} = (e_1, \dots, e_n, \mathrm{i}e_1, \dots, \mathrm{i}e_n)$$

is a basis of V as a real one. In fact, every $v \in V$ may uniquely be expanded as $\lambda^i e_i$ with $\lambda^i \in \mathbb{C}$. Writing $\lambda^i = a^i + \mathrm{i}b^i$, with a^i and b^i real, we get the expansion $v = a^i e_i + b^i \mathrm{i}e_i$, so $\mathcal{B}_\mathbb{R}$ is a system of generators over $\mathbb{R}$.[9] Moreover, if $a^i e_i + b^i \mathrm{i}e_i = 0$ for $a^i, b^i \in \mathbb{R}$, then $\lambda^i e_i = 0$; linear independence of $\mathcal{B}_\mathbb{C}$ over $\mathbb{C}$ implies, for all i, $\lambda^i = 0$ and, therefore, $a^i = b^i = 0$, which is linear independence of $\mathcal{B}_\mathbb{R}$ over $\mathbb{R}$. Therefore, $\mathcal{B}_\mathbb{R}$ is basis. We conclude that $\dim_\mathbb{C} V = n$ implies that $\dim_\mathbb{R} V = 2n$; i.e.,

$$\dim_\mathbb{R} V = 2 \dim_\mathbb{C} V.$$

One can prove that a linearly independent collection $(e_1, \dots, e_n)$ in an n-dimensional space is automatically a basis. This in particular implies the following

PROPOSITION 1.49. *If V is a finite-dimensional vector space and W is a subspace of V of the same dimension, then $W = V$.*

REMARK 1.50 (Direct sums and bases). There is a strong relationship between the notions of direct sums and bases. Namely, by Definition 1.21, a collection $(v_1, \dots, v_n)$ is a basis of V iff $V = \bigoplus_{i=1}^n \mathbb{K}v_i$.

REMARK 1.51 (Union of bases). Another relation is the following. If $V = \bigoplus_{i=1}^k W_i$ and $\mathcal{B}_i = (v_{i,j})_{j=1,\dots,d_i}$ is a basis of W_i, then

$$\mathcal{B} := \cup_{i=1}^k \mathcal{B}_i = (v_{i,j})_{\substack{i=1,\dots,k \\ j=1,\dots,d_i}}$$

is a basis of V. As a consequence,

$$\dim \bigoplus_{i=1}^k W_i = \sum_{i=1}^k \dim W_i. \tag{1.2}$$

PROOF. By Definition 1.21, every $v \in V$ uniquely decomposes as $v = w_1 + \dots + w_k$ with $w_i \in W_i$. By definition of basis, every w_i uniquely decomposes as $w_i = \sum_{j=1}^{d_i} \alpha_{i,j} v_{i,j}$. Therefore, v uniquely decomposes as $v = \sum_{i=1}^k \sum_{j=1}^{d_i} \alpha_{i,j} v_{i,j}$. □

[8] Here i denotes the imaginary unit, and $\mathrm{i}e_j$ is the scalar multiplication of the scalar i with the vector e_j.

[9] A standard convention in mathematics is to use italic characters for variables and roman characters for constants. The imaginary unit, being a constant, is then denoted by i, whereas i is a variable, like, e.g., the index in e_i. By handwriting it is however better to avoid using a variable i when the imaginary unit also appers.

REMARK 1.52 (The basis isomorphism). A basis $\mathcal{B} = (e_1, \dots, e_n)$ of V determines a linear map

$$\phi_{\mathcal{B}} \colon \mathbb{K}^n \to V$$

by

$$\phi_{\mathcal{B}} \begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix} := v^i e_i.$$

In particular, we have

$$\phi_{\mathcal{B}} \boldsymbol{e}_i = e_i$$

for all i. Note that $\phi_{\mathcal{B}}$ is an isomorphism with inverse $\phi_{\mathcal{B}}^{-1} \colon V \to \mathbb{K}^n$ the map that sends $v \in V$ to the column vector with components the components of v in the basis $\mathcal{B}$. The vector $\phi_{\mathcal{B}}^{-1} v$ is called the coordinate vector of v. Note that we then have

$$\dim V = n \iff V \cong \mathbb{K}^n.$$

REMARK 1.53. If V and W have the same finite dimension, then, by composition of the above, we get an isomorphism $V \to W$. Note that such an isomorphism depends on the choice of bases.

REMARK 1.54 (Bases and frames). When we define the isomorphism of Remark 1.52, the order in which we take the basis vectors $(e_1, \dots, e_n)$ matters. This is an additional choice to just the notion of a basis (which is by definition a collection, i.e., a set, of linearly independent generators). A basis with a choice of ordering is more precisely called a frame. According to a general habit, we will be sloppy about it and speak of a basis also when we actually mean a frame as, e.g., in Remark 1.52.

An immediate consequence of the basis isomorphism is the following

PROPOSITION 1.55. *Every subspace of a finite-dimensional vector space has a complement.*

PROOF. Let V be n-dimensional, and let Z be a subspace of V. By choosing a basis $\mathcal{B}$, we have the isomorphism $\phi_{\mathcal{B}} \colon \mathbb{K}^n \to V$. Then $W := \phi_{\mathcal{B}}^{-1}(Z)$ is a subspace of $\mathbb{K}^n$. By Lemma 1.27, it has a complement W'. Finally, $\phi_{\mathcal{B}}(W')$ is a complement of Z. □

DIGRESSION 1.56. By the axiom of choice, one can show that a subspace of any vector space has a complement.

DEFINITION 1.57 (Change of basis). If $\mathcal{B}$ and $\mathcal{B}'$ are bases of an n-dimensional space V, the composition

$$\phi_{\mathcal{B}'\mathcal{B}} := \phi_{\mathcal{B}'}^{-1} \phi_{\mathcal{B}} \in \operatorname{Aut}(\mathbb{K}^n)$$

is called the corresponding change of basis.

REMARK 1.58. If you have a vector $v \in V$, then $\phi_{\mathcal{B}}^{-1}v$ is the column vector of its components in the basis $\mathcal{B}$. The column vector $\phi_{\mathcal{B}'}^{-1}v$ of its components in the basis $\mathcal{B}'$ is then related to $\phi_{\mathcal{B}}^{-1}v$ by

$$\phi_{\mathcal{B}'}^{-1}v = \phi_{\mathcal{B}'\mathcal{B}}\phi_{\mathcal{B}}^{-1}v.$$

Therefore, $\phi_{\mathcal{B}'\mathcal{B}}$ maps the coordinate vector in the $\mathcal{B}$ basis to the coordinate vector in the $\mathcal{B}'$ basis. (A more descriptive, but also more cumbersome, notation would be $\phi_{\mathcal{B}'\leftarrow\mathcal{B}}$ instead of $\phi_{\mathcal{B}'\mathcal{B}}$.)

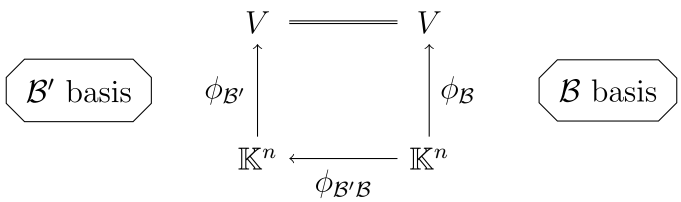

REMARK 1.59 (The dual basis). A basis $\mathcal{B} = (e_1, \dots, e_n)$ of V allows defining uniquely the components v^i of any vector v. The map

$$\begin{array}{rccc} e^i\colon & V & \to & \mathbb{K} \\ & v & \mapsto & v^i \end{array}$$

is linear for every i. The collection $\mathcal{B}^* := (e^1, \dots, e^n)$ of linear functionals is called the dual basis of V^*—more precisely, the basis of V^* dual to $\mathcal{B}$—and satisfies, by definition,

$$e^i(e_j) = \delta^i_j := \begin{cases} 1 & \text{if } i = j, \\ 0 & \text{if } i \neq j, \end{cases}$$

where δ^i_j is called the Kronecker delta. It is indeed a basis of V^*. In fact, if $\alpha := \alpha_i e^i = 0$, then $0 = \alpha(e_j) = \alpha_j$ for every j, so $(e^1, \dots, e^n)$ is linearly independent. Moreover, given any $\alpha \in V^*$, we get

$$\alpha(v) = \alpha(v^i e_i) = v^i \alpha(e_i) = e^i(v)\alpha(e_i) = (\alpha(e_i)e^i)(v),$$

which shows that $(e^1, \dots, e^n)$ is a system of generators and that, in particular, we can compute the component α_i of α as $\alpha(e_i)$.

EXAMPLE 1.60. The basis $(\boldsymbol{e}^1, \dots, \boldsymbol{e}^n)$ of $(\mathbb{K}^n)^*$ is the dual basis to $(\boldsymbol{e}_1, \dots, \boldsymbol{e}_n)$.

More generally, we have proved the

PROPOSITION 1.61. *If V is a finite-dimensional vector space, then*

$$\dim V^* = \dim V.$$

This also implies that $\dim V^{**} = \dim V$. As a consequence of Remark 1.36 and of Proposition 1.49, we get the

PROPOSITION 1.62. *If V is a finite-dimensional vector space, then*

$$V^{**} = V.$$

REMARK 1.63 (Canonical and noncanonical maps). A linear map is called canonical if it does not depend on any additional structure. For example, if W is a subspace of V, the inclusion map is canonical. If $V = W_1 \oplus W_2$, the projections from V to W_1 and W_2 are also canonical. If W is a subspace of V, we can always find a complement W', so we can write $V = W \oplus W'$ and, therefore, get a projection $V \to W$. This projection is not canonical because it depends on the choice of a complement. Similarly, we saw that that every element V defines a linear functional on V^*, so we have a canonical inclusion map of V into V^{**}. If V is finite dimensional, we then have a canonical isomorphism between V and V^{**}: we therefore write $V = V^{**}$. On the other hand, by Proposition 1.61 and Remark 1.53, we also have an isomorphism between V and V^*, but this is not canonical because it depends on the choice of a basis. Explicitly, the map $V \to V^*$ sends $v^i e_i$ to $\sum_{i=1}^n v^i e^i$ (not that we cannot use Einstein's convention in this case).

EXAMPLE 1.64. With respect to the standard basis, the isomorphism $\mathbb{K}^n \xrightarrow{\sim} (\mathbb{K}^n)^*$, as at the end of the previous remark, is the transposition map $\boldsymbol{v} \mapsto \boldsymbol{v}^{\mathsf{T}}$:

$$\begin{pmatrix} v^1 \\ \vdots \\ v^n \end{pmatrix} \mapsto (v^1, \ldots, v^n).$$

1.5. Representing matrices

If $F\colon V \to W$ is a linear map and $(e_i)_{i\in S}$ is a basis of V, then the values of F on the e_is are enough to reconstruct F; in fact, every $v \in V$ is uniquely expanded as $v^i e_i$, so by linearity we get

$$F(v) = v^i F(e_i).$$

Vice versa, we can define a linear map $F\colon V \to W$ by specifying $w_i := F(e_i) \in W$ for all $i \in S$.

If $(\bar{e}_{\bar{\imath}})_{\bar{\imath}\in\bar{S}}$ is a basis of W, we may expand $F(e_i)$ as

$$F(e_i) = A_i^{\bar{\imath}}\, \bar{e}_{\bar{\imath}}, \tag{1.3}$$

where, by Einstein's convention, a sum over $\bar{\imath}$ is understood. The scalars $A_i^{\bar{\imath}}$ are the entries of the representing matrix $\boldsymbol{A}$ of F. On a generic vector

$v = v^i e_i \in V$ we then get

$$F(v) = v^i A_i^{\bar{\imath}}\, \bar{e}_{\bar{\imath}}.$$

We now assume that V and W are finite-dimensional: say, $\dim V = n$ and $\dim W = m$. Then

$$\boldsymbol{A} = (A_i^{\bar{\imath}})_{\substack{i=1,\dots,n\\ \bar{\imath}=1,\dots,m}} = \begin{pmatrix} A_1^1 & A_2^1 & \cdots & A_n^1 \\ A_1^2 & A_2^2 & \cdots & A_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ A_1^m & A_2^m & \cdots & A_n^m \end{pmatrix}$$

is called an $m \times n$ matrix.

Remark 1.65. Note that we consistently put the indices in upper and lower positions. We will also encounter matrices with only lower indices representing bilinear forms.

Remark 1.66 (Normal form). Given a linear map $F\colon V \to W$ between finite-dimensional vector spaces, one can always find bases $e_1, \dots, e_n$ of V and $\bar{e}_1, \dots, \bar{e}_m$ of W such that the representing matrix of F reads

$$\boldsymbol{A} = \begin{pmatrix} \mathbf{1}_r & \mathbf{0}_{r,n-r} \\ \mathbf{0}_{m-r,r} & \mathbf{0}_{m-r,n-r} \end{pmatrix},$$

where $\mathbf{1}_r$ denotes the $r \times r$ identity matrix and $\mathbf{0}_{i,j}$ denotes the $i \times j$ zero matrix. Here $r = \dim \operatorname{im} F$ is called the rank of F. From this presentation it also follows that the kernel of F corresponds, under the isomorphism to $\mathbb{K}^n$ induced by the basis, to vectors of the form

$$\begin{pmatrix} 0 \\ \vdots \\ 0 \\ v^{r+1} \\ \vdots \\ v^n \end{pmatrix},$$

which show that $\dim \ker F = n - r$. We thus get the dimension formula

$$\dim \ker F + \dim \operatorname{im} F = \dim V \tag{1.4}$$

for any linear map $F\colon V \to W$.

If $F\colon V \to W$ is an isomorphism, then $\ker F = 0$ (see Remark 1.30) and $\operatorname{im} F = W$, so we get $\dim V = \dim W$. If, on the other hand, V and W have the same finite dimension n, then each of them is isomorphic to $\mathbb{K}^n$ by the basis isomorphism of Remark 1.52. Therefore,

Proposition 1.67. *Two finite-dimensional vector spaces are isomorphic iff they have the same dimension.*

1.5.1. Operations. If $\boldsymbol{B}$ is the representing matrix of another linear map $G\colon V \to W$, with respect to the same bases, the representing matrix of $F+G$ is $\boldsymbol{A}+\boldsymbol{B}$, where addition is defined entry-wise:

$$(A+B)_i^{\bar{\imath}} = A_i^{\bar{\imath}} + B_i^{\bar{\imath}}.$$

If G is instead a linear map $W \to Z$, the collection $(\tilde{e}_{\tilde{\imath}})_{\tilde{\imath}=1,\dots,k}$ is a basis of Z, and $\boldsymbol{B}$ is the corresponding $k \times m$ representing matrix of G, then the representing matrix of GF is the $k \times n$ matrix $\boldsymbol{BA}$, where the matrix multiplication is defined by[10]

$$(BA)_i^{\tilde{\imath}} = B_{\bar{\imath}}^{\tilde{\imath}} A_i^{\bar{\imath}}.$$

1.5.2. Duals. If $F^*\colon W^* \to V^*$ is the dual map, see Remark 1.39, of a linear map $F\colon V \to W$ with representing matrix $\boldsymbol{A}$, then we get, for $\alpha = \alpha_{\bar{\jmath}}\,\bar{e}^{\bar{\jmath}} \in W^*$ and $v = v^i e_i \in V$,

$$F^*(\alpha)(v) = \alpha(Fv) = \alpha_{\bar{\jmath}}\,\bar{e}^{\bar{\jmath}}(v^i A_i^{\bar{\imath}}\,\bar{e}_{\bar{\imath}}) = \alpha_{\bar{\imath}}\,A_i^{\bar{\imath}}\,v^i,$$

where $\bar{e}^1,\dots,\bar{e}^m$ is the dual basis to $\bar{e}_1,\dots,\bar{e}_m$. To compute the ith component of $F^*\alpha$, we evaluate on $v = e_i$ getting $(F^*\alpha)_i = \alpha_{\bar{\imath}}\,A_i^{\bar{\imath}}$, so

$$F^*(\alpha) = \alpha_{\bar{\imath}} A_i^{\bar{\imath}}\,e^i.$$

To get the representing matrix of F^*, we then compute

$$F^*(\bar{e}^{\bar{\imath}}) = A_i^{\bar{\imath}}\,e^i. \tag{1.5}$$

The difference between (1.3) and (1.5) is that in the former we sum over the upper index whereas in the latter we sum over the lower index of the matrix $\boldsymbol{A}$. In the usual notation with two lower indices, this correponds to summing over the first or second index, and the two representing matrices are related by the exchange of the indices—the operation known as transposition: if $\boldsymbol{A}$ is the representing matrix of F, with respect to some bases, then $\boldsymbol{A}^\mathsf{T}$ is the representing matrix of F^* with respect to the dual bases.

Remark 1.68. When working with vector spaces of row or column vectors, the standard basis is always assumed, unless otherwise specified. A linear map between such subspaces is then always understood as the corresponding matrix. We would in particular write $\boldsymbol{A}\colon \mathbb{K}^n \to \mathbb{K}^m$ to denote the linear map with representing matrix $\boldsymbol{A}$ with respect to the standard bases, i.e., the map $\boldsymbol{v} \mapsto \boldsymbol{Av}$, where we use matrix multiplication. Note that the jth row of $\boldsymbol{A}$ is equal to $\boldsymbol{Ae}_j$. One can also easily see that the dual map, acting on row vectors, is instead given by $\boldsymbol{\alpha} \mapsto \boldsymbol{\alpha A}$.

[10] Comparing with the usual way of writing a product of matrices, we see here that the upper index is the first index and the lower index is the second index.

1.5.3. Change of basis. If we want to keep track of the chosen bases, we need a more descriptive notation. Let $\mathcal{B} = (e_1, \dots, e_n)$ be the chosen basis of V and $\bar{\mathcal{B}} = (\bar{e}_1, \dots, \bar{e}_m)$ the chosen basis of W. The representing matrix of a linear map $F\colon V \to W$ with respect to the bases $\mathcal{B}$ and $\bar{\mathcal{B}}$ is then denoted by $F_{\bar{\mathcal{B}}\mathcal{B}}$.

As we identify linear maps $\mathbb{K}^n \to \mathbb{K}^m$ with their representing matrices with respect to the standard bases, we may regard $F_{\bar{\mathcal{B}}\mathcal{B}}$ as an $m \times n$ matrix or, equivalently, as a linear map. In the latter case, we have, in terms of the basis isomorphisms of Remark 1.52,

$$F_{\bar{\mathcal{B}}\mathcal{B}} = \phi_{\bar{\mathcal{B}}}^{-1} F \phi_{\mathcal{B}}.$$

If we choose another basis $\mathcal{B}'$ of V and another basis $\bar{\mathcal{B}}'$ of W, we then have, using the notation of Definition 1.57,

$$F_{\bar{\mathcal{B}}'\mathcal{B}'} = \phi_{\bar{\mathcal{B}}'}^{-1} F \phi_{\mathcal{B}'} = \phi_{\bar{\mathcal{B}}'}^{-1} \phi_{\bar{\mathcal{B}}} F_{\bar{\mathcal{B}}\mathcal{B}} \phi_{\mathcal{B}}^{-1} \phi_{\mathcal{B}'} = \phi_{\bar{\mathcal{B}}'\bar{\mathcal{B}}} F_{\bar{\mathcal{B}}\mathcal{B}} \phi_{\mathcal{B}\mathcal{B}'},$$

i.e.,

$$F_{\bar{\mathcal{B}}'\mathcal{B}'} = \phi_{\bar{\mathcal{B}}'\bar{\mathcal{B}}} F_{\bar{\mathcal{B}}\mathcal{B}} \phi_{\mathcal{B}\mathcal{B}'}.$$

It is easy to remember this formula, for it looks similar to the formula for matrix product, with indices replaced by bases.

Remark 1.69 (Equivalence of matrices). The formula for the change of bases motivates the following definition of equivalence of matrices. Two $m \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are called equivalent if there is an invertible $m \times m$ matrix $\boldsymbol{T}$ and an invertible $n \times n$ matrix $\boldsymbol{S}$ such that

$$\boldsymbol{A} = \boldsymbol{T}\boldsymbol{B}\boldsymbol{S}.$$

By this definition, any two representing matrices of the same linear map are equivalent, as we may see by setting $\boldsymbol{A} = F_{\bar{\mathcal{B}}'\mathcal{B}'}$, $\boldsymbol{B} = F_{\bar{\mathcal{B}}\mathcal{B}}$, $\boldsymbol{S} = \phi_{\mathcal{B}\mathcal{B}'}$, and $\boldsymbol{T} = \phi_{\bar{\mathcal{B}}'\bar{\mathcal{B}}}$. Note that, explicitly, we have

$$A^{\bar{\imath}'}_{i'} = T^{\bar{\imath}'}_{\bar{\imath}} B^{\bar{\imath}}_{i} S^{i}_{i'}.$$

Also useful are the formulae

$$e'_{i'} = S^i_{i'} e_i \quad \text{and} \quad \bar{e}_{\bar{\imath}} = T^{\bar{\imath}'}_{\bar{\imath}} \bar{e}'_{\bar{\imath}'}.$$

Let us prove the first (the second is analogous):

$$S^i_{i'} e_i = S^i_{i'} \phi_{\mathcal{B}} \boldsymbol{e}_i = \phi_{\mathcal{B}}(S^i_{i'} \boldsymbol{e}_i) = \phi_{\mathcal{B}}(\phi_{\mathcal{B}\mathcal{B}'} \boldsymbol{e}_{i'}) = \phi_{\mathcal{B}'} \boldsymbol{e}_{i'} = e'_{i'},$$

where we also used the linearity of $\phi_{\mathcal{B}}$ and the definition $\phi_{\mathcal{B}\mathcal{B}'} = \phi_{\mathcal{B}}^{-1} \phi_{\mathcal{B}'}$.

1.5.4. Endomorphisms. If F is an endomorphism of a finite-dimensional vector space V, one usually chooses the same basis (say, $\mathcal{B}$) for V as source and target space. In this case, the representing matrix of V, now a square matrix,[11] with respect to the basis $\mathcal{B}$ is written $F_\mathcal{B}$ and we have

$$F_\mathcal{B} = \phi_\mathcal{B}^{-1} F \phi_\mathcal{B}.$$

If we pass to another basis (say, $\mathcal{B}'$), we then have

$$F_{\mathcal{B}'} = \phi_{\mathcal{B}'\mathcal{B}} F_\mathcal{B} \phi_{\mathcal{B}\mathcal{B}'}.$$

Observing that the isomorphisms $\phi_{\mathcal{B}'\mathcal{B}}$ and $\phi_{\mathcal{B}\mathcal{B}'}$ are inverse to each other, we can also write

$$F_{\mathcal{B}'} = \phi_{\mathcal{B}\mathcal{B}'}^{-1} F_\mathcal{B} \phi_{\mathcal{B}\mathcal{B}'}.$$

REMARK 1.70 (Similarity of matrices). The formula for the change of basis for the representing matrix of an endomorphism motivates the following definition of similarity of matrices. Two $n \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ are called similar if there is an invertible $n \times n$ matrix $\boldsymbol{S}$ such that

$$\boldsymbol{A} = \boldsymbol{S}^{-1} \boldsymbol{B} \boldsymbol{S}.$$

By this definition, any two representing matrices of the same endomorphism are similar.

1.5.5. Bilinear forms. A bilinear form on a $\mathbb{K}$-vector space V is a map $B\colon V \times V \to \mathbb{K}$ that is linear in both arguments; viz.,

$$B(\lambda_1 v_1 + \lambda_2 v_2, w) = \lambda_1 B(v_1, w) + \lambda_2 B(v_2, w),$$
$$B(v, \lambda_1 w_1 + \lambda_2 w_2) = \lambda_1 B(v, w_1) + \lambda_2 B(v, w_2).$$

If $\mathcal{B} = \{e_1, \dots, e_n\}$ is a basis of V, the representing matrix $\boldsymbol{B}$ of a bilinear form B has the entries

$$B_{ij} := B(e_i, e_j). \tag{1.6}$$

Note that, consistently with the r.h.s., we use lower indices for the entries of the representing matrix.

Now consider a new basis $\mathcal{B}' = \{e'_1, \dots, e'_n\}$ and denote by $\boldsymbol{S}$ the matrix representing, in the standard basis of $\mathbb{K}^n$, the change of basis $\phi_{\mathcal{B}\mathcal{B}'}$; i.e., as we showed above,

$$e'_{i'} = S^i_{i'} e_i.$$

Denoting by $\boldsymbol{B}'$ the representing matrix of B in the basis $\mathcal{B}'$, we then have

$$B'_{i'j'} = S^i_{i'} B_{ij} S^j_{j'}.$$

[11] i.e., a matrix with the same number of rows and columns

REMARK 1.71 (Congruency of matrices). The formula for the change of basis of a bilinear form motivates the following definition of congruency of matrices. Two $n \times n$ matrices $\boldsymbol{B}$ and $\boldsymbol{B}'$ are called congruent if there is an invertible $n \times n$ matrix $\boldsymbol{S}$ such that

$$\boldsymbol{B}' = \boldsymbol{S}^{\mathsf{T}} \boldsymbol{B} \boldsymbol{S},$$

where $^{\mathsf{T}}$ denotes transposition. By this definition, any two representing matrices of the same bilinear form are congruent.[12]

1.6. Traces and determinants

Trace and determinant are particularly important functions on square matrices which we review here.

The trace of an $n \times n$ matrix $\boldsymbol{A}$ is the sum of its diagonal elements. If we write $\boldsymbol{A} = (A^i_j)$, then

$$\operatorname{tr} \boldsymbol{A} = A^i_i,$$

where we used Einstein's convention. If we we write $\boldsymbol{A} = A_{ij}$, then we have to write the sum symbol explicitly: $\operatorname{tr} \boldsymbol{A} = \sum_{i=1}^n A_{ii}$. (The fact that this second notation is incompatible with Einstein's convention is related to the fact that the trace of a bilinear form is not a natural operation.)

Immediate properties of the trace are

(1) $\operatorname{tr} \boldsymbol{A}^{\mathsf{T}} = \operatorname{tr} \boldsymbol{A}$ for every $n \times n$ matrix, and
(2) $\operatorname{tr} \boldsymbol{AB} = \operatorname{tr} \boldsymbol{BA}$ for all $n \times n$ matrices $\boldsymbol{A}, \boldsymbol{B}$.

The second property implies $\operatorname{tr} \boldsymbol{S}^{-1}\boldsymbol{B}\boldsymbol{S} = \operatorname{tr} \boldsymbol{B}$, so similar matrices—see Remark 1.70—have the same trace. Therefore, we can make the following

DEFINITION 1.72. The trace of an endomorphism F of a finite-dimensional vector space V is the trace of any of its representing matrices: $\operatorname{tr} F := \operatorname{tr} F_{\mathcal{B}}$, where $\mathcal{B}$ is any basis of V.

Congruent matrices may on the other hand have different traces. For this reason the trace of a bilinear form is not a well-defined concept, as it depends on the explicit choice of a basis.[13]

We conclude with a few additional properties of the trace:

---

[12] Note that in the explicit formula (1.6) it is the upper index of the first $\boldsymbol{S}$ that is the same as the first index of $\boldsymbol{B}$, whereas in the usual product of matrices—see footnote 10—it should be the lower index to be involved. It is for this reason that the first $\boldsymbol{S}$ is actually transposed.

[13] It is well-defined if we may restrict to a special class of bases that are related to each other by an orthogonal transformation, i.e., if we only allow congruences $\boldsymbol{B}' = \boldsymbol{S}^{\mathsf{T}}\boldsymbol{B}\boldsymbol{S}$ with $\boldsymbol{S}\boldsymbol{S}^{\mathsf{T}} = \mathbf{1}$.

(1) The trace is linear: $\operatorname{tr}(\lambda \boldsymbol{A}+\lambda \boldsymbol{B}) = \lambda \operatorname{tr} \boldsymbol{A} + \mu \operatorname{tr} \boldsymbol{B}$ for all $n \times n$ matrices $\boldsymbol{A}, \boldsymbol{B}$ and for all scalars λ, μ.
(2) $\operatorname{tr} \mathbf{1}_n = n$ if $\mathbf{1}_n$ is the identity matrix on $\mathbb{K}^n$.
(3) $\operatorname{tr}\colon \operatorname{End}(V) \to \mathbb{K}$ is a linear map: $\operatorname{tr}(\lambda F + \lambda G) = \lambda \operatorname{tr} F + \mu \operatorname{tr} G$ for all endomorphisms F, G of the finite-dimensional space V and for all scalars λ, μ.
(4) $\operatorname{tr} \operatorname{Id}_V = \dim V$.

The **determinant** of a square matrix can be uniquely characterized by some properties or can be, equivalently, defined by an explicit formula. The defining properties are the following:

(1) The determinant is linear with respect to every column of the matrix.
(2) The determinant vanishes if any two column of the matrix are equal.
(3) The determinant of the identity matrix is 1.

One can show that, for every n, there is a unique map $\operatorname{Mat}_{n\times n}(\mathbb{K}) \to \mathbb{K}$ satisfying these three properties, and this map is called the determinant. In words, one says that the determinant is the unique alternating multilinear normalized map on the columns of a square matrix. Some derived properties are the following:

(D.1) Properties (1) and (2) above hold taking rows instead of colums.
(D.2) $\det \boldsymbol{A}^\mathsf{T} = \det \boldsymbol{A}$ for every $n \times n$ matrix $\boldsymbol{A}$.
(D.3) If two columns (or two rows) are exchanged the determinant changes sign.
(D.4) If one adds to a column a linear combination of the other columns (or to a row a linear combination of the other rows) the determinant does not change.
(D.5) $\det(\lambda \boldsymbol{A}) = \lambda^n \det \boldsymbol{A}$ for every $n \times n$ matrix $\boldsymbol{A}$ and every scalar λ.
(D.6) The determinant of a diagonal matrix or of an upper triangular matrix or of a lower triangular matrix is equal to the product of its diagonal elements.
(D.7) $\det \boldsymbol{AB} = \det \boldsymbol{A} \det \boldsymbol{B}$ for all $n \times n$ matrices $\boldsymbol{A}, \boldsymbol{B}$.
(D.8) $\det \boldsymbol{A} \neq 0$ iff $\boldsymbol{A}$ is invertible. By the previous property, we also see that in this case, $\det \boldsymbol{A}^{-1} = (\det \boldsymbol{A})^{-1}$.
(D.9) A collection $\boldsymbol{v}_1, \dots, \boldsymbol{v}_n$ of vectors of $\mathbb{K}^n$ is a basis iff the determinant of the matrix $\boldsymbol{S}$ whose ith column is $\boldsymbol{v}_i$ is different from zero. (Note that $\boldsymbol{v}_i = \boldsymbol{S}\boldsymbol{e}_i$ for all i.)
(D.10) The determinant of a block-diagonal matrix is the product of the determinants of its blocks: $\det \left(\begin{smallmatrix} \boldsymbol{A} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{D} \end{smallmatrix}\right) = \det \boldsymbol{A} \det \boldsymbol{D}$, where $\boldsymbol{A}$ and $\boldsymbol{D}$ are square matrices.

(D.11) More generally, $\det\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{B}\\ \boldsymbol{0} & \boldsymbol{D}\end{smallmatrix}\right) = \det \boldsymbol{A} \det \boldsymbol{D}$, where $\boldsymbol{A}$ and $\boldsymbol{D}$ are square matrices.

The last property is derived from the others by the following remarks. We should distinguish the case when $\boldsymbol{A}$ is invertible and when it is not. In the second case, there is some nonzero vector $\boldsymbol{v}$ in the kernel of $\boldsymbol{A}$. By completing $\boldsymbol{v}$ with zeros, we get a nonzero vector in the kernel of $\det\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{B}\\ \boldsymbol{0} & \boldsymbol{D}\end{smallmatrix}\right)$, which is then also not invertible. In this case, both sides of the equality vanish by (D.8). If, on the other hand, $\boldsymbol{A}$ is invertible, we can write $\left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{B}\\ \boldsymbol{0} & \boldsymbol{D}\end{smallmatrix}\right) = \left(\begin{smallmatrix}\boldsymbol{A} & \boldsymbol{0}\\ \boldsymbol{0} & \boldsymbol{D}\end{smallmatrix}\right)\left(\begin{smallmatrix}\boldsymbol{1} & \boldsymbol{A}^{-1}\boldsymbol{B}\\ \boldsymbol{0} & \boldsymbol{1}\end{smallmatrix}\right)$. The determinant of the first matrix on the right hand side is $\det \boldsymbol{A} \det \boldsymbol{D}$ by (D.10), whereas the determinant of the second is 1 by (D.6).

The determinant of an $n \times n$ matrix $\boldsymbol{A} = A^i_j$ can also be explicitly computed by the Leibniz formula

$$\det \boldsymbol{A} = \sum_{\sigma \in S_n} \operatorname{sgn}\sigma A^1_{\sigma(1)} \dots A^n_{\sigma(n)}, \tag{1.7}$$

where $\operatorname{sgn}\sigma$ is the sign of the permutation σ.[14] In particular,

$$\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad - bc.$$

The determinant can also be computed in terms of the Laplace expansion along the ith row

$$\det \boldsymbol{A} = \sum_{j=1}^{n} (-1)^{i+j} A^i_j d^i_j,$$

where d^i_j is the determinant of the matrix obtained by removing the ith row and the jth column from $\boldsymbol{A}$. For example, the Laplace expansion of a 3×3 matrix along the first row is

$$\det \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} = a \det \begin{pmatrix} e & f \\ h & i \end{pmatrix} - b \det \begin{pmatrix} d & f \\ g & i \end{pmatrix} + c \det \begin{pmatrix} d & e \\ g & h \end{pmatrix}.$$

An analogus formula for the expansion along a column also exists as a consequence of (D.2).

The determinant may also be used to compute the inverse of an invertible matrix as

$$\boldsymbol{A}^{-1} = \frac{1}{\det \boldsymbol{A}} \operatorname{adj} \boldsymbol{A},$$

[14]Every permutation may be written, in a nonunique way, as a product of transpositions, i.e., permutations that exchange exactly two elements. The parity of the number of occurring transpositions does not depend on the decomposition. The sign of a permutation is then defined as -1 to the number of occurring transpositions.

where adj $\boldsymbol{A}$ denotes the adjugate matrix of $\boldsymbol{A}$, i.e., the matrix whose (i, j) entry is $(-1)^{i+j}$ times the determinant d_i^j of the matrix obtained by removing the the jth row and and ith column from $\boldsymbol{A}$. This formula is theoretically important—it shows, e.g., that the entries of $\boldsymbol{A}^{-1}$ are rational functions of the entries of $\boldsymbol{A}$—but practically not so useful, apart from the 2×2 case, which is also easy to remember:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1} = \frac{1}{ad - bc} \begin{pmatrix} d & -b \\ -c & a \end{pmatrix}. \tag{1.8}$$

Namely, apart from dividing by the determinant, we just have to swap the diagonal entries and change the sign of the off diagonal entries.

Determinants may also be used, via property (D.8), to establish whether a linear map $F\colon V \to W$ between finite-dimensional vector spaces is invertible. By Proposition 1.67, we know that a necessary condition is that V and W have the same dimension, say, n. A representing matrix of F is then an $n \times n$ matrix, which by (D.8) is invertible iff its determinant is different from zero.[15] On the other hand, F is invertible iff any of its representing matrices is so. Therefore, we have the

Proposition 1.73. *A linear map $F\colon V \to W$ between finite-dimensional vector spaces is invertible iff* $\dim V = \dim W$ *and the determinant of any of its representing matrices is different from zero.*

It follows from (D.7) and (D.8) that $\det \boldsymbol{S}^{-1}\boldsymbol{B}\boldsymbol{S} = \det \boldsymbol{B}$, so similar matrices—see Remark 1.70—have the same trace. Therefore, we can make the following

Definition 1.74. The **determinant of an endomorphism** F of a finite-dimensional vector space V is the determinant of any of its representing matrices: $\det F := \det F_{\mathcal{B}}$, where $\mathcal{B}$ is any basis of V.

In particular, an endomorphism is invertible iff its determinant is different from zero.

Remark 1.75 (Discriminants). Congruent matrices may on the other hand have different determinants. We see however, from (D,2) and (D.7) that $\boldsymbol{B}' = \boldsymbol{S}^{\mathsf{T}}\boldsymbol{B}\boldsymbol{S}$ implies $\det \boldsymbol{B}' = (\det \boldsymbol{S})^2 \det \boldsymbol{B}$, so the determinant changes by a factor that is the square of a nonzero scalar. The **discriminant** of a matrix is by definition its determinant up to such a

[15]Note that this statment is independent of the chosen bases. In fact, any two representing matrices of F are equivalent $n \times n$ matrices $\boldsymbol{A}$ and $\boldsymbol{B}$, as in Remark 1.69, i.e., $\boldsymbol{A} = \boldsymbol{T}\boldsymbol{B}\boldsymbol{S}$, where $\boldsymbol{S}$ and $\boldsymbol{T}$ are invertible. We then have, by (D.7), $\det \boldsymbol{A} = c \det \boldsymbol{B}$, where $c = \det \boldsymbol{S} \det \boldsymbol{T}$ is a nonzero number.

factor. It follows that congruent matrices have the same discriminant and that we may define the discriminant of a bilinear form as the discriminant of any of its representing matrices. If we work over $\mathbb{C}$, where every element is a square, the only meaningful statement we can make is whether the discriminant is equal to zero or different from zero. Over $\mathbb{R}$, we can refine this and speak of strictly positive, strictly negative or zero discriminants.

CHAPTER 2

Linear ODEs and Diagonalization of Endomorphisms

In this chapter we discuss the problem of bringing an endomorphism, via a suitable choice of basis, to a diagonal representing matrix, whenever possible. We start by motivating this problem with the study of systems of linear ordinary differential equations with constant coefficients.

2.1. Differential equations

An ordinary differential equation (ODE) is an equation whose unknown is a differentiable function of one variable that appears in the equation together with its derivatives.

Newton's equation $F = ma$ is an example of an ODE. In this case, the unknown is a path $x(t)$.[1] In normal form (i.e., with the highest derivative set in evidence on the left hand side) the equation reads

$$\ddot{x} = \frac{1}{m} F(x, \dot{x}, t).$$

A solution is a specific function $x(t)$ that satisfies the equation for all t in some open interval (possibly the whole of $\mathbb{R}$). In the one-dimensional case, this is a single equation, but if we consider the problem in three space dimensions, we get a system of ODEs—one equation for each component. We also get a system is we describe the interaction of several particles.

The order of an ODE (or of a system of ODEs) corresponds to highest derivative occurring in it. For example, Newton's equation (like many fundamental equations in physics) is a second-order ODE. It is possible to reduce the order by the following trick, which we illustrate in the case of Newton's equation. Namely, we introduce the momentum $p := mv$. We can then rewrite Newton's equation as the first-order

[1] By path we mean a (twice) differentiable map whose domain is an interval.

system

$$\begin{cases} \dot{x} = \dfrac{p}{m} \\ \dot{p} = F\left(x, \dfrac{p}{m}, t\right) \end{cases}$$

where now the pair (x, p) is regarded as the unknown.

The Cauchy problem for a system of first-order ODEs consists of the system together the specification of the variables at some initial time.

The theory of ODEs is discussed in the analysis classes, and there it is proved, under some mild conditions, that a Cauchy problem has a unique solution (in an open interval around the initial time).

We will consider here only linear (systems of) first-order ODEs with constant coefficients, where methods of linear algebra can be used.

2.1.1. Linear ODEs with constant coefficients. ODEs are called linear if the unknown and its derivatives appear linearly. The ODE is called homogenous if there is no term independent of them, inhonogeneous otherwise.

A linear first-order ODE is then the equation $\dot{x} = ax$, when homogeneous, and $\dot{x} = ax + b$, when inhomogeneous, where a and b are given functions of t. We say that the equation has constant coefficients if a is constant.

NOTATION 2.1. It is a standard practice in the theory of ODEs to write (t) after a variable to specify that it is not assumed to be constant. If (t) does not appear, the variable is assumed to be a constant. The unknown function is written without (t) in the equation, as the notation $x(t)$ is reserved for writing a solution. Therefore, a linear first-order ODE constant coefficients is written as

$$\dot{x} = ax + b(t). \tag{2.1}$$

In the homogeneous case—i.e., when $b(t) = 0$—we write

$$\dot{x} = ax. \tag{2.2}$$

EXAMPLE 2.2 (Growth processes). The homogenous equation $\dot{x} = ax$ with constant a describes a growth process where the growth $\dot{x}$ is proportional to the quantity x itself (properly speaking, we have a growth when $a > 0$ and a decay when $a < 0$). Such equation is widely used: e.g., in economics to describe capital growth by compound interest, in biology to describe growth (or decline) of a population, in physics to describe radioactive decay.

To solve (2.2), we introduce $y(t) := \mathrm{e}^{-at}x(t)$. Differentianting we get $\dot{y} = \mathrm{e}^{-at}(\dot{x} - ax)$. Therefore, x is a solution to (2.2) iff $\dot{y} = 0$, i.e., $y = c$, where c is a constant. We then get the general solution $x(t) = \mathrm{e}^{at}c$. We can also rephrase this as the

PROPOSITION 2.3. *The Cauchy problem*

$$\begin{cases} \dot{x} &= ax \\ x(0) &= x_0 \end{cases}$$

for a homogenous linear ODE with constant coefficient a has the unique solution

$$x(t) = \mathrm{e}^{at} x_0. \tag{2.3}$$

The solution is defined for all $t \in \mathbb{R}$.

By the same trick, we may also study the associated nonhomogenous equation (2.1), where $b(t)$ is not necessarily assumed to be constant. Namely, we write again $y(t) := \mathrm{e}^{-at}x(t)$. In this case, x is a solution to (2.1) iff $\dot{y} = \mathrm{e}^{-at}\, b(t)$, so we can get y, and hence x, by integrating $\mathrm{e}^{-at}b(t)$. Namely, we have $y(t) = c + \int_0^t \mathrm{e}^{-as}\, b(s)\, \mathrm{d}s$, where c is a constant. Therefore, we have the general solution

$$x(t) = \mathrm{e}^{at}c + \int_0^t \mathrm{e}^{a(t-s)}\, b(s)\, \mathrm{d}s. \tag{2.4}$$

This leads to the

PROPOSITION 2.4. *The Cauchy problem*

$$\begin{cases} \dot{x} &= ax + b(t) \\ x(0) &= x_0 \end{cases}$$

for a nonhomogenous linear ODE with constant coefficient a has the unique solution

$$x(t) = \mathrm{e}^{at}x_0 + \int_0^t \mathrm{e}^{a(t-s)}\, b(s)\, \mathrm{d}s. \tag{2.5}$$

The solution is defined for all $t \in \mathbb{R}$.

2.1.2. Systems of linear ODEs with constant coefficients. A system of the form

$$\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{b}(t), \tag{2.6}$$

where $\boldsymbol{A}$ is a given $n \times n$ matrix (with constant entries), called the coefficient matrix, and $\boldsymbol{b}$ is a given map from an open interval to $\mathbb{R}^n$, is called a system of n linear ODEs with constant coefficients. For $\boldsymbol{b}(t)$ the zero map, we have the system

$$\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x}, \tag{2.7}$$

which is called homogeneous.

EXAMPLE 2.5 (Harmonic oscillator). Consider Newton's equation in one dimension with force $F = -kx$, where k is a given positive constant. The second-order ODE $m\ddot{x} = -kx$ may be rewritten as the system

$$\begin{cases} \dot{x} = \dfrac{p}{m} \\ \dot{p} = -kx \end{cases}$$

which can be brought in matrix form $\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x}$ by setting $\boldsymbol{x} = \left(\begin{smallmatrix} x \\ p \end{smallmatrix}\right)$ and

$$\boldsymbol{A} = \begin{pmatrix} 0 & \frac{1}{m} \\ -k & 0 \end{pmatrix}.$$

EXAMPLE 2.6 (Homogenous nth-order linear ODE with constant coefficients). Consider the ODE

$$x^{(n)} + a_1 x^{(n-1)} + \cdots + a_{n-1}\dot{x} + a_n x = 0, \tag{2.8}$$

where $a_1, \ldots, a_n$ are given constants. This ODE can be rewritten as a system by defining

$$\boldsymbol{x} := \begin{pmatrix} x \\ \dot{x} \\ \vdots \\ x^{(n-2)} \\ x^{(n-1)} \end{pmatrix}.$$

In fact, we have

$$\dot{\boldsymbol{x}} = \begin{pmatrix} \dot{x} \\ \ddot{x} \\ \vdots \\ x^{(n-1)} \\ x^{(n)} \end{pmatrix} = \begin{pmatrix} \dot{x} \\ \ddot{x} \\ \vdots \\ x^{(n-1)} \\ -a_1 x^{(n-1)} - \cdots - a_{n-1}\dot{x} - a_n x \end{pmatrix}.$$

Therefore, the ODE (2.8) is equivalent to the system (2.7) with

$$\boldsymbol{A} = \begin{pmatrix} 0 & 1 & & & \\ 0 & 0 & 1 & & \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ 0 & \cdots & \cdots & 0 & 1 \\ -a_n & -a_{n-1} & \cdots & -a_2 & -a_1 \end{pmatrix}, \tag{2.9}$$

where all the nondisplayed entries are equal to zero.

EXAMPLE 2.7 (Infinite-dimensional examples). In physics one also studies equations involving functions of several variables together with their partial derivatives—these are called partial differential equations

(PDEs). Several important PDEs in physics are linear. For example, the wave equation

$$\frac{1}{c^2}\frac{\partial^2\psi}{\partial t^2} = \Delta\psi,$$

the heat equation

$$\frac{\partial\psi}{\partial t} = \alpha\Delta\psi,$$

and the Schrödinger equation

$$\frac{\mathrm{i}}{\hbar}\frac{\partial\psi}{\partial t} = -\frac{\hbar^2}{2m}\Delta\psi + V\psi.$$

In these examples, ψ is a function (real in the first two cases and complex in the third) of the time variable t and the space variables x, y, z; $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ is the Laplace operator; c, α, $\hbar$, and m are real constants (respectively: velocity, diffusivity, Planck constant, mass); and V is a real function of the space variables (the potential). Each of these equations may be viewed as a system of linear ODEs with constant coefficients in the infinite-dimensional vector space $C^2(\mathbb{R}^3)$ of twice differentiable functions in the space variables. The unknown $\psi(t; x, y, z)$ is then viewed as map $\mathbb{R} \to C^2(\mathbb{R}^3)$, $t \mapsto \psi_t$, with $\psi_t(x, y, z) := \psi(t; x, y, z)$. The techniques we present in this section for finite systems of linear ODEs with constant coefficients may be extended to these infinite-dimensional systems, but we will not do it here.

Example 2.8 (Diagonal case). The homogeneous system (2.6) may be easily solved if the coefficient matrix is diagonal. Namely, suppose $\boldsymbol{A} = \boldsymbol{D}$ with $\boldsymbol{D}$ diagonal with diagonal entries $\lambda_1, \dots, \lambda_n$:

$$\boldsymbol{D} = \begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix}. \tag{2.10}$$

The system then splits into n independent equations

$$\dot{x}^1 = \lambda_1 x^1, \dots, \dot{x}^n = \lambda_n x^n.$$

The ith equation has the general solution $x^i(t) = \mathrm{e}^{\lambda_i t} c^i$, where c^i is a constant. It then follows that the associated Cauchy problem with initial condition $\boldsymbol{x}(0) = \boldsymbol{x}_0$ has the unique solution $x^i(t) = \mathrm{e}^{\lambda_i t} x^i_0$, which is defined for all $t \in \mathbb{R}$, for $i = 1, \dots, n$. If we denote by $\mathrm{e}^{\boldsymbol{D}t}$ the diagonal matrix with diagonal entries $\mathrm{e}^{\lambda_1 t}, \dots, \mathrm{e}^{\lambda_n t}$,

$$\mathrm{e}^{\boldsymbol{D}t} = \begin{pmatrix} \mathrm{e}^{\lambda_1 t} & & \\ & \ddots & \\ & & \mathrm{e}^{\lambda_n t} \end{pmatrix}, \tag{2.11}$$

we can write the unique solution as

$$\boldsymbol{x}(t) = \mathrm{e}^{\boldsymbol{D}t}\boldsymbol{x}_0. \tag{2.12}$$

DIGRESSION 2.9 (Upper triangular case). The homogenous case when the coefficient matrix is upper triangular may also be easily solved. To illustrate the idea, we consider the two-dimensional case with $\boldsymbol{A} = \left(\begin{smallmatrix} \lambda_1 & \beta \\ 0 & \lambda_2 \end{smallmatrix}\right)$. We then have the two equations

$$\dot{x}^1 = \lambda_1 x^1 + \beta x^2, \quad \dot{x}^2 = \lambda_2 x^2.$$

The second equation is independent of x^1 and has the general solution $x^2(t) = \mathrm{e}^{\lambda_2 t} c^2$, where c^2 is a constant. We may plug this solution into the first equation getting the nonhomogenous equation

$$\dot{x}^1 = \lambda_1 x^1 + \beta\, \mathrm{e}^{\lambda_2 t} c^2,$$

which can be solved using (2.4) with $b(t) = \mathrm{e}^{\lambda_2 t}\beta c^2$. If $\lambda_1 \neq \lambda_2$, we then get

$$x^1(t) = \mathrm{e}^{\lambda_1 t} c^1 + \frac{\mathrm{e}^{\lambda_2 t}}{\lambda_2 - \lambda_1}\beta c^2,$$

where c^1 is a new constant. If $\lambda_1 = \lambda_2 = \lambda$, we get instead

$$x^1(t) = \mathrm{e}^{\lambda t} c^1 + t\, \mathrm{e}^{\lambda t}\beta c^2.$$

Note that in the "degenerate case" when $\lambda_1 = \lambda_2$, the solution does not only depend on exponential functions but also has a factor t. The general case, with $\boldsymbol{A}$ upper triangular, is solved similarly. One solves the equations iteratively from the last equation, which only depends on the last component of $\boldsymbol{x}$, to the first. Every time one inserts the solution into the previous equation, which then turns out to be a nonhomogeneous linear ODE that can be solved by (2.4). By induction one sees that the inhomogenous term $b(t)$ is a linear combination of products of exponential and polynomials. In conclusion, the general solution will also be given by a linear combination of products of exponentials and polynomials. If the diagonal entries are all different, the general solution is simply a linear combination of exponentials.

2.1.3. The matrix exponential. Following the examples of the solutions (2.3) and (2.12), we now want to get a general solution to (2.7) in the form of an exponential.

The exponential of a square matrix $\boldsymbol{A}$ is defined by extending to matrices the usual series defining the exponential of a real or complex number:

$$\mathrm{e}^{\boldsymbol{A}} := \sum_{k=0}^{\infty} \frac{1}{k!}\boldsymbol{A}^k = \boldsymbol{1} + \boldsymbol{A} + \frac{1}{2}\boldsymbol{A}^2 + \cdots .$$

One can easily see that the series converges for any matrix $\boldsymbol{A}$ and, moreover, that the power series

$$\mathrm{e}^{\boldsymbol{A}t} = \sum_{k=0}^{\infty} \frac{t^k}{k!}\boldsymbol{A}^k = \mathbf{1} + t\boldsymbol{A} + \frac{t^2}{2}\boldsymbol{A}^2 + \cdots .$$

has infinite radius of convergence. It follows that it defines a smooth (i.e., infinitely often continuously differentiable) function and that taking a derivative commutes with the sum, so we get

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{e}^{\boldsymbol{A}t} = \mathrm{e}^{\boldsymbol{A}t}\boldsymbol{A} = \boldsymbol{A}\mathrm{e}^{\boldsymbol{A}t}. \tag{2.13}$$

Moreover, from

$$\mathrm{e}^{\boldsymbol{A}0} = \mathbf{1},$$

we see that $\boldsymbol{U}(t) := \mathrm{e}^{\boldsymbol{A}t}$ is the unique solution to the matrix Cauchy problem

$$\begin{cases} \dot{\boldsymbol{U}} &= \boldsymbol{A}\boldsymbol{U} \\ \boldsymbol{U}(0) &= \mathbf{1} \end{cases}$$

Remark 2.10. The matrix exponential has the following properties, which can be easily proved:

(1) As in the case of the exponential of a number,

$$\mathrm{e}^{\boldsymbol{A}(t+s)} = \mathrm{e}^{\boldsymbol{A}t}\,\mathrm{e}^{\boldsymbol{A}s} \tag{2.14}$$

for all $t, s \in \mathbb{C}$. In particular, taking $s = -t$, we see that

$$\mathbf{1} = \mathrm{e}^{\boldsymbol{A}t}\mathrm{e}^{-\boldsymbol{A}t},$$

so $\mathrm{e}^{\boldsymbol{A}t}$ is always invertible and its inverse is $\mathrm{e}^{-\boldsymbol{A}t}$.

(2) If $\boldsymbol{A}$ and $\boldsymbol{B}$ commute (i.e., $\boldsymbol{AB} = \boldsymbol{BA}$), then in the product of two exponentials we can rearrange the factors. Therefore, we have

$$\mathrm{e}^{\boldsymbol{A}+\boldsymbol{B}} = \mathrm{e}^{\boldsymbol{A}}\,\mathrm{e}^{\boldsymbol{B}} \quad \text{if } \boldsymbol{AB} = \boldsymbol{BA}. \tag{2.15}$$

Note that this equation does not hold if $\boldsymbol{A}$ and $\boldsymbol{B}$ do not commute because on the right hand side all powers of $\boldsymbol{A}$ comes to the left and all powers of $\boldsymbol{B}$ to the right, whereas on the left hand side powers of $\boldsymbol{A}$ and $\boldsymbol{B}$ come in all possible orders.

(3) For every invertible matrix $\boldsymbol{S}$, we have

$$\mathrm{e}^{\boldsymbol{S}^{-1}\boldsymbol{AS}} = \boldsymbol{S}^{-1}\mathrm{e}^{\boldsymbol{A}}\boldsymbol{S}. \tag{2.16}$$

This follows from $(\boldsymbol{S}^{-1}\boldsymbol{AS})^k = \boldsymbol{S}^{-1}\boldsymbol{A}^k\boldsymbol{S}$, which is easily proved for all k.

One more interesting property of the matrix exponential is given by the following

PROPOSITION 2.11. *Let* $\boldsymbol{A}$ *be a square matrix. Then*

$$\det \mathrm{e}^{\boldsymbol{A}t} = \mathrm{e}^{t \operatorname{tr} \boldsymbol{A}}$$

for every t.

PROOF. We consider the function $d(t) := \det \mathrm{e}^{\boldsymbol{A}t}$ and compute its derivative. Using (2.14) and the multiplicativity of the determinant—i.e, property (D.7) on page 27—we have $d(t+h) = d(t)d(h)$. Therefore, using $d(0) = 1$, we have

$$\dot{d}(t) = \lim_{h\to 0} \frac{d(t+h) - d(t)}{h} = \lim_{h\to 0} \frac{d(h) - 1}{h} d(t) = ad(t)$$

with $a := \dot{d}(0)$. By Proposition 2.3, we then have $d(t) = \mathrm{e}^{at}$. To complete the proof, we only have to show that $a = \operatorname{tr} \boldsymbol{A}$. This can be done explicitly by using the Leibniz formula (1.7):

$$\det \mathrm{e}^{\boldsymbol{A}h} = \det(\boldsymbol{1} + h\boldsymbol{A} + O(h^2)) = 1 + h \operatorname{tr} \boldsymbol{A} + O(h^2).$$

□

We may use the matrix exponential to solve any homogeneous linear system of ODEs with constant coefficients $\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x}$ by the same trick we used in the case of a single equation. Namely, we introduce $\boldsymbol{y} := \mathrm{e}^{-\boldsymbol{A}t}\boldsymbol{x}$. Differentianting, thanks to (2.13), we get $\dot{\boldsymbol{y}} = \mathrm{e}^{-\boldsymbol{A}t}(\dot{\boldsymbol{x}} - \boldsymbol{A}\boldsymbol{x})$. Therefore, $\boldsymbol{x}$ is a solution to (2.7) iff $\dot{\boldsymbol{y}} = 0$, i.e., $\boldsymbol{y} = \boldsymbol{c}$, where $\boldsymbol{c}$ is a constant vector. We then get the general solution $\boldsymbol{x}(t) = \mathrm{e}^{\boldsymbol{A}t}\boldsymbol{c}$. We can also rephrase this as the

PROPOSITION 2.12. *The Cauchy problem*

$$\begin{cases} \dot{\boldsymbol{x}} & = & \boldsymbol{A}\boldsymbol{x} \\ \boldsymbol{x}(0) & = & \boldsymbol{x}_0 \end{cases}$$

for a homogenous linear system of ODEs with constant coefficients has the unique solution

$$\boldsymbol{x}(t) = \mathrm{e}^{\boldsymbol{A}t}\boldsymbol{x}_0. \tag{2.17}$$

The solution is defined for all $t \in \mathbb{R}$.

By the same trick, we may also study the associated nonhomogenous equation (2.6), where $\boldsymbol{b}(t)$ is a (not necessarily constant) map from an interval to $\mathbb{R}^n$. Namely, we write again $\boldsymbol{y} := \mathrm{e}^{-\boldsymbol{A}t}\boldsymbol{x}$. In this case, $\boldsymbol{x}$ is a solution to (2.6) iff $\dot{\boldsymbol{y}} = \mathrm{e}^{-\boldsymbol{A}t}\,\boldsymbol{b}(t)$, so we can get $\boldsymbol{y}$, and hence $\boldsymbol{x}$, by integrating $\mathrm{e}^{-\boldsymbol{A}t}\,\boldsymbol{b}(t)$. Namely, we have $\boldsymbol{y}(t) = \boldsymbol{c} + \int_0^t \mathrm{e}^{-\boldsymbol{A}s}\,\boldsymbol{b}(s)\,\mathrm{d}s$, where $\boldsymbol{c}$ is a constant vector, and the integral is computed componentwise. Therefore, we have the general solution

$$\boldsymbol{x}(t) = \mathrm{e}^{\boldsymbol{A}t}\boldsymbol{c} + \int_0^t \mathrm{e}^{\boldsymbol{A}(t-s)}\,\boldsymbol{b}(s)\,\mathrm{d}s. \tag{2.18}$$

2.1.4. Computing the matrix exponential. The practical problem consists in computing the exponential of a given matrix $\boldsymbol{A}$. The simplest case is when the matrix $\boldsymbol{A}$ is diagonal, $\boldsymbol{A} = \boldsymbol{D}$ with $\boldsymbol{D}$ as in (2.10). In fact, we have

$$\boldsymbol{D}^k = \begin{pmatrix} \lambda_1^k & & \\ & \ddots & \\ & & \lambda_n^k \end{pmatrix},$$

so $\mathrm{e}^{\boldsymbol{D}t}$ is as in (2.11). This way we recover the solution discussed in Example 2.8.

Thanks to (2.13), we can also explictly compute the exponential of a diagonalizable matrix, i.e., a square matrix $\boldsymbol{A}$ that is similar to a diagonal matrix $\boldsymbol{D}$. Writing

$$\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \boldsymbol{D}$$

for some invertible matrix $\boldsymbol{S}$, we get the solution to the associated Cauchy problem in the form

$$\boldsymbol{x}(t) = \boldsymbol{S}^{-1} \begin{pmatrix} \mathrm{e}^{\lambda_1 t} & & \\ & \ddots & \\ & & \mathrm{e}^{\lambda_n t} \end{pmatrix} \boldsymbol{S}\boldsymbol{x}_0,$$

where $\lambda_1, \dots, \lambda_n$ are the diagonal elements of $\boldsymbol{D}$. In the next sections we will develop methods to determine the scalars $\lambda_1, \dots, \lambda_n$ and the matrix $\boldsymbol{S}$, whenever possible.

REMARK 2.13. Not every matrix is diagonalizable. Consider, e.g., $\boldsymbol{A} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$. Pick an invertible matrix $\boldsymbol{S} = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. Computing $\boldsymbol{S}^{-1}$ as in (1.8), we then get

$$\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \frac{1}{ad - bc} \begin{pmatrix} dc & d^2 \\ -c^2 & -cd \end{pmatrix}.$$

Since $\boldsymbol{S}$ is invertible, the entries c and d cannot be both zero, so the right and side cannot be a diagonal matrix.

DIGRESSION 2.14. By Digression 2.9, and by (2.13), we can also explicitly compute the exponential of a matrix that is similar to an upper triangular matrix. This turns out to be always possible if we work over complex numbers. (See Section 2.4.)

EXAMPLE 2.15. Consider again the nondiagonalizable matrix $\boldsymbol{A} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$. We can compute its exponential explicitly as follows. We easily see that $\boldsymbol{A}^2 = \boldsymbol{0}$, which in turn implies $\boldsymbol{A}^n = 0$ for all $n > 1$. Therefore,

$$\mathrm{e}^{\boldsymbol{A}t} = \boldsymbol{1} + t\boldsymbol{A} = \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}.$$

This exponential actually solves the problem of the free particle

$$m\ddot{x} = 0,$$

which is the same as the harmonic oscillator of Example 2.5 but with $k = 0$. The coefficient matrix is in this case $\boldsymbol{A} = \begin{pmatrix} 0 & \frac{1}{m} \\ 0 & 0 \end{pmatrix}$, i.e., $\frac{1}{m}$ times the matrix considered above. In this case, we have

$$\mathrm{e}^{\boldsymbol{A}t} = \mathbf{1} + t\boldsymbol{A} = \begin{pmatrix} 1 & \frac{t}{m} \\ 0 & 1 \end{pmatrix}.$$

The solution to the Cauchy problem is then

$$\begin{pmatrix} x(t) \\ p(t) \end{pmatrix} = \mathrm{e}^{\boldsymbol{A}t} \begin{pmatrix} x_0 \\ p_0 \end{pmatrix} = \begin{pmatrix} 1 & \frac{t}{m} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_0 \\ p_0 \end{pmatrix} = \begin{pmatrix} x_0 + \frac{t}{m} p_0 \\ p_0 \end{pmatrix}.$$

This yields the usual formula

$$x(t) = x_0 + \frac{p_0}{m} t.$$

2.2. Diagonalization of matrices

Suppose $\boldsymbol{A}$ is diagonalizable, i.e., $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \boldsymbol{D}$ for some invertible matrix $\boldsymbol{S}$ and a diagonal matrix $\boldsymbol{D}$. If λ_i is the ith diagonal entry of $\boldsymbol{D}$, then we have $\boldsymbol{D}\boldsymbol{e}_i = \lambda_i \boldsymbol{e}_i$. Denoting by $\boldsymbol{v}_i$ the ith column of $\boldsymbol{S}$, i.e., $\boldsymbol{v}_i = \boldsymbol{S}\boldsymbol{e}_i$, we get

$$\boldsymbol{A}\boldsymbol{v}_i = \boldsymbol{A}\boldsymbol{S}\boldsymbol{e}_i = \boldsymbol{S}\boldsymbol{D}\boldsymbol{e}_i = \lambda_i \boldsymbol{S}\boldsymbol{e}_i = \lambda_i \boldsymbol{v}_i.$$

This motivates the following

Definition 2.16 (Eigenvectors and eigenvalues). A nonzero vector $\boldsymbol{v}$ is called an eigenvector of a square matrix $\boldsymbol{A}$ if there is a scalar λ, called the eigenvalue to the eigenvector $\boldsymbol{v}$, such that

$$\boldsymbol{A}\boldsymbol{v} = \lambda \boldsymbol{v}. \tag{2.19}$$

We then have the

Theorem 2.17. *A square matrix is diagonalizable iff it admits a basis of eigenvectors.*

Proof. If $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \boldsymbol{D}$ is diagonal, then, by the above discussion, the vectors $\boldsymbol{v}_i = \boldsymbol{S}\boldsymbol{e}_i$ are eigenvectors. They are a basis because $\boldsymbol{S}$ is invertible. (Note that $\boldsymbol{S}$ defines the change of basis from the standard basis $(\boldsymbol{e}_1, \dots, \boldsymbol{e}_n)$ to the basis $(\boldsymbol{v}_1, \dots, \boldsymbol{v}_n)$.)

On the other hand, if $(\boldsymbol{v}_1, \dots, \boldsymbol{v}_n)$ is a basis of eigenvectors, we define $\boldsymbol{S}$ as the matrix whose ith column is $\boldsymbol{v}_i$. It is invertible by property (D.9) on page 27. We then have

$$\boldsymbol{D}\boldsymbol{e}_i = \boldsymbol{D}\boldsymbol{S}^{-1}\boldsymbol{v}_i = \boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{v}_i = \lambda_i \boldsymbol{S}^{-1}\boldsymbol{v}_i = \lambda_i \boldsymbol{e}_i,$$

which shows that $\boldsymbol{D}$ is diagonal with diagonal entries $\lambda_1, \dots, \lambda_n$. $\square$

REMARK 2.18 (Linear systems of ODEs). Back to our problem $\dot{\boldsymbol{x}} = \boldsymbol{A}\boldsymbol{x}$, assuming we have a basis of eigenvectors $(\boldsymbol{v}_1, \dots, \boldsymbol{v}_n)$ of $\boldsymbol{A}$, we have the expansion $\boldsymbol{x}(t) = \xi^i(t)\boldsymbol{v}_i$, with uniquely determined scalars $\xi^1(t), \dots, \xi^n(t)$ for each t. The system then becomes $\dot{\xi}^1 = \lambda_1\xi^1, \dots, \dot{\xi}^n = \lambda_n\xi^n$, which is solved by $\xi^i(t) = \mathrm{e}^{\lambda_i t}\xi_0^i$, with $(\xi_0^1, \dots, \xi_0^n)$ the components of the expansion of $\boldsymbol{x}_0$: $\boldsymbol{x}_0 = \xi_0^i\boldsymbol{v}_i$. Therefore, we get the unique solution to the Cauchy problem in the form

$$\boldsymbol{x}(t) = \sum_{i=1}^{n} \mathrm{e}^{\lambda_i t}\xi_0^i\boldsymbol{v}_i. \tag{2.20}$$

REMARK 2.19 (Choice of field). For our application to linear systems of ODE we assume $\boldsymbol{A}$ to have real (or complex) entries. On the other hand, the general problem of diagonalization, Theorem 2.17, and the rest of the discussion make sense for every ground field.

Now note that we can rewrite the eigenvector equation (2.19) as $(\boldsymbol{A} - \lambda\mathbf{1})\boldsymbol{v} = 0$. This shows the following

LEMMA 2.20. *A scalar λ is an eigenvalue of the square matrix $\boldsymbol{A}$ iff $\boldsymbol{A} - \lambda\mathbf{1}$ is not invertible.*

PROOF. $\boldsymbol{A}-\lambda\mathbf{1}$ is invertible iff its kernel is different from zero. This happens iff there is a nonzero vector $\boldsymbol{v}$ such that $(\boldsymbol{A} - \lambda\mathbf{1})\boldsymbol{v} = 0$. $\square$

It follows, from property (D.8) on page 27, that the eigenvalues of $\boldsymbol{A}$ are precisely the solutions to $\det(\boldsymbol{A} - \lambda\mathbf{1}) = 0$. This motivates considering the function

$$P_{\boldsymbol{A}} := \det(\boldsymbol{A} - \lambda\mathbf{1}).$$

If $\boldsymbol{A}$ is an $n \times n$ matrix, by the Leibniz formula (1.7) we see that $P_{\boldsymbol{A}}$ is a polynomial of degree n,

$$P_A = b_0\lambda^n + b_1\lambda^{n-1} + \cdots + b_n.$$

In particular, $b_0 = (-1)^n$, $b_1 = (-1)^{n-1}\operatorname{tr}\boldsymbol{A}$, and $b_n = \det\boldsymbol{A}$.

DEFINITION 2.21 (Characteristic polynomial). The polynomial $P_{\boldsymbol{A}}$ is called the characteristic polynomial of the square matrix $\boldsymbol{A}$.

We may summarize the previous discussion as the

PROPOSITION 2.22. *The eigenvalues of a square matrix are the roots of its characteristic polynomial.*

In the quest for the diagonalization of $\boldsymbol{A}$, we first compute its eigenvalues by this proposition. Next, we proceed to the determination of its eigenvectors. That is, for each root λ of the characteristic polynomial of $\boldsymbol{A}$, we consider the system $\boldsymbol{A}\boldsymbol{v} = \lambda\boldsymbol{v}$ of n linear equations (with unknown the components $v^1, \dots, v^n$ of the vector $\boldsymbol{v}$). We know that this system has nontrivial solutions because $\boldsymbol{A} - \lambda\mathbf{1}$ is not invertible.

Note that for every eigenvalue there are infinitely many eigenvectors. For example, if $\boldsymbol{v}$ is a λ-eigenvalue, then so is $a\boldsymbol{v}$ for every scalar $a \neq 0$. More generally, if $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$ are λ-eigenvalues, then so is any nonzero linear combination of them.

In order to diagonalize a matrix, we have to determine all its eigenvalues, but there is no need to find all the corresponding eigenvectors: it is enough to find a basis of eigenvectors (if possible).

Example 2.23. Let $\boldsymbol{A} = \begin{pmatrix} 0 & b \\ c & 0 \end{pmatrix}$. Its characteristic polynomial is $P_{\boldsymbol{A}} = \det \begin{pmatrix} -\lambda & b \\ c & -\lambda \end{pmatrix} = \lambda^2 - bc$. Assuming $bc > 0$, we have the two real distinct roots $\lambda_\pm = \pm\sqrt{bc}$. The eigenvector equation $\boldsymbol{A}\boldsymbol{v} = \lambda_+\boldsymbol{v}$ is then the system

$$\begin{cases} bv^2 = \sqrt{bc}\, v^1 \\ cv^1 = \sqrt{bc}\, v^2 \end{cases}$$

The first equation yields the relation $v^2 = \sqrt{\frac{c}{b}}\, v^1$. The second equation does not yield any new independent condition (this is a consequence of $\boldsymbol{A}-\lambda_+\mathbf{1}$ being not invertible). Therefore, we have a 1-parameter family of solutions (we can choose $v^1 \in \mathbb{R}$ as the parameter). For example, for $v^1 = 1$ we have the eigenvector $\boldsymbol{v}_+ = \begin{pmatrix} 1 \\ \sqrt{\frac{c}{b}} \end{pmatrix}$. A similar computation yields the eigenvector $\boldsymbol{v}_- = \begin{pmatrix} 1 \\ -\sqrt{\frac{c}{b}} \end{pmatrix}$ to the eigenvalue λ_-.[2] One can easily check that $(\boldsymbol{v}_+, \boldsymbol{v}_-)$ is a basis of $\mathbb{R}^2$. The transformation matrix and its inverse are then

$$\boldsymbol{S} = (\boldsymbol{v}_+ \ \boldsymbol{v}_-) = \begin{pmatrix} 1 & 1 \\ \sqrt{\frac{c}{b}} & -\sqrt{\frac{c}{b}} \end{pmatrix}, \quad \boldsymbol{S}^{-1} = \frac{1}{2} \begin{pmatrix} 1 & \sqrt{\frac{b}{c}} \\ 1 & -\sqrt{\frac{b}{c}} \end{pmatrix}.$$

One can then explicitly verify that $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \begin{pmatrix} \sqrt{bc} & 0 \\ 0 & -\sqrt{bc} \end{pmatrix}$.

---

[2] If b and c are both positive, we could also pick $v^1 = \sqrt{b}$, getting $v^2 = \sqrt{c}$ and $\boldsymbol{v}_\pm = \begin{pmatrix} \sqrt{b} \\ \pm\sqrt{c} \end{pmatrix}$.

EXAMPLE 2.24. Consider the matrix $\boldsymbol{A}$ of (2.9) associated to a homogenous nth-order linear ODE with constant coefficients as in Example 2.6. It is a good exercise to show that in this case

$$P_{\boldsymbol{A}} = (-1)^n(\lambda^n + a_1\lambda^{n-1} + \cdots + a_{n-1}\lambda + a_n).$$

Therefore, the characteristic equation $P_{\boldsymbol{A}} = 0$ may be obtained from (2.8) by formally substituting λ^k to $x^{(k)}$ for $k = 0, \dots, n$. Equivalently, it may be obtained by inserting into (2.8) the ansatz $x(t) = \mathrm{e}^{\lambda t}$.

To be sure that the roots of the characteristic polynomial exist, we assume from now that we work over $\mathbb{C}$ and make use of the

THEOREM 2.25 (Fundamental theorem of algebra). *A nonconstant complex polynomial has a root. As a consequence, it splits into a product of linear factors.*

The characteristic polynomial $P_{\boldsymbol{A}}$ of a complex $n \times n$ matrix $\boldsymbol{A}$ factorizes as

$$P_{\boldsymbol{A}} = (-1)^n\,(\lambda - \lambda_1)^{s_1} \cdots (\lambda - \lambda_k)^{s_k},$$

where $\lambda_1, \dots, \lambda_k$ are the pairwise distinct roots of $P_{\boldsymbol{A}}$. Note that $k \leq n$. The exponent s_i is called the algebraic multiplicity of λ_i. Note that we have $s_1 + \cdots + s_k = n$.

REMARK 2.26 (Linear systems of ODEs). The coefficient matrix $\boldsymbol{A}$ of a linear system of ODEs is usually assumed to be real, and we are interested in a real solution $\boldsymbol{x}(t)$. The trick is to regard $\boldsymbol{A}$ as a complex matrix; assuming it then to be diagonalizable, we may find a basis of eigenvectors and proceed as in Remark 2.18. The unique solution to the Cauchy problem is still given by (2.20), i.e.,

$$\boldsymbol{x}(t) = \sum_{i=1}^{n} \mathrm{e}^{\lambda_i t}\xi_0^i \boldsymbol{v}_i,$$

where now the λ_is, ξ_0^is, and $\boldsymbol{v}_i$s may be complex. If the initial condition $\boldsymbol{x}_0$ is real, then the unique solution is also real, which ensures that the sum of complex vectors on the right hand side of (2.20) yields a real vector. It is always possible to rearrange this sum of complex exponentials times complex vectors into a real sum involving real exponentials and trigonometric functions as we explain in the next example and, more generally, in Section 2.2.1.

EXAMPLE 2.27 (Harmonic oscillator). In the Example 2.5 of the harmonic oscillator, we have $\boldsymbol{A} = \begin{pmatrix} 0 & \frac{1}{m} \\ -k & 0 \end{pmatrix}$. The characteristic polynomial is

$$P_{\boldsymbol{A}} = \det(\boldsymbol{A} - \lambda\mathbf{1}) = \begin{pmatrix} -\lambda & \frac{1}{m} \\ -k & -\lambda \end{pmatrix} = \lambda^2 + \omega^2$$

with $\omega := \sqrt{\frac{k}{m}}$. The complex eigenvalues are then $\pm\mathrm{i}\omega$. To find the eigenvectors, we then study first the equation $\boldsymbol{A}\boldsymbol{v} = \mathrm{i}\omega\boldsymbol{v}$. Writing $\boldsymbol{v} = \left(\begin{smallmatrix} a \\ b \end{smallmatrix}\right)$, we get the equation

$$\frac{b}{m} = \mathrm{i}\omega a$$

(with the second equation in the system a multiple of this one). By choosing $a = 1$, we get the eigenvector $\boldsymbol{v} = \left(\begin{smallmatrix} 1 \\ \mathrm{i}m\omega \end{smallmatrix}\right)$. Similarly, one sees that $\bar{\boldsymbol{v}} = \left(\begin{smallmatrix} 1 \\ -\mathrm{i}m\omega \end{smallmatrix}\right)$ is an eigenvector for $-\mathrm{i}\omega$. (This is a general fact: if $\boldsymbol{v}$ is an eigenvector with eigenvalue λ for a real matrix $\boldsymbol{A}$, then $\bar{\boldsymbol{v}}$ is an eigenvector for $\bar{\lambda}$; this follows from taking the complex conjugation $\boldsymbol{A}\bar{\boldsymbol{v}} = \bar{\lambda}\bar{\boldsymbol{v}}$ of the equation $\boldsymbol{A}\boldsymbol{v} = \lambda\boldsymbol{v}$. See also Remark 2.36.) Since $(\boldsymbol{v}, \bar{\boldsymbol{v}})$ is a basis of $\mathbb{C}^2$, the matrix $\boldsymbol{A}$ is diagonalizable. The general real solution $\boldsymbol{x}$ of the associated linear system of ODEs has then the form

$$\boldsymbol{x}(t) = z\mathrm{e}^{\mathrm{i}\omega t}\begin{pmatrix} 1 \\ \mathrm{i}m\omega \end{pmatrix} + \bar{z}\mathrm{e}^{-\mathrm{i}\omega t}\begin{pmatrix} 1 \\ -\mathrm{i}m\omega \end{pmatrix}$$

for some complex constant z. In particular, the first component is

$$x(t) = z\mathrm{e}^{\mathrm{i}\omega t} + \bar{z}\mathrm{e}^{-\mathrm{i}\omega t} = A\cos(\omega t + \alpha)$$

if we write $z = \frac{A}{2}\mathrm{e}^{\mathrm{i}\alpha}$.

REMARK 2.28. Note that a nonzero vector cannot be the eigenvector of two different eigenvalues. In fact, if $\boldsymbol{A}\boldsymbol{v} = \lambda\boldsymbol{v}$ and $\boldsymbol{A}\boldsymbol{v} = \mu\boldsymbol{v}$, we get by taking the difference that $(\lambda - \mu)\boldsymbol{v} = 0$. If $\lambda \neq \mu$, we then get $\boldsymbol{v} = 0$.

Another important observation is the following

LEMMA 2.29. *A collection $(\boldsymbol{v}_1, \dots, \boldsymbol{v}_m)$ of eigenvectors of an $n \times n$ matrix $\boldsymbol{A}$ corresponding to pairwise distinct eigenvalues is linearly independent.*

PROOF. Suppose $\sum_{i=1}^m \alpha_i \boldsymbol{v}_i = 0$. Pick $k \in \{1, \dots, m\}$. If we apply $\prod_{j\neq k}(\boldsymbol{A} - \lambda_j \mathbf{1})$ to the sum, all the terms with $i \neq k$ are killed. On the other hand,

$$\prod_{j\neq k}(\boldsymbol{A} - \lambda_j \mathbf{1})\boldsymbol{v}_k = \prod_{j\neq k}(\lambda_k - \lambda_j)\boldsymbol{v}_k.$$

Therefore, $\alpha_k \prod_{j\neq k}(\lambda_k - \lambda_j)\boldsymbol{v}_k = 0$. Since $\boldsymbol{v}_k \neq 0$ and $\lambda_j \neq \lambda_k$ for all $j \neq k$, we get $\alpha_k = 0$. Repeating this argument for each $k \in \{1, \dots, m\}$, we see that every α_k has to vanish. □

This immediately implies the following useful criterion.

PROPOSITION 2.30. *If the $n \times n$ matrix $\boldsymbol{A}$ has n pairwise distinct eigenvalues, then it is diagonalizable.*

PROOF. Let $\lambda_1, \dots, \lambda_n$ be the pairwise distinct eigenvalues of $\boldsymbol{A}$. Choose an eigenvector $\boldsymbol{v}_i$ for each eigenvalue λ_i. By Lemma 2.29 this is a basis. □

To study the diagonalization procedure in general, we need the following

DEFINITION 2.31. Let λ be an eigenvalue of $\boldsymbol{A}$. The space

$$\mathrm{Eig}(\boldsymbol{A}, \lambda) := \ker(\boldsymbol{A} - \lambda \mathbf{1})$$

is called the eigenspace of $\boldsymbol{A}$ associated with λ.[3] The dimension

$$d := \dim \mathrm{Eig}(\boldsymbol{A}, \lambda)$$

is called the geometric multiplicity of λ.

REMARK 2.32. One can show that $d \leq s$ for every eigenvalue, where s is the algebraic multiplicity and d is the geometric multiplicity. (We will prove this as Corollary 2.52 in Section 2.4.1.)

We have the following generalization of Proposition 2.30.

THEOREM 2.33. *Let $\lambda_1, \dots, \lambda_k$ be the pairwise distinct eigenvalues of the $n \times n$ matrix $\boldsymbol{A}$ and let d_i denote the geometric multiplicity of λ_i. Then $\boldsymbol{A}$ is diagonalizable iff $d_1 + \cdots + d_k = n$. In this case, we have*

$$\mathbb{K}^n = \mathrm{Eig}(\boldsymbol{A}, \lambda_1) \oplus \cdots \oplus \mathrm{Eig}(\boldsymbol{A}, \lambda_k).$$

PROOF. We use the criterion of Remark 1.22 to show that the sum $\mathrm{Eig}(\boldsymbol{A}, \lambda_1) + \cdots + \mathrm{Eig}(\boldsymbol{A}, \lambda_k)$ is direct.

Suppose we have $\boldsymbol{v}_1 + \cdots + \boldsymbol{v}_k = 0$, $\boldsymbol{v}_i \in \mathrm{Eig}(\boldsymbol{A}, \lambda_i)$. If some of the $\boldsymbol{v}_i$s were different from zero, then we would have a zero linear combination of eigenvectors corresponding to pairwise distinct eigenvalues, which is in contradiction with Lemma 2.29. Therefore, the zero vector has a unique decomposition and the sum is direct.

By (1.2), the direct sum $\mathrm{Eig}(\boldsymbol{A}, \lambda_1) \oplus \cdots \oplus \mathrm{Eig}(\boldsymbol{A}, \lambda_k)$ has dimension $d_1 + \cdots + d_k$. By Proposition 1.49, it is then the whole space $\mathbb{K}^n$ iff $d_1 + \cdots + d_k = n$.

In this case, for $i = 1, \dots, k$ let $(\boldsymbol{v}_{i,j})_{j=1,\dots,d_i}$ be a basis of $\mathrm{Eig}(\boldsymbol{A}, \lambda_i)$. By Remark 1.51, the union of these bases is a basis of the whole space. Since every $\boldsymbol{v}_{i,j}$ is an eigenvector, $\boldsymbol{A}$ is then diagonalizable by Theorem 2.17. □

[3] In principle we may define $\mathrm{Eig}(\boldsymbol{A}, \lambda)$ for any scalar λ. However, if λ is not an eigenvalue, we have $\mathrm{Eig}(\boldsymbol{A}, \lambda) = 0$.

Digression 2.34. By Remark 2.32, one then also has that $\boldsymbol{A}$ is diagonalizable iff the geometric multiplicity of every eigenvalue is equal to its algebraic multiplicity.

The procedure for the diagonalization of a square matrix $\boldsymbol{A}$ is then the following:

Step 1. Find all the pairwise distinct roots $\lambda_1, \ldots, \lambda_k$ of the characteristic polynomial $P_{\boldsymbol{A}}$.

Step 2. For every root λ_i choose a basis of $\mathrm{Eig}(\boldsymbol{A}, \lambda_i)$ and use it compute the dimension d_i.

Step 3. If $d_1 + \cdots + d_k = n$, then we have found a basis of eigenvectors and $\boldsymbol{A}$ is diagonalizable.

Remark 2.35. We have seen in Remark 2.13 the example $\boldsymbol{A} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$ of a nondiagonalizable matrix. Let us see what goes wrong with the diagonalization procedure. The characteristic polynomial is $P_{\boldsymbol{A}} = \det \boldsymbol{A} = \left(\begin{smallmatrix} -\lambda & 1 \\ 0 & -\lambda \end{smallmatrix}\right) = \lambda^2$. We therefore only have the eigenvalue $\lambda = 0$, which comes with algebraic multiplicity 2. An eigenvector $\left(\begin{smallmatrix} a \\ b \end{smallmatrix}\right)$ must then satisfy $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)\left(\begin{smallmatrix} a \\ b \end{smallmatrix}\right) = 0$, i.e., $b = 0$. The eigenspace of $\lambda = 0$ is then the span of $\left(\begin{smallmatrix} 1 \\ 0 \end{smallmatrix}\right)$, which shows that the geometric multiplicity is 1. In particular, $\mathrm{Eig}\left(\left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right), 0\right) \subsetneq \mathbb{C}^2$.

2.2.1. Digression: The real case. Suppose now that $\boldsymbol{A}$ is a real $n \times n$ matrix—e.g., the coefficient matrix of a real system of ODEs. To proceed, we regard it as a complex matrix. Its eigenvalues may then be complex numbers.

Remark 2.36. Since $P_{\boldsymbol{A}}$ is in this case a real polynomial, any complex root comes with its complex conjugate root. Therefore, if λ is not real, we also have a distinct eigenvalue $\bar{\lambda}$. Suppose that $\boldsymbol{z} \in \mathbb{C}^n$ is an eigenvector to the eigenvalue λ, i.e., $\boldsymbol{A}\boldsymbol{z} = \lambda \boldsymbol{z}$. By taking complex conjugation, we get $\boldsymbol{A}\bar{\boldsymbol{z}} = \bar{\lambda}\bar{\boldsymbol{z}}$, so $\bar{\boldsymbol{z}}$ is an eigenvector to the eigenvalue $\bar{\lambda}$.

Now suppose that $\boldsymbol{A}$ is diagonalizable as a complex matrix. We want to show that one can find a convenient basis of real vectors associated to the complex eigenvectors so that we can bring $\boldsymbol{A}$ in a convenient normal form.

To proceed let us introduce the following notation. We denote by $\mathrm{Eig}_{\mathbb{C}}(\boldsymbol{A}, \lambda) \subseteq \mathbb{C}^n$ the eigenspace to the eigenvalue λ of $\boldsymbol{A}$ as a complex matrix. If λ is real, we denote by $\mathrm{Eig}_{\mathbb{R}}(\boldsymbol{A}, \lambda) \subseteq \mathbb{R}^n$ the eigenspace to the eigenvalue λ of $\boldsymbol{A}$ as a real matrix.

If λ is a real eigenvalue, one can show that $\dim_{\mathbb{R}} \mathrm{Eig}_{\mathbb{R}}(\boldsymbol{A}, \lambda) = \dim_{\mathbb{C}} \mathrm{Eig}_{\mathbb{C}}(\boldsymbol{A}, \lambda)$.

If λ is not real, then we write $\lambda = \alpha + \mathrm{i}\beta$, with α and β its real and imaginary parts. Let $\boldsymbol{v}$ be a λ-eigenvector. Then, by Remark 2.36, $\bar{\boldsymbol{v}}$ is a $\bar{\lambda}$-eigenvector. As $\boldsymbol{v}$ and $\bar{\boldsymbol{v}}$ belong to different eigenspaces, they are linearly independent. Now let $\boldsymbol{u}$ and $\boldsymbol{w}$ be the real and imaginary parts of $\boldsymbol{v}$, i.e., $\boldsymbol{v} = \boldsymbol{u} + \mathrm{i}\boldsymbol{w}$ with $\boldsymbol{u}$ and $\boldsymbol{w}$ real vectors. We also have $\bar{\boldsymbol{v}} = \boldsymbol{u} - \mathrm{i}\boldsymbol{w}$. Note that $(\boldsymbol{u}, \boldsymbol{w})$ is a basis of $\mathrm{Span}_{\mathbb{C}}\{\boldsymbol{v}, \bar{\boldsymbol{v}}\}$. Moreover, we have

$$\begin{aligned}\boldsymbol{A}\boldsymbol{u} &= \boldsymbol{A}\frac{\boldsymbol{v}+\bar{\boldsymbol{v}}}{2} = \frac{\lambda\boldsymbol{v}+\bar{\lambda}\bar{\boldsymbol{v}}}{2} = \alpha\boldsymbol{u} - \beta\boldsymbol{w},\\ \boldsymbol{A}\boldsymbol{w} &= \boldsymbol{A}\frac{\boldsymbol{v}-\bar{\boldsymbol{v}}}{2\mathrm{i}} = \frac{\lambda\boldsymbol{v}-\bar{\lambda}\bar{\boldsymbol{v}}}{2\mathrm{i}} = \beta\boldsymbol{u} + \alpha\boldsymbol{w}.\end{aligned}$$

Therefore, the restriction of $\boldsymbol{A}$ to $\mathrm{Span}_{\mathbb{R}}\{\boldsymbol{u}, \boldsymbol{w}\}$ has the representing matrix $\left(\begin{smallmatrix}\alpha & \beta\\ -\beta & \alpha\end{smallmatrix}\right)$ in the basis $(\boldsymbol{u}, \boldsymbol{w})$.

If we now do this for every nonreal eigenvalue $\lambda_i = \alpha_i + \mathrm{i}\beta_i$, α_i and β_i real, passing from the basis $(\boldsymbol{v}_{i,1}, \dots, \boldsymbol{v}_{i,d_i}, \bar{\boldsymbol{v}}_{i,1}, \dots, \bar{\boldsymbol{v}}_{i,d_i})$ of $\mathrm{Eig}_{\mathbb{C}}(\boldsymbol{A}, \lambda) \oplus \mathrm{Eig}_{\mathbb{C}}(\boldsymbol{A}, \bar{\lambda})$ to the basis $(\boldsymbol{u}_{i,1}, \boldsymbol{w}_{i,1}, \dots, \boldsymbol{u}_{i,d_i}, \boldsymbol{w}_{i,d_i})$, with $\boldsymbol{u}_{i,j}$ and $\boldsymbol{w}_{i,j}$ the real and imaginary parts of $\boldsymbol{v}_{i,j}$, we get the

PROPOSITION 2.37. *Let $\boldsymbol{A}$ be a real $n \times n$ matrix that is diagonalizable as a complex matrix. Then there is an invertible real matrix $\boldsymbol{S}$ such that $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ has the block form*

$$\begin{pmatrix}\lambda_1 &&&&&\\ &\ddots&&&&\\ &&\lambda_r&&&\\ &&&\boldsymbol{B}_1&&\\ &&&&\ddots&\\ &&&&&\boldsymbol{B}_s\end{pmatrix},$$

where $\lambda_1, \dots, \lambda_r$ are the real eigenvalues of $\boldsymbol{A}$, and the $\boldsymbol{B}_j$s are 2×2 blocks of the form

$$\boldsymbol{B}_j = \begin{pmatrix}\alpha_j & \beta_j\\ -\beta_j & \alpha_j\end{pmatrix}, \qquad \alpha_j, \beta_j \in \mathbb{R},$$

corresponding each to a pair $(\lambda_j, \bar{\lambda}_j)$ of conjugate nonreal eigenvalues.

REMARK 2.38. If we want to apply this result to the solution of linear system of ODEs with constant coefficients, we have to compute the exponential of t times the block matrix in the proposition. To do so, it is enough to compute the exponential of each block. For each of the real eigenvalues, we then simply get the 1×1 block $\mathrm{e}^{\lambda_i t}$. For each block $\boldsymbol{B}_j$ we observe that $\boldsymbol{B}_j = \alpha_j\boldsymbol{1} + \beta_j\left(\begin{smallmatrix}0 & 1\\ -1 & 0\end{smallmatrix}\right)$. By (2.15), we then get

$$\mathrm{e}^{\boldsymbol{B}_j t} = \mathrm{e}^{\alpha_j t}\mathrm{e}^{\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)\beta_j t}.$$

It is a useful exercise to check that[4]

$$\mathrm{e}^{\left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)x} = \begin{pmatrix} \cos x & \sin x \\ -\sin x & \cos x \end{pmatrix}.$$

Therefore,

$$\mathrm{e}^{\boldsymbol{B}_j t} = \mathrm{e}^{\alpha_j t}\begin{pmatrix} \cos(\beta_j t) & \sin(\beta_j t) \\ -\sin(\beta_j t) & \cos(\beta_j t) \end{pmatrix}.$$

This shows that the solution to a linear system of ODEs with constant coefficients whose coefficient matrix is diagonalizable as a complex matrix can be written in terms of real exponentials and trigonometric functions as announced at the end of Remark 2.26.

2.3. Diagonalization of endomorphisms

The problem of diagonalization and the results we have discussed for matrices can be generalized to endomorphisms. We present them here, also as an occasion to recapitulate what we have seen.

Definition 2.39. An endomorphism F of a vector space V is called diagonalizable if there is a basis $\mathcal{B}$ such that $F_{\mathcal{B}}$ is a diagonal matrix.

Note that if v is an element of such a basis, we have $v \neq 0$ and $Fv = \lambda v$ for some scalar λ.

Definition 2.40 (Eigenvectors and eigenvalues). A nonzero vector v is called an eigenvector of an endomorphism F if there is a scalar λ, called the eigenvalue to the eigenvector v, such that $Fv = \lambda v$.

We clearly have the

Theorem 2.41. *An endomorphism F of V is diagonalizable iff V admits a basis of eigenvectors of F.*

If V is finite-dimensional, by Definition 1.74, we may define the characteristic polynomial of an endomorphism F as

$$P_F := \det(F - \lambda\,\mathrm{Id}).$$

We then have the

Lemma 2.42. *The eigenvalues of an endomoprhism of a finite-dimensional space are the roots of its characteristic polynomial.*

[4]This can either be done by explicit resumming the exponential series or by solving the harmonic oscillator with $k = m = 1$.

We also have the following generalization of Lemma 2.29 (with essentially the same proof).

LEMMA 2.43. *A collection of eigenvectors of an endomorphism corresponding to pairwise distinct eigenvalues is linearly independent.*

This implies again the

PROPOSITION 2.44. *If an endomorphism of an n-dimensional space has n pairwise distinct eigenvalues, then it is diagonalizable.*

We also associate to an eigenvalue λ of $F \in \operatorname{End}(V)$ its eigenspace

$$\operatorname{Eig}(F, \lambda) := \ker(F - \lambda \operatorname{Id})$$

and its geometric multiplicity

$$d := \dim \operatorname{Eig}(F, \lambda).$$

We have $\operatorname{Eig}(F, \lambda) \cap \operatorname{Eig}(F, \mu) = 0$ for $\lambda \neq \mu$ and the following

THEOREM 2.45. *Let $\lambda_1, \dots, \lambda_k$ be the pairwise distinct eigenvalues of an endomorphism F of an n-dimensional space V and let d_i denote the geometric multiplicity of λ_i. Then F is diagonalizable iff $d_1 + \cdots + d_k = n$. In this case, we have*

$$V = \operatorname{Eig}(\boldsymbol{A}, \lambda_1) \oplus \cdots \oplus \operatorname{Eig}(\boldsymbol{A}, \lambda_k).$$

If P_F splits into linear factor (e.g., if the ground field is $\mathbb{C}$) as

$$P_F = (-1)^n (\lambda - \lambda_1)^{s_1} \cdots (\lambda - \lambda_k)^{s_k},$$

then s_i is called the algebraic multiplicity of λ_i.

REMARK 2.46. For every eigenvalue λ_i, we have $d_i \leq s_i$. Therefore, F is diagonalizable iff $d_i = s_i$ for every i.

In applications it is often important to diagonalize two different endomorphisms at the same time. Of course one has first of all to assume that each of them is diagonalizable. We say that two diagonalizable endomorphisms F and G on a vector space V are simultaneously diagonalizable if they possess a common basis of eigenvectors.

PROPOSITION 2.47 (Simultaneous diagonalization). *Two diagonalizable endomorphisms F and G on a vector space V are simultaneously diagonalizable iff they commute, i.e., $FG = GF$.*

PROOF. See Exercise 2.10 □

In the case of matrices, the above proposition reads more explictly as follows.

COROLLARY 2.48. *Two diagonalizable matrices* $\boldsymbol{A}$ *and* $\boldsymbol{B}$ *commute (i.e.,* $\boldsymbol{AB} = \boldsymbol{BA}$*) iff there is an invertible matrix* $\boldsymbol{S}$ *such that* $\boldsymbol{S}^{-1}\boldsymbol{AS} = \boldsymbol{D}$ *and* $\boldsymbol{S}^{-1}\boldsymbol{BS} = \boldsymbol{D}'$ *where* $\boldsymbol{D}$ *and* $\boldsymbol{D}'$ *are diagonal matrices.*

2.3.1. The spectral decomposition. Suppose $F \in \operatorname{End}(V)$ is diagonalizable. If we decompose $V = \operatorname{Eig}(\boldsymbol{A}, \lambda_1) \oplus \cdots \oplus \operatorname{Eig}(\boldsymbol{A}, \lambda_k)$, we have, as in every direct sum, a unique decomposition of every $v \in V$ as $v = w_1 + \cdots + w_k$ with $w_i \in \operatorname{Eig}(\boldsymbol{A}, \lambda_i)$. We let $P_i \colon V \to V$ be the linear map that assigns to a vector v its ith component w_i. The P_is are a complete system of mutually transversal projections; i.e.,

$$P_i^2 = P_i \ \forall i, \qquad P_iP_j = P_jP_i = 0 \ \forall i \neq j, \quad \sum_{i=1}^{k} P_i = \operatorname{Id}.$$

Since the ith component of a vector is an eigenvector to λ_i, we have $FP_iv = Fw_i = \lambda_i w_i = \lambda_i P_i v$ for every $v \in V$. As an identity of maps, this reads $FP_i = \lambda_i P_i$. Summing over i, we then get

$$F = \sum_{i=1}^{k} \lambda_i P_i. \tag{2.21}$$

The set $\{\lambda_1, \dots, \lambda_k\}$ of the pairwise distinct eigenvalues of F is called its spectrum and (2.21) is called the spectral decomposition of F.[5]

2.3.2. The infinite-dimensional case. In Example 2.7, we have seen that some important PDEs in physics are linear and can be viewed as linear ODEs on some infinite-dimensional space.

We will not treat this case here in general, but we will consider an example: the wave equation on a one-dimensional space interval, a.k.a. the vibrating string. Namely, we want to study the PDE

$$\frac{1}{c^2}\frac{\partial^2\psi}{\partial t^2} = \frac{\partial^2\psi}{\partial x^2},$$

where the unknown ψ is a function on $\mathbb{R} \times [0, L] \ni (t, x)$.[6] We assume that the string endpoints are fixed:

$$\psi(t, 0) = \psi(t, L) = 0 \quad \text{for all } t$$

[5] This terminology ultimately comes from physics, namely from the fact that the spectral lines of an atom are computed in quantum mechanics by taking differences of eigenvalues of a certain endomorphism, the Hamiltonian operator of the electrons in the atom.

[6] More generally, we could consider an interval $[a, b]$, but translating it to $[0, L]$, with $L = b - a$ its length, simplifies the discussion.

(this models, e.g., the string of a musical instrument). We then introduce the infinite-dimensional vector space

$$V := \{\phi \in C^\infty([0,L]) \mid \phi(0) = \phi(L) = 0\}$$

and regard ψ as a map $\mathbb{R} \to V$.

The right hand side of the wave equation uses the endomorphism $F := \frac{\mathrm{d}^2}{\mathrm{d}x^2}$ of V. We now want to find its eigenvectors, i.e., $\phi \in V \setminus \{0\}$ such that

$$\phi'' = \lambda\phi$$

for some complex scalar λ. Denoting by $\pm\alpha$ the two square roots of λ, we see that the general solution to this equation is

$$\phi(x) = A\mathrm{e}^{\alpha x} + B\mathrm{e}^{-\alpha x},$$

where A and B are complex constants. The endpoint conditions $\phi(0) = \phi(L) = 0$ amount to the linear system

$$\left\{\begin{array}{ccccc} A & + & B & = & 0 \\ A\mathrm{e}^{\alpha L} & + & B\mathrm{e}^{-\alpha L} & = & 0 \end{array}\right.$$

which has a nontrivial solution ($B = -A \neq 0$) iff the coefficient matrix $\begin{pmatrix} 1 & 1 \\ \mathrm{e}^{\alpha L} & \mathrm{e}^{-\alpha L} \end{pmatrix}$ is degenerate. Since its determinant is $\mathrm{e}^{-\alpha L} - \mathrm{e}^{\alpha L}$, we may have a nontrivial solution iff $\mathrm{e}^{2\alpha L} = 1$, i.e., $\alpha = \frac{\mathrm{i}\pi k}{L}$ with k an integer.

The case $k = 0$ yields $\phi = 0$, which is not an eigenvector. For $k \neq 0$, we take $A = \frac{1}{2\mathrm{i}}$ (and hence $B = -\frac{1}{2\mathrm{i}}$), so we have the real eigenvector

$$\phi_k(x) = \sin\left(\frac{\pi k x}{L}\right)$$

corresponding to the eigenvalue $\lambda = -\frac{\pi^2 k^2}{L^2}$. Note that $\phi_{-k} = -\phi_k$, so they are not linearly independent. Therefore, we only consider $k > 0$.

By Lemma 2.43, the collection $(\phi_k)_{k \in \mathbb{Z}_{>0}}$ is linearly independent in V. Therefore, on the subspace V' spanned by this collection, we have that the set of eigenvalues (a.k.a. the spectrum) is

$$\left\{-\frac{\pi^2 k^2}{L^2},\ k \in \mathbb{N}_{>0}\right\}.$$

If the two initial conditions, $\psi|_{t=0}$ and $\frac{\partial\psi}{\partial t}|_{t=0}$, for the wave equation are linear combinations of the ϕ_ks, we may then write the unique solution to the Cauchy problem as a linear combination of the ϕ_ks with time-dependent coefficients. In fact, suppose

$$\psi|_{t=0} = \sum_{k=1}^{\infty} b_{k0}\phi_k,$$
$$\frac{\partial\psi}{\partial t}|_{t=0} = \sum_{k=1}^{\infty} v_{k0}\phi_k,$$

where only finitely many of the b_{k0}s and of the v_{k0}s are different from zero. We can then consider a solution of the form

$$\psi = \sum_{k=1}^{\infty} b_k\phi_k,$$

where the coefficients b_k are now functions of time, and the sum is restricted to the ks for which $b_{k0} \neq 0$ or $v_{k0} \neq 0$. The wave equation then yields separate ODEs for each of these ks, which we can assemble into the Cauchy problems

$$\begin{cases} \ddot{b}_k &= -\frac{\pi^2 c^2 k^2}{L^2} b_k \\ b_k(0) &= b_{k0} \\ \dot{b}_k(0) &= v_{k0} \end{cases}$$

Interestingly, it turns out that one can also make sense of infinite linear combinations $\sum_{k=1}^{\infty} b_k \sin\left(\frac{\pi k x}{L}\right)$, where (b_k) is a sequence of real numbers with appropriate decaying conditions for $k \to \infty$. This is an example of Fourier series. We will return to this in Example **??**, where we will also learn a method to compute the coefficients b_k of an expansion.

2.4. Trigonalization

Even though not every matrix can be diagonalized, it turns out complex matrices can be brought to a nice upper triangular form. More precisely, we have the

THEOREM 2.49. *Let $\boldsymbol{A}$ be a $n \times n$ matrix whose characteristic polynomial splits into linear factors (e.g., a complex matrix),*

$$P_{\boldsymbol{A}} = (-1)^n (\lambda - \lambda_1)^{s_1} \cdots (\lambda - \lambda_k)^{s_k}. \tag{2.22}$$

Then there is an invertible matrix $\boldsymbol{S}$ such that

$$\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \boldsymbol{D} + \boldsymbol{N},$$

where $\boldsymbol{D}$ is a diagonal matrix with the eigenvalues of $\boldsymbol{A}$ as its diagonal entries, $\boldsymbol{N}$ is an upper triangular matrix with zeros on the diagonal, and $\boldsymbol{D}\boldsymbol{N} = \boldsymbol{N}\boldsymbol{D}$.

Before we prove the theorem, let us see its consequences for the exponential of a real or complex matrix $\boldsymbol{A}$ (and hence for the associated system of ODEs). By (2.16), we have that $\mathrm{e}^{\boldsymbol{A}} = \boldsymbol{S}\mathrm{e}^{\boldsymbol{D}+\boldsymbol{N}}\boldsymbol{S}^{-1}$. By (2.15), we have $\mathrm{e}^{\boldsymbol{D}+\boldsymbol{N}} = \mathrm{e}^{\boldsymbol{D}}\mathrm{e}^{\boldsymbol{N}}$. We already know how to compute the exponential of a diagonal matrix, so we are only left with the exponential of $\boldsymbol{N}$. Observe that $\boldsymbol{N}$ applied to a vector whose last $k < n$ components are equal to 0 yields a vector whose last $k+1$ components are equal to 0. Therefore, $\boldsymbol{N}^n$ applied to any vector yields the zero vector. In conclusion, $\boldsymbol{N}^m = 0$ for all $m \geq n$. One says that $\boldsymbol{N}$ is nilpotent (meaning that it has a vanishing power). It follows that $\mathrm{e}^{\boldsymbol{N}}$ is a finite sum. These results do not change if we multiply $\boldsymbol{A}$ by t, so we get

$$\mathrm{e}^{\boldsymbol{A}t} = \boldsymbol{S}^{-1}\mathrm{e}^{\boldsymbol{D}t}\left(\sum_{r=0}^{n-1}\frac{t^r}{r!}\boldsymbol{N}^r\right)\boldsymbol{S}.$$

In particular, this means that a solution of the associated system of ODEs is a combination of exponentials and polynomials.

Also note that $\mathrm{e}^{\boldsymbol{N}}$ is an upper triangular matrix with 1s on the diagonal, so $\det \mathrm{e}^{\boldsymbol{N}} = 1$. Therefore, $\det \mathrm{e}^{\boldsymbol{A}t} = \det \mathrm{e}^{\boldsymbol{D}t}$. Since $\boldsymbol{D}$ is diagonal, we obviously have $\det \mathrm{e}^{\boldsymbol{D}t} = \mathrm{e}^{t \operatorname{tr} \boldsymbol{D}}$. On the other hand, we have $\operatorname{tr} \boldsymbol{A} = \operatorname{tr} \boldsymbol{D}$. In conclusion,

$$\det \mathrm{e}^{\boldsymbol{A}t} = \mathrm{e}^{t \operatorname{tr} \boldsymbol{A}}$$

for every t. This is a purely algebraic proof of Proposition 2.11. (In the theorem we assume that $P_{\boldsymbol{A}}$ splits into linear factors, which might not be the case if $\boldsymbol{A}$ is real. In this case, we can however regard $\boldsymbol{A}$ as a complex matrix, apply the theorem, and get the last identity; finally, we observe that both the left and the right hand side are defined over $\mathbb{R}$.)

2.4.1. Proof of Theorem 2.49. The eigenspace of $\boldsymbol{A}$ associated with the eigenvalue λ may be viewed as the largest subspace on which the restriction of $\boldsymbol{A} - \lambda\mathbf{1}$ is zero. Since we are looking for a basis in which the representing matrix minus the diagonal matrix with the eigenvalues as its diagonal entries is nilpotent, we define the generalized eigenspace associated with the eigenvalue λ as

$$\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda] := \{\boldsymbol{v} \mid \exists m \in \mathbb{N}\ (\boldsymbol{A} - \lambda\mathbf{1})^m\boldsymbol{v} = 0\}.$$

More precisely, we say that $\boldsymbol{v}$ is a generalized λ-eigenvector of rank $m > 0$ if $(\boldsymbol{A} - \lambda\mathbf{1})^m\boldsymbol{v} = 0$ but $(\boldsymbol{A} - \lambda\mathbf{1})^{m-1}\boldsymbol{v} \neq 0$. Note that every nonzero vector in $\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda]$ is a generalized λ-eigenvector with a well-defined rank. In particular, an eigenvector in the original sense is a generalized eigenvector of rank 1, so

$$\operatorname{Eig}(\boldsymbol{A},\lambda) \subseteq \widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda].$$

Also observe that $\boldsymbol{A}-\lambda\mathbf{1}$ applied to a generalized λ-eigenvector of rank $m>1$ yields a generalized λ-eigenvector of rank $m-1$.

Finally, note that $\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda]$ is $\boldsymbol{A}$-invariant, so also $(\boldsymbol{A}-\lambda\mathbf{1})$-invariant. We denote by $\boldsymbol{N}_\lambda$ the restriction of $\boldsymbol{A}-\lambda\mathbf{1}$ to $\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda]$. By the definition of generalized eigenspace, $\boldsymbol{N}_\lambda$ is nilpotent.

Let $(\boldsymbol{v}_1,\dots,\boldsymbol{v}_d)$ be a basis of $\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda]$. We order the basis in such a way that the rank of $\boldsymbol{v}_j$ is less than or equal to the rank of $\boldsymbol{v}_{j+1}$ for $j=1,\dots,d-1$. It follows that $\boldsymbol{N}_\lambda$ is represented in this basis by an upper triangular matrix with zeros on the diagonal. Moreover, $\boldsymbol{N}_\lambda$ clearly commutes with $\lambda\operatorname{Id}_{\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda]}$. We are then done after proving the following

Proposition 2.50. *Under the assumptions of Theorem 2.49:*

(1) the algebraic multiplicity s_i of λ_i is equal to the dimension δ_i of $\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_i]$ for every i, and
(2) we have the decomposition $\mathbb{K}^n=\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_1]\oplus\cdots\oplus\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_k]$.

In fact, it is enough to choose a basis of each generalized eigenspace, ordered by rank as above. In the basis of $\mathbb{K}^n$ given by the union of these bases, $\boldsymbol{A}$ is represented by a matrix $\boldsymbol{D}+\boldsymbol{N}$ as in the theorem.[7]

We start by proving the

Lemma 2.51. *The sum $\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_1]+\cdots+\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_k]$ is direct.*

Proof. We use the criterion of Remark 1.22. Suppose we have $\boldsymbol{v}_1+\cdots+\boldsymbol{v}_k=0$, $\boldsymbol{v}_i\in\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_i]$.

Suppose, by contradiction, that some $\boldsymbol{v}_i$ is different from zero, and let $m_i>0$ be its rank. Since $\boldsymbol{w}_i:=(\boldsymbol{A}-\lambda_i\mathbf{1})^{m_i-1}\boldsymbol{v}_i\neq 0$ is an eigenvector for λ_i, we have

$$\prod_{j=1}^{k}(\boldsymbol{A}-\lambda_j\mathbf{1})^{m_j-1}\boldsymbol{v}_i=\prod_{j\neq i}(\boldsymbol{A}-\lambda_j\mathbf{1})^{m_j-1}\boldsymbol{w}_i=\prod_{j\neq i}(\lambda_i-\lambda_j)^{m_j-1}\boldsymbol{w}_i=:\boldsymbol{z}_i.$$

Since $\prod_{j\neq i}(\lambda_i-\lambda_j)^{m_j-1}$ is different from zero, we get that $\boldsymbol{z}_i\neq 0$.

If we now apply $\prod_{j=1}^{k}(\boldsymbol{A}-\lambda_j\mathbf{1})^{m_j}$ to $\boldsymbol{v}_1+\cdots+\boldsymbol{v}_k=0$, we get $\sum_{r:\boldsymbol{v}_r\neq 0}\boldsymbol{z}_r=0$. Since the $\boldsymbol{z}_r$s in the sum are also different from zero, this is in contradiction with Lemma 2.29.

[7]In particular, the restriction of $\boldsymbol{D}$ to $\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_j]$ is $\lambda_j\operatorname{Id}_{\widetilde{\operatorname{Eig}}[\boldsymbol{A},\lambda_j]}$, whereas the restriction of $\boldsymbol{N}$ is $\boldsymbol{N}_{\lambda_j}$.

Therefore, the zero vector has a unique decomposition and the sum is direct.

∎

Proof of Proposition 2.50. Note that part (1) of the statement implies part (2) using Lemma 2.51. In fact,

$$\dim(\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_1]\oplus\cdots\oplus\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_k])=\delta_1+\cdots+\delta_k=s_1+\ldots s_k=n,$$

so the direct sum is the whole space $\mathbb{K}^n$.

Therefore, we will only prove statement (1) by induction on the dimension n. For $n=1$ there is nothing to prove, as in this case $\boldsymbol{A}=(\lambda)$ and $\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda]=\mathrm{Eig}(\boldsymbol{A},\lambda)=\mathbb{K}$.

Next, we assume we have proved (1), and hence (2), for dimensions up to $n-1$. We pick an $i\in\{1,\ldots,k\}$, denote by $\boldsymbol{A}_i$ the restriction of $\boldsymbol{A}$ to $\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]$, which is $\boldsymbol{A}$-invariant, and choose a complement W to $\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]$ in $\mathbb{K}^n$. With respect to the decomposition $\mathbb{K}^n=\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]\oplus W$, $\boldsymbol{A}$ has the form

$$\begin{pmatrix}\boldsymbol{A}_i & \boldsymbol{B}\\ \boldsymbol{0} & \boldsymbol{C}\end{pmatrix},$$

where $\boldsymbol{B}$ is the composition of the restriction of $\boldsymbol{A}$ to W with the projection to $\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]$ and $\boldsymbol{C}$ is the composition of the restriction of $\boldsymbol{A}$ to W with the projection to W.

By property (D.11) of the determinant, we have $P_{\boldsymbol{A}}=P_{\boldsymbol{A}_i}P_{\boldsymbol{C}}$. If we take a rank-ordered basis of $\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]$, we have that $\boldsymbol{A}_i$ is represented by an upper triangular matrix with diagonal entries equal to λ_i. Therefore, by property (D.6), $P_{\boldsymbol{A}_i}=(-1)^{\delta_i}(\lambda-\lambda_i)^{\delta_i}$. In conclusion, $P_{\boldsymbol{A}}=(-1)^{\delta_i}(\lambda-\lambda_i)^{\delta_i}P_{\boldsymbol{C}}$. Comparing with (2.22), we have that

$$P_{\boldsymbol{C}}=(-1)^{\dim W}(\lambda-\lambda_1)^{s_1'}\cdots(\lambda-\lambda_k)^{s_k'}$$

with $s_j'=s_j$ for $j\neq i$ and $s_i'=s_i-\delta_i$. We claim that $s_i'=0$, which in particular implies $\delta_i=s_i$. Since this can be done for every i, this completes the proof of the proposition.

We now prove the claim that $s_i'=0$. Since $\dim W<n$ and $P_{\boldsymbol{C}}$ splits into linear factors, we may apply the induction hypothesis, so[8] $W=\widetilde{\mathrm{Eig}}[\boldsymbol{C},\lambda_1]\oplus\cdots\oplus\widetilde{\mathrm{Eig}}[\boldsymbol{C},\lambda_k]$.

Consider the space $V=\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]\oplus\widetilde{\mathrm{Eig}}[\boldsymbol{C},\lambda_i]$. For $\boldsymbol{v}$ in the first summand, we have $\boldsymbol{A}\boldsymbol{v}\in\widetilde{\mathrm{Eig}}[\boldsymbol{A},\lambda_i]$; for $\boldsymbol{v}$ in the second summand, we have $\boldsymbol{A}\boldsymbol{v}=\boldsymbol{B}\boldsymbol{v}+\boldsymbol{C}_i\boldsymbol{v}$, where $\boldsymbol{C}_i$ denotes the restriction of $\boldsymbol{C}$ to $\widetilde{\mathrm{Eig}}[\boldsymbol{C},\lambda_i]$, which is a $\boldsymbol{C}$-invariant subspace. Therefore, the restriction

[8] We actually choose a basis to identify W with $\mathbb{K}^{\dim W}$ and apply statement (2) to the representing matrix of $\boldsymbol{C}$.

$\boldsymbol{A}_V$ of $\boldsymbol{A}$ to V has the following form with respect to the decomposition:[9]

$$\begin{pmatrix} \boldsymbol{A}_i & \boldsymbol{B} \\ \boldsymbol{0} & \boldsymbol{C}_i \end{pmatrix}.$$

One easily proves, by induction on r, that

$$(\boldsymbol{A}_V - \lambda_i \boldsymbol{1})^r = \begin{pmatrix} \boldsymbol{A}_i - \lambda_i \boldsymbol{1} & \boldsymbol{B} \\ \boldsymbol{0} & \boldsymbol{C}_i - \lambda_i \boldsymbol{1} \end{pmatrix}^r = \begin{pmatrix} (\boldsymbol{A}_i - \lambda_i \boldsymbol{1})^r & \boldsymbol{B}_r \\ \boldsymbol{0} & (\boldsymbol{C}_i - \lambda_i \boldsymbol{1})^r \end{pmatrix}$$

for some matrix $\boldsymbol{B}_r$. If $\boldsymbol{v} \in \widetilde{\mathrm{Eig}}[\boldsymbol{C}, \lambda_i]$ has rank r, we then have $(\boldsymbol{A} - \lambda_i \boldsymbol{1})^r \boldsymbol{v} = \boldsymbol{B}_r \boldsymbol{v}$. Since this is now an element of $\widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$, there is an s such that $(\boldsymbol{A} - \lambda_i \boldsymbol{1})^s \boldsymbol{B}_r \boldsymbol{v} = 0$. This means that $(\boldsymbol{A} - \lambda_i \boldsymbol{1})^{r+s} \boldsymbol{v} = 0$, i.e., $\boldsymbol{v} \in \widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$. Finally, since $\widetilde{\mathrm{Eig}}[\boldsymbol{C}, \lambda_i]$ belongs to a complement of $\widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$, the only vector in their intersection is 0. We have thus proved that $\widetilde{\mathrm{Eig}}[\boldsymbol{C}, \lambda_i] = 0$ and, hence, that $s'_i = 0$. □

Another interesting consequence of part (1) of Proposition 2.50 is the following result, announced in Remark 2.32.

COROLLARY 2.52. *The geometric multiplicity of every eigenvalue is less than or equal to its algebraic multiplicity.*

PROOF. Since $\mathrm{Eig}(\boldsymbol{A}, \lambda_i) \subseteq \widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$, we have $\dim \mathrm{Eig}(\boldsymbol{A}, \lambda_i) \leq \delta_i = s_i$. □

2.5. Digression: The Jordan normal form

By choosing a more suitable basis of each $\widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$, the strictly upper triangular matrix $\boldsymbol{N}$ of Theorem 2.49 may be put in a "canonical" form that can be more easily dealt with.

For this we need the notion of a Jordan block (after the French mathematician CAMILLE JORDAN). For a scalar λ and a positive integer m, the Jordan block $\boldsymbol{J}_{\lambda,m}$ is the $m \times m$ upper triangular matrix whose diagonal entries are equal to λ, the entries right above the diagonal are equal to 1, and all other entries vanish; e.g.,

$$\boldsymbol{J}_{\lambda,1} = \begin{pmatrix} \lambda \end{pmatrix}, \quad \boldsymbol{J}_{\lambda,2} = \begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}, \quad \boldsymbol{J}_{\lambda,3} = \begin{pmatrix} \lambda & 1 & 0 \\ 0 & \lambda & 1 \\ 0 & 0 & \lambda \end{pmatrix},$$

[9] By abuse of notation, we keep writing $\boldsymbol{B}$, but what appears here is actually the restriction of $\boldsymbol{B}$ to $\widetilde{\mathrm{Eig}}[\boldsymbol{C}, \lambda_i]$.

and

$$\boldsymbol{J}_{\lambda,4} = \begin{pmatrix} \lambda & 1 & 0 & 0 \\ 0 & \lambda & 1 & 0 \\ 0 & 0 & \lambda & 1 \\ 0 & 0 & 0 & \lambda \end{pmatrix}.$$

A Jordan matrix is a block diagonal matrix whose diagonal blocks are Jordan blocks.

THEOREM 2.53. *Let $\boldsymbol{A}$ be a matrix as in Theorem 2.49. Then there is a basis in which $\boldsymbol{A}$ is represented by a Jordan matrix. More precisely, for each eigenvalue λ_i, the generalized eigenspace $\widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$ has a basis in which the restriction of $\boldsymbol{A}$ is represented by a Jordan matrix of the form*

$$\begin{pmatrix} \boldsymbol{J}_{\lambda_i,m_1} & & \\ & \ddots & \\ & & \boldsymbol{J}_{\lambda_i,m_k} \end{pmatrix},$$

with $0 < m_1 \leq \cdots \leq m_k$ and $m_1 + \cdots + m_k = s_i$.

To prove the theorem, it is enough to prove the statement for each generalized eigenspace separately. By definition of generalized eigenspace, the restriction $\boldsymbol{N}_i$ of $\boldsymbol{A} - \lambda\mathbf{1}$ to $\widetilde{\mathrm{Eig}}[\boldsymbol{A}, \lambda_i]$ is nilpotent. It follows that it is enough to prove the following

PROPOSITION 2.54. *Let N be a nilpotent operator on an n-dimensional vector space V. Then there is a basis of V in which N is represented by a Jordan matrix of the form*

$$\begin{pmatrix} \boldsymbol{J}_{0,m_1} & & \\ & \ddots & \\ & & \boldsymbol{J}_{0,m_k} \end{pmatrix},$$

with $0 < m_1 \leq \cdots \leq m_k$ and $m_1 + \cdots + m_k = n$.

To prove the proposition, we need some preliminary remarks. The first remark is that a nilpotent operator is not invertible, so it has a nonzero kernel.

We say that a vector v in V has rank $m > 0$ if $N^m v = 0$ but $N^{m-1} v \neq 0$. For a vector v of rank m we define

$$v_j := N^{m-j} v$$

for $j = 1, \ldots, m$. The vectors $v_1, \ldots, v_m$ are called a Jordan chain (note that $v_m = v$ and that v_1 is in the kernel of N). We denote by $\mathrm{Jor}(v)$ their span:

$$\mathrm{Jor}(v) := \mathrm{Span}\{v_1, \ldots, v_m\}.$$

The next remark is that $\mathrm{Jor}(v)$ is N-invariant and that the vectors $v_1, \dots, v_m$ form a basis. That they generate $\mathrm{Jor}(v)$ is obvious by definition, so we only have to check that they are linearly independent. Suppose $\alpha^i v_i = 0$ for some scalars $\alpha^1, \dots, \alpha^m$. Applying N^{m-1} to this linear combination yields $\alpha^m N^{m-1} v = 0$, which implies $\alpha^m = 0$. If we then apply N^{m-2} to the linear combination, knowing that $\alpha^m = 0$ we get $\alpha^{m-1} N^{m-1} v = 0$ which implies $\alpha^{m-1} = 0$, and so on. In particular, this shows that $\dim \mathrm{Jor}(v) = m$.

The final remark is that, by construction, the restriction of N to $\mathrm{Jor}(v)$ in the basis $(v_1, \dots, v_m)$ is represented by the Jordan block $\boldsymbol{J}_{0,m}$.

The strategy to prove Proposition 2.54 consists then in decomposing V into spans of Jordan chains.

Note that the vector v of rank m we started with to define $\mathrm{Jor}(v)$ might be in the image of N, say, $v = Nw$. In this case, we may extend the Jordan chain $\mathrm{Jor}(v)$ to the Jordan chain $\mathrm{Jor}(w) \supsetneq \mathrm{Jor}(v)$ (note that $w_i = v_i$ for $j = 1, \dots, m$ and that $w_{m+1} = w$). If v is not in the image of N, we say that $v_1, \dots, v_m$ is a maximal Jordan chain and that v is a lead vector for it. Proposition 2.54 is then a consequence of the following

Lemma 2.55. *Let N be a nilpotent operator on an n-dimensional vector space V. Then there is a collection $v_{(1)}, \dots, v_{(k)}$ of lead vectors*[10] *of ranks $m_1, \dots, m_k$, respectively, such that $V = \mathrm{Jor}(v_{(1)}) \oplus \cdots \oplus \mathrm{Jor}(v_{(k)})$.*

Note that we may arrange the lead vectors $v_{(1)}, \dots, v_{(k)}$ so that $0 < m_1 \leq \cdots \leq m_k$ as in Proposition 2.54.

Proof of the lemma. We prove the lemma by induction on the dimension n of V. If $n = 0$, there is nothing to prove.

Next, assume we have proved the lemma up to dimension $n - 1$. Let $W \subseteq V$ be the image of N. Note that W is N-invariant. By the dimension formula (1.4), we have $\dim W = n - \dim \ker N < n$, since N has nonzero kernel, so we can apply the induction assumption to W. Namely, we can find vectors $v_{(1)}, \dots, v_{(k)}$ in W such that $W = \mathrm{Jor}(v_{(1)}) \oplus \cdots \oplus \mathrm{Jor}(v_{(k)})$. (The $v_{(i)}$s are lead vectors in W but not in V.)

We now have two cases to consider. The first case is when $W \cap \ker N = 0$.[11] In this case $V = W \oplus \ker N$. A basis $(z_{(1)}, \dots, z_{(l)})$ of $\ker N$ produces the decomposition $\ker N = \mathrm{Jor}(z_{(1)}) \oplus \mathrm{Jor}(z_{(l)})$ (note

[10] We write $v_{(i)}$ for the lead vectors and $v_{(i),j}$ for the vectors in the corresponding Jordan chain.

[11] This case in particular happens when $n = 1$.

that $\mathrm{Jor}(z_{(i)}) = \mathbb{K}z_{(i)}$, for $z_{(i)}$ is in the kernel of N). This concludes the proof in this case.

The other case is when $W \cap \ker N \neq 0$. Let W' be a complement of $W + \ker N$ in V. In particular, the restriction of N to W' is injective. Therefore, we have uniquely determined $w_{(1)}, \dots, w_{(k)}$ in W' satisfying $v_{(i)} = Nw_{(i)}$ for $i = 1, \dots, k$.

We claim that $(w_{(1)}, \dots, w_{(k)})$ is a basis of W'. This completes the proof of the lemma, since it yields the decomposition $V = \mathrm{Jor}(w_{(1)}) \oplus \dots \oplus \mathrm{Jor}(w_{(k)}) \oplus \mathrm{Jor}(z_{(1)}) \oplus \mathrm{Jor}(z_{(l)})$, where $(z_{(1)}, \dots, z_{(l)})$ is a basis of a complement of W in $W + \ker N$.

To prove the claim, take some v in W'. We have to show that it has a unique decomposition in $w_{(1)}, \dots, w_{(k)}$. Since Nv is in W, we may expand it in the basis $(N^{j_i} v_{(i)})_{i=1,\dots,k,\ j_i=0,\dots,m'_i-1}$, where m'_i is the rank of $v_{(i)}$ in W. Note that all these vectors but $v_{(1)}, \dots, v_{(k)}$ are in the image of N^2. Therefore, we have uniquely determined scalars $\alpha^1, \dots, \alpha^k$ and some vector w such that $Nv = \sum_i \alpha^i v_{(i)} + N^2 w$. Setting $\widetilde{v} := v - \sum_i \alpha^i w_{(i)}$, we get $N\widetilde{v} = N^2 w$, so $\widetilde{v} - Nw \in \ker N$ and hence $\widetilde{v} \in W + \ker N$. Since, however, $\widetilde{v} \in W'$, which is a complement of $W + \ker N$, we get $\widetilde{v} = 0$ and hence $v = \sum_i \alpha^i w_{(i)}$. □

Exercises for Chapter 2

2.1. Applying the formula $\mathrm{e}^{\boldsymbol{A}t} = \boldsymbol{S}\mathrm{e}^{\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}t}\boldsymbol{S}^{-1}$, compute $\mathrm{e}^{\boldsymbol{A}t}$ for $\boldsymbol{A} = \left(\begin{smallmatrix} 0 & 2 \\ 8 & 0 \end{smallmatrix}\right)$ using $\boldsymbol{S} = \left(\begin{smallmatrix} 1 & 1 \\ 2 & -2 \end{smallmatrix}\right)$.

2.2. Let $\boldsymbol{A} = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$.
   (a) Compute $\boldsymbol{A}^n$ for all integers $n > 0$.
   Hint: Distinguish the cases n even and n odd.
   (b) Using the above result and writing the exponential series as

$$\mathrm{e}^{\boldsymbol{A}t} = \sum_{s=0}^{\infty} \frac{1}{(2s)!} t^{2s} \boldsymbol{A}^{2s} + \sum_{s=0}^{\infty} \frac{1}{(2s+1)!} t^{2s+1} \boldsymbol{A}^{2s+1},$$

   show that

$$\mathrm{e}^{\boldsymbol{A}t} = \begin{pmatrix} \cos t & \sin t \\ -\sin t & \cos t \end{pmatrix}.$$

   Hint: Use the series expansions of sin and cos.

2.3. Let $\boldsymbol{A} = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$ and $\boldsymbol{B} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$. Compute $\mathrm{e}^{\boldsymbol{A}t}$, $\mathrm{e}^{\boldsymbol{B}t}$, $\mathrm{e}^{\boldsymbol{A}t}\mathrm{e}^{\boldsymbol{B}t}$, $\mathrm{e}^{\boldsymbol{B}t}\mathrm{e}^{\boldsymbol{A}t}$, and $\mathrm{e}^{(\boldsymbol{A}+\boldsymbol{B})t}$. (Hint: Proceed as in the previous exercise.)

2.4. Determine the characteristic polynomial of the following matrices, and find their eigenvalues and a basis of eigenvectors:

$$\boldsymbol{A} = \begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}, \quad \boldsymbol{B} = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad \boldsymbol{C} = \begin{pmatrix} 1 & -1 & -1 \\ -1 & 3 & -1 \\ -1 & -1 & 3 \end{pmatrix}.$$

2.5. Let $\boldsymbol{A} = \begin{pmatrix} 2 & 1 \\ 2 & 3 \end{pmatrix}$.

(a) Find the similarity transformation that diagonalizes $\boldsymbol{A}$, i.e., find a matrix $\boldsymbol{S}$ such that $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S} = \boldsymbol{D}$ is a diagonal matrix and compute $\boldsymbol{D}$ explicitly.

(b) Using your results just obtained, find the solution to the Cauchy problem given by

$$\frac{\mathrm{d}\boldsymbol{x}}{\mathrm{d}t} = \boldsymbol{A}\boldsymbol{x}, \qquad \boldsymbol{x}(0) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

2.6. In this exercise we prove that $P_{\boldsymbol{A}}(\boldsymbol{A}) = \boldsymbol{0}$ for any 2×2 matrix,[12] where $P_{\boldsymbol{A}}(\lambda)$ is the characteristic polynomial of $\boldsymbol{A}$ in λ. For this we can proceed as follows:

(a) Show that for any 2×2 matrix $\boldsymbol{A}$ the characteristic polynomial can be written as

$$P_{\boldsymbol{A}}(\lambda) = \lambda^2 - \operatorname{tr}(\boldsymbol{A})\lambda + \det(\boldsymbol{A}).$$

(b) We will now interpret it as a polynomial in the matrix $\boldsymbol{A}$. Show that

$$\boldsymbol{A}^2 - \operatorname{tr}(\boldsymbol{A})\,\boldsymbol{A} + \det(\boldsymbol{A})\,\boldsymbol{1} = \boldsymbol{0}.$$

2.7. Let J be an endomorphism of V satisfying $J^2 = \mathrm{id}$.

(a) Show that $\lambda = \pm 1$ are the only possible eigenvalues of J.

(b) Using

$$v = \frac{v + Jv}{2} + \frac{v - Jv}{2},$$

show that J is diagonalizable.

[12]The result actually holds for matrices of any size and is known as the Cayley–Hamilton theorem.

2.8. Motivated by the study of the vibrating string where the endpoints are free to slide frictionless in the vertical direction, we consider the endomorphism $F := \frac{\mathrm{d}^2}{\mathrm{d}x^2}$ of the vector space

$$V := \{\phi \in C^\infty([0,\ell]) \mid \phi'(0) = \phi'(\ell) = 0\}.$$

Find all eigenvalues of F and a corresponding linearly independent system of eigenvectors.

2.9. Let

$$\boldsymbol{A} = \begin{pmatrix} 5 & -1 & -2 \\ -1 & 5 & 2 \\ 0 & 0 & 6 \end{pmatrix}.$$

(a) Find all eigenvalues of $\boldsymbol{A}$.
(b) Find linearly independent eigenvectors corresponding to all eigenvalues.
(c) Find an invertible matrix $\boldsymbol{S}$ such that $\boldsymbol{S}^{-1}\boldsymbol{A}\boldsymbol{S}$ is upper triangular.
Hint: Find an appropriate basis for each generalized eigenspace.

2.10. The goal of this exercise is to show that two diagonalizable endomorphisms F and G on a vector space[13] V are simultaneously diagonalizable—i.e., possess a common basis of eigenvectors—iff they commute—i.e., $FG = GF$.
(a) Assume that F and G have a a common basis of eigenvectors. Show that they commute.
(b) Now assume that F and G commute.
(i) Show that every eigenspace of F is G-invariant.
(ii) Let $(\lambda_1, \dots, \lambda_k)$ be the pairwise distinct eigenvalues of F. Let v be an eigenvector of G with eigenvalue μ. Let $v = v_1 + \cdots + v_k$ be the unique decomposition of v with $v_i \in \mathrm{Eig}(F, \lambda_i)$. Show that $Gv_i = \mu v_i$ for every i.
Hint: Use point (a) and the uniqueness of the decomposition.
(iii) Show that $\mathrm{Eig}(G,\mu) = \bigoplus_{i=1}^k (\mathrm{Eig}(F,\lambda_i) \cap \mathrm{Eig}(G,\mu))$.[14]
Hint: Use point (b).

[13]For notational simplicity, we assume V to be finite-dimensional.
[14]Some of these intersections might be the zero space.

(iv) Conclude that

$$V = \bigoplus_{\substack{i=1,\dots,k \\ j=1,\dots,l}} \operatorname{Eig}(F, \lambda_i) \cap \operatorname{Eig}(G, \mu_j),$$

where $(\mu_1, \dots, \mu_l)$ are the pairwise distinct eigenvalues of G.

(v) Conclude that F and G have a common basis of eigenvectors.

Index